

# Elastic MapReduce EMR on CVM 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

#### STencent Cloud

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

#### EMR on CVM Operation Guide

Planning Cluster

**Business Evaluation** 

**Cluster Types** 

Cross-AZ Cluster

Cross-AZ Cluster Deployment

Cross-AZ High Availability

**Configuring Cluster** 

**Creating Cluster** 

Administrative rights

CAM Overview

**Role Authorization** 

Collaborator/Sub-user Permissions

CAM-Enabled EMR API Authorization Granularity Details

Authentication Granularity Scheme

Cluster COS Service Role

Setting Tags

Bootstrap Actions

Software Configuration

Mounting CHDFS Instance

Unified Management of Hive Metadata

Setting Security Groups

**Component Configuration Sharing** 

Managing Cluster

Instance Information

Node Specification Management

Checking and Updating Public IP

Logging in to Clusters

**Cluster Scale-Out** 

**Cluster Scale-in** 

**Releasing Now** 

Auto Scaling

Overview

Custom Scaling Configuration

Managed Scaling Configuration

Viewing Scaling Records

Graceful Scale-In

**Changing Configurations** 

Automatic Replacement

Exporting Software Configuration

**Cluster Scripts** 

**Cluster Termination** 

**Operation Logs** 

Task Center

**Disk Management** 

Scaling up Cloud Disks

Disk Update Check

**Repairing Disks** 

Managing Service

Managing Users

Adding Components

**Restarting Service** 

Starting/Stopping Services

WebUI Access

Role Management

**Client Management** 

Configuration Management

**Configuration Management** 

**Configuration Status** 

**Configuration Rollback** 

**Configuration Group Management** 

Service List

YARN Resource Scheduling

Overview

Configuring Fair Scheduler

Configuring Capacity Scheduler

Label Management

Viewing Scheduling History

HBase RIT Fixing

**Component Port Information** 

RSS Uniffle Operation Guide

Monitoring and Alarms

**Cluster Overview** 

Node Status Service Status **Cluster Event** Log **Application Analysis** StarRocks Query Management Hive Data Table Analysis YARN Job Query HDFS File Storage Analysis Impala Query Management **HBase Table Analysis** Kudu Table Analysis **Cluster Inspection** Java Analysis **Monitoring Metrics** Node Monitoring Metrics **HDFS Monitoring Metrics** YARN Monitoring Metrics ZooKeeper Monitoring Metrics **HBase Monitoring Metrics Hive Monitoring Metrics** Spark Monitoring Metrics **Presto Monitoring Metrics Trino Monitoring Metrics ClickHouse Monitoring Metrics Druid Monitoring Metrics** Kudu Monitoring Metrics Alluxio Monitoring Metrics PrestoSQL Monitoring Metrics Impala Monitoring Metrics **Ranger Monitoring Metrics COSRanger Monitoring Metrics Doris Monitoring Metrics** Kylin Monitoring Metrics **Zeppelin Monitoring Metrics Oozie Monitoring Metrics** Storm Monitoring Metrics Livy Monitoring Metrics

Kyuubi Monitoring Metrics StarRocks Monitoring Metrics Kafka Monitoring Metrics Alarm Configurations Alarm Records

## EMR on CVM Operation Guide Planning Cluster Business Evaluation

Last updated : 2025-01-20 17:34:28

### Selecting a Cluster Type

Elastic MapReduce (EMR) provides six types of clusters for you to choose from based on your business needs. Hadoop cluster: Based on open-source Hadoop and the components that form a Hadoop ecosystem, it offers five use cases, including the default use case, ZooKeeper, HBase, Presto, and Kudu, and provides big data solutions for massive data storage, offline/real-time data analysis, streaming data compute, and machine learning. Druid cluster: Druid is a high-performance real-time analytics database. It supports big data queries in milliseconds and multiple data ingestion methods. It is suitable for real-time big data query scenarios.

ClickHouse cluster: ClickHouse is a column-oriented database management system. It is suitable for data warehouse analysis scenarios such as real-time wide table analysis, real-time BI report analysis, and user behavior analysis. Doris cluster: Doris is an MPP analytical database product that supports sub-second queries on PB-level, structured data. It is compatible with MySQL protocol and uses the standard SQL syntax. It is suitable for historical report analysis, real-time data analysis, real-time data analysis, interactive data analysis, etc.

Kafka cluster: Kafka is a distributed, partitioned, multi-replica, and multi-subscriber message processing system based on ZooKeeper coordination. It is suitable for asynchronous processing, message communication, and streaming data receiving and distribution.

StarRocks cluster: StarRocks adopts full vectorization technology. It supports extremely fast and unified OLAP databases. It is suitable for many data analysis scenarios, such as multidimensional, real-time, and high-concurrency analysis.

### Selecting Billing Mode

Billing mode for EMR clusters:

**Pay-as-you-go:** All nodes in a cluster are charged on a pay-as-you-go basis. This is suitable for clusters that exist for a short time or periodically.

#### Note:

Select the shutdown mode with caution when shutting down a pay-as-you-go EMR cluster node in the CVM console, because EMR nodes do not support the "no charges when shut down" mode.

### Selecting Model and Specification

EMR offers a wide variety of CVM models, including EMR Standard, EMR Compute, EMR High IO, EMR MEM Optimized, and EMR Big Data. If you need the CPM model, submit a ticket to us.

You can choose a model based on your business needs and budget.

If you require low latency for offline compute, we recommend you select a model with local disks or the Big Data model.

If you need to use the real-time database HBase, we recommend you select the EMR High IO model with local SSD disks for optimal performance.

#### Node specification recommendations

EMR offers five types of nodes for your choice based on the cluster type.

Cluster Type	Use Case	Node Type	Recommended Specification
Hadoop	Default use case	Master	Master node: We recommend you select an instance specification with a large memory size (at least 16 GB) and use cloud disks for high stability.
		Core	If most of your data is stored on COS, core nodes will function in a way similar to task nodes and should have a capacity of at least 500 GB. Core nodes cannot be elastically scaled. If your architecture does not use COS, core nodes are responsible for processing cluster compute and storage tasks, and three-replica backup is enabled by default. When estimating the data disk capacity, you need to consider the capacity for storing three replicas. In this case, the Big Data model is recommended.
		Task	If your architecture does not use COS, task nodes are not required. If most of your data is stored on COS, task nodes can be used as elastic compute resources and deployed as needed.
		Common	Common node: It is mainly used as ZooKeeper nodes. You need to select a specification of at least 2 cores, 4 GB memory, and 100 GB cloud disk capacity to meet the requirements.
		Router	Router node: It is mainly used to relieve the load of master nodes and as a task submitter. Therefore, we



		recommend you select a model with a large memory size, preferably not lower than the specification of master nodes.
ZooKeeper	Common	Common node: It is mainly used as ZooKeeper nodes. You need to select a specification of at least 2 cores, 4 GB memory, and 100 GB cloud disk capacity to meet the requirements.
	Master	Master node: We recommend you select an instance specification with a large memory size (at least 16 GB) and use cloud disks for high stability.
	Core	If most of your data is stored on COS, core nodes will function in a way similar to task nodes and should have a capacity of at least 500 GB. Note: Core nodes cannot be elastically scaled. If your architecture does not use COS, core nodes are responsible for processing cluster compute and storage tasks.
HBase	Task	If your architecture does not use COS, task nodes are not required. If most of your data is stored on COS, task nodes can be used as elastic compute resources and deployed as needed.
	Common	Common node: It is mainly used as ZooKeeper nodes. You need to select a specification of at least 2 cores, 4 GB memory, and 100 GB cloud disk capacity to meet the requirements.
	Router	Router node: It is mainly used to relieve the load of master nodes and as a task submitter. Therefore, we recommend you select a model with a large memory size, preferably not lower than the specification of master nodes.
Kudu	Master	Master node: We recommend you select an instance specification with a large memory size (at least 16 GB) and use cloud disks for high stability.
	Core	If most of your data is stored on COS, core nodes will function in a way similar to task nodes and should have a capacity of at least 500 GB. Note: Core nodes cannot be elastically scaled.

		If your architecture does not use COS, core nodes are responsible for processing cluster compute and storage tasks, and three-replica backup is enabled by default. When estimating the data disk capacity, you need to consider the capacity for storing three replicas. In this case, the Big Data model is recommended.
	Task	If your architecture does not use COS, task nodes are not required. If most of your data is stored on COS, task nodes can be used as elastic compute resources and deployed as needed.
	Common	Common node: It is mainly used as ZooKeeper nodes. You need to select a specification of at least 2 cores, 4 GB memory, and 100 GB cloud disk capacity to meet the requirements.
	Router	Router node: It is mainly used to relieve the load of master nodes and as a task submitter. Therefore, we recommend you select a model with a large memory size, preferably not lower than the specification of master nodes.
Presto	Master	Master node: We recommend you select an instance specification with a large memory size (at least 16 GB) and use cloud disks for high stability.
	Core	If most of your data is stored on COS, core nodes will function in a way similar to task nodes and should have a capacity of at least 500 GB. Note: Core nodes cannot be elastically scaled. If your architecture does not use COS, core nodes are responsible for processing cluster compute and storage tasks, and three-replica backup is enabled by default. When estimating the data disk capacity, you need to consider the capacity for storing three replicas. In this case, the Big Data model is recommended.
	Task	If your architecture does not use COS, task nodes are not required. If most of your data is stored on COS, task nodes can be used as elastic compute resources and deployed as needed.



		Common	Common node: It is mainly used as ZooKeeper nodes. You need to select a specification of at least 2 cores, 4 GB memory, and 100 GB cloud disk capacity to meet the requirements.
		Router	Router node: It is mainly used to relieve the load of master nodes and as a task submitter. Therefore, we recommend you select a model with a large memory size, preferably not lower than the specification of master nodes.
ClickHouse	Default use case	Core	Core node: We recommend you select a model with high CPU and a large memory size. Because data may be lost if a local disk is corrupted, cloud disks are recommended.
		Common	Common node: The CPU and memory configuration should be at least 4 cores and 16 GB.
Kafka	Default use case	Core	Core node: We recommend you select a model with high CPU and a large memory size. Because data may be lost if a local disk is corrupted, cloud disks are recommended.
		Common	Common node: The CPU and memory configuration should be at least 4 cores and 16 GB.
Doris	Default use case	Master	Master node: We recommend you select an instance specification with a large memory size (at least 16 GB) and store all the metadata of master nodes in the memory.
		Core	Core node: We recommend you select an instance specification with a large memory size (at least 8 GB) and use cloud SSD for better IO performance and stability.
		Router	Router node: The frontend module is deployed here for high read/write availability. Therefore, we recommend you select a model with a large memory size, preferably not less than that of master nodes.
Druid	Default use case	Master	Master node: We recommend you select an instance specification with a large memory size (at least 16 GB) and use SSD disks for better IO performance.
		Core	Core node: We recommend you select an instance



			specification with a large memory size (at least 8 GB) and use cloud SSD for better IO performance and stability.
		Task	If your architecture does not use COS, task nodes are not required. If most of your data is stored on COS, task nodes can be used as elastic compute resources and deployed as needed.
		Common	Common node: It is mainly used as ZooKeeper nodes. We recommend you select the specification of 2 cores, 4 GB memory, and 100 GB cloud disk capacity to meet the requirements.
		Router	Router node: It is mainly used to relieve the load of master nodes and as a task submitter. Therefore, we recommend you select a model with a large memory size, preferably not lower than the specification of master nodes.
StarRocks	Default use case	Master	Master node: We recommend you select an instance specification with a large memory size (at least 16 GB) and store all the metadata of master nodes in the memory.
		Core	Core node: We recommend you select an instance specification with a large memory size (at least 8 GB) and use cloud SSD for better IO performance and stability.
		Router	Router node: The frontend module is deployed here for high read/write availability. Therefore, we recommend you select a model with a large memory size, preferably not less than that of master nodes.

#### Note:

Different cluster types have different requirements for the node specification. Currently, the system automatically recommends the configuration that meets the cluster's requirements by default. You can adjust the model specification based on your business needs, and the recommended model is for reference only.

Core nodes cannot be elastically scaled. If your architecture does not use COS, core nodes are responsible for processing cluster compute and storage tasks, and three-replica backup is enabled by default. When estimating the data disk capacity, you need to consider the capacity for storing three replicas. In this case, the Big Data model is recommended.

### Network and Security

To ensure the network security, the EMR cluster is placed in a VPC, and a security group policy is added to the VPC. In addition, to ensure easy access to the WebUI of Hadoop, a public IP is enabled for one of the master nodes and the node is billed by traffic. A public IP is not enabled for router nodes by default. However, you can bind a router node to an EIP in the CVM console to enable a public IP for it.

#### Note:

A public IP is enabled for master nodes when a cluster is created. You can disable it as needed.

Enabling a public IP for master nodes is mainly for SSH login and component WebUI access.

Master nodes with a public IP enabled are billed by traffic with a bandwidth of up to 5 Mbps. You can adjust the network in the console after creating a cluster.

## **Cluster Types**

Last updated : 2023-12-27 09:59:01

EMR supports six cluster types and their respective use cases and defines five node types. Different cluster types and their respective use cases support different node types, number of deployed nodes, and deployed services. You can select the most appropriate cluster type and use case based on your business needs when creating a cluster.

#### Note

ClickHouse and Doris cluster types are not available by default. To use them, submit a ticket for application.

### **Cluster Type Description**

#### Hadoop cluster

Use Case	Description	Node Deployment
Default	Based on open-source Hadoop and the components that form a Hadoop ecosystem, it provides big data solutions for massive data storage, offline/real-time data analysis, streaming data computing, and machine learning.	Master node: It is a management node that ensures the scheduling of the cluster works properly. Processes such as NameNode, ResourceManager, and HMaster are deployed here. The number of master nodes is 1 in non-HA mode and 2 in HA mode.Note: If Kudu is deployed, the cluster supports only the HA mode, and there are 3 master nodes. Core node: It is a compute and storage node. All your data in HDFS is stored in core nodes. Therefore, core nodes cannot be scaled in once scaled out to ensure data security. Processes such as DataNode, NodeManager, and RegionServer are deployed here. The number of core nodes is ≥ 2 in non-HA mode and ≥ 3 in HA mode. Task node: It is a node for computing only and does not store any data. The computed data comes from a core nodes and can be scaled in or out as needed. Processes such as NodeManager and PrestoWork are deployed here. The number of task nodes can be changed at any time to scale the cluster, with a minimum value of 0. Common node: It provides data sharing and syncing and HA fault tolerance services for the master nodes in an HA cluster. Distributed coordinator components such as ZooKeeper and JournalNode are deployed

		<ul> <li>here. The number of common nodes is 0 in non-HA mode and ≥ 3 in HA mode.</li> <li>Router node: It is used to share the load of a master node or as the task submitter of the cluster. It can be scaled in or out at any time. Hadoop packages, including software programs and processes such as Hive, Hue, and Spark, are deployed here. The number of router nodes can be changed at any time, with a minimum value of 0.</li> </ul>
ZooKeeper	It is suitable for creating a distributed, high-availability coordination service for large clusters.	Common node: Distributed coordinator components such as ZooKeeper are deployed here. The number of deployed common nodes must be odd and be at least three. Common nodes support only the HA mode.
HBase	It is suitable for storing massive amounts of unstructured or semi- structured data. It provides a high- reliability, high-performance, column-oriented, scalable distributed storage system that supports real-time data read/write.	Master node: It is a management node that ensures the scheduling of the cluster works properly. Processes such as NameNode, ResourceManager, and HMaster are deployed here. The number of master nodes is 1 in non-HA mode and 2 in HA mode. Core node: It is a compute and storage node. All your data in HDFS is stored in core nodes. Therefore, core nodes cannot be scaled in once scaled out to ensure data security. Processes such as DataNode, NodeManager, and RegionServer are deployed here. The number of core nodes is $\geq 2$ in non-HA mode and $\geq 3$ in HA mode. Task node: It is a node for computing only and does not store any data. The computed data comes from a core node or COS. Therefore, task nodes are often elastic nodes and can be scaled in or out as needed. Processes such as NodeManager are deployed here. The number of task nodes can be changed at any time to scale the cluster, with a minimum value of 0. Common node: It provides data sharing and syncing and HA fault tolerance services for the master nodes in an HA cluster. Distributed coordinator components such as ZooKeeper and JournalNode are deployed here. The number of common nodes is 0 in non-HA mode and $\geq 3$ in HA mode. Router node: It is used to share the load of a master node or as the task submitter of the cluster. It can be scaled in or out at any time. The number of router nodes can be changed at any time of o.

Presto	It provides an open-source distributed SQL query engine for quick query and analysis of massive amounts of data. It is suitable for interactive analytical queries.	Master node: It is a management node that ensures the scheduling of the cluster works properly. Processes such as NameNode and ResourceManager are deployed here. The number of master nodes is 1 in non-HA mode and 2 in HA mode. Core node: It is a compute and storage node. All your data in HDFS is stored in core nodes. Therefore, core nodes cannot be scaled in once scaled out to ensure data security. Processes such as DataNode and NodeManager are deployed here. The number of core nodes is $\geq 2$ in non-HA mode and $\geq 3$ in HA mode. Task node: It is a node for computing only and does not store any data. The computed data comes from a core node or COS. Therefore, task nodes are often elastic nodes and can be scaled in or out as needed. Processes such as NodeManager and PrestoWork are deployed here. The number of task nodes can be changed at any time to scale the cluster, with a minimum value of 0. Common node: It provides data sharing and syncing and HA fault tolerance services for the master nodes in an HA cluster. Distributed coordinator components such as ZooKeeper and JournalNode are deployed here. The number of common nodes is 0 in non-HA mode and $\geq 3$ in HA mode. Router node: It is used to share the load of a master node or as the task submitter of the cluster. It can be scaled in or out at any time. The number of router nodes can be changed at any time. The number of router nodes can be changed at any time. The number of router nodes can be changed at any time. The number of router nodes can be changed at any time. The number of router nodes can be changed at any time, with a minimum value of 0.
Kudu	It provides a distributed and scalable columnar storage manager and supports random reads/writes and OLAP analysis to process frequently updated data.	Master node: It is a management node that ensures the scheduling of the cluster works properly. Processes such as NameNode and ResourceManager are deployed here. The number of master nodes is 1 in non-HA mode and 2 in HA mode. Core node: It is a compute and storage node. All your data in HDFS is stored in core nodes. Therefore, core nodes cannot be scaled in once scaled out to ensure data security. The number of core nodes is $\geq 2$ in non-HA mode and $\geq 3$ in HA mode. Task node: It is a node for computing only and does not store any data. The computed data comes from a core nodes and can be scaled in or out as needed. The



	number of task nodes can be changed at any time to scale the cluster, with a minimum value of 0. Common node: It provides data sharing and syncing and HA fault tolerance services for the master nodes in an HA cluster. Distributed coordinator components such as ZooKeeper and JournalNode are deployed here. The number of common nodes is 0 in non-HA mode and $\geq$ 3 in HA mode. Router node: It is used to share the load of a master node or as the task submitter of the cluster. It can be scaled in or out at any time. The number of router nodes can be changed at any time, with a minimum value of 0.
--	---

#### **Druid cluster**

Use Case	Description	Node Deployment
Default	It supports high-performance real- time analysis, big data queries in milliseconds, and multiple data ingestion methods. It is suitable for real-time big data query scenarios.	Master node: It is a management node that ensures the scheduling of the cluster works properly. Processes such as NameNode and ResourceManager are deployed here. The number of master nodes is 1 in non-HA mode and 2 in HA mode. Core node: It is a compute and storage node. All your data in HDFS is stored in core nodes. Therefore, core nodes cannot be scaled in once scaled out to ensure data security. Processes such as DataNode and NodeManager are deployed here. The number of core nodes is $\geq 2$ in non-HA mode and $\geq 3$ in HA mode. Task node: It is a node for computing only and does not store any data. The computed data comes from a core node or COS. Therefore, task nodes are often elastic nodes and can be scaled in or out as needed. Processes such as NodeManager are deployed here. The number of task nodes can be changed at any time to scale the cluster, with a minimum value of 0. Common node: It provides data sharing and syncing and HA fault tolerance services for the master nodes in an HA cluster. Distributed coordinator components such as ZooKeeper and JournalNode are deployed here. The number of z in HA mode. Router node: It is used to share the load of a master node or as the task submitter of the cluster. It can be scaled in or out

	at any time. The number of router nodes can be changed at any time, with a minimum value of 0.
--	--

#### ClickHouse cluster

Use Case	Description	Node Deployment
Default	It provides a column-oriented database management system. It is suitable for data warehouse analysis scenarios such as real-time wide table analysis, real-time BI report analysis, and user behavior analysis.	Core node: It is a compute and storage node. ClickHouseServer is deployed here. Common node: It provides data sharing and syncing and HA fault tolerance services for the master nodes in an HA cluster. Distributed coordinator components such as ZooKeeper are deployed here. The number of common nodes is 0 in non-HA mode and ≥ 3 in HA mode.

#### **Doris cluster**

Use Case	Description	Node Deployment
Default	It is an MPP analytical database product that supports sub-second queries on PB-scale, structured data. It is compatible with MySQL protocol and uses the standard SQL syntax. It is suitable for historical report analysis, real-time data analysis, interactive data analysis, etc.	Master node: It is a frontend module that provides the WebUI feature. Processes such as FE Follower and Broker are deployed here. The number of master nodes is $\geq$ 1 in non-HA mode and $\geq$ 3 in HA mode. Core node: It is a backend module that provides the data storage feature. Processes such as BE and Broker are deployed here. The number of core nodes is $\geq$ 3. Router node: It is a frontend module that helps achieve high read/write availability. Processes such as FE Observer and Broker are deployed here. Router nodes can be scaled out but not scaled in.

#### Kafka cluster

Use Case	Description	Node Deployment
Default	It is a distributed, partitioned, multi-replica,	Core node: It is a backend module for data

System based on Zookeeper coordination.       Broker         It is suitable for asynchronous processing,       core no         message communication, and streaming       in HA m         data receiving and distribution.       Common         and syr       services         cluster.       0 in nor	odes is $\geq$ 1 in non-HA mode or $\geq$ 2 mode. on node: It provides data sharing ncing and HA fault tolerance es for the core nodes in an HA . The number of common nodes is n-HA mode or $\geq$ 3 in HA mode.
--	---

#### StarRocks cluster

Use Case	Description	Node Deployment
Default	StarRocks adopts full vectorization technology. It supports extremely fast and unified OLAP databases. It is suitable for many data analysis scenarios, such as multidimensional, real-time, and high-concurrency analysis.	Master node: It is a frontend module that provides the WebUI feature. Processes such as FE Follower and Broker are deployed here. The number of master nodes is $\geq$ 1 in non-HA mode and $\geq$ 3 in HA mode. Core node: It is a backend module that provides the data storage feature. Processes such as BE and Broker are deployed here. The number of core nodes is $\geq$ 3. Task node: It is a node for computing only and does not store any data. The computed data comes from a core node or COS. Therefore, task nodes are often elastic nodes and can be scaled in or out as needed. Processes such as Compute Node are deployed here. The number of task nodes can be changed at any time to scale the cluster, with a minimum value of 0. Router node: It is a frontend module that helps achieve high read/write availability. Processes such as FE Observer and Broker are deployed here. Router nodes can be scaled out but not scaled in.

## Cross-AZ Cluster Cross-AZ Cluster Deployment

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

### Multi-AZ Cluster Deployment

EMR by default recommends deploying in a single AZ. However, a single AZ may face risks such as complete AZ failure or insufficient resources. To enhance AZ-level reliability and resource flexibility, in Hadoop cluster scenes, you can configure a multi-AZ deployment policy. Based on deployment characteristics, the policies are divided into split policy and balanced policy.

Split policy: This policy enhances resource scalability at the AZ level by deploying services with resources from two AZs when a cluster is created. You need to configure two AZs for cluster deployment: one primary AZ and one secondary AZ. If the secondary AZ encounters an exception, the primary AZ can continue to provide services normally.



Balanced policy: By deploying across mutually backed-up AZs, the cluster can continue to operate normally even if one AZ fails. Since cross-AZ data distribution capabilities are not yet supported, AZ-level disaster recovery is currently implemented using the EMR on COS deployment scheme. This requires configuring the cluster across three AZs: one

### 🔗 Tencent Cloud

primary AZ, one secondary AZ, and one balanced AZ. The primary and secondary AZs have mutually backed-up resources, so a failure in either the primary or secondary AZ does not affect the operation of the other.



#### Note

Cross-AZ network incurs usage fees. Since December 2022, this has been a paid service. For more details, see VPC Purchase Guide.

#### **Split Policy**

The following document explains how to create and manage clusters based on the split policy.

#### **Creating Clusters**

Log in to the EMR Console, and on the EMR on CVM cluster list page, click Create Cluster.

1. Software configuration step: **Select Hadoop Cluster Type**, and select the required cluster scene. For other information, see the single cluster purchase page, and see the Creating Cluster.

2. Region and hardware configuration step: In the cross-AZ information item, select the Cross-AZ option. In the

deployment policy item, select Split Policy. For the AZ information item, choose two AZs as needed, where the first row is the primary AZ, and the second row is the secondary AZ. Select subnets according to the chosen

#### AZs.

For node configuration items, set the node specifications and quantities for each selected AZ. In the split policy, the high availability option is enabled by default and cannot be disabled.

3. For basic configuration and to confirm the configuration details, see the single-cluster purchase page. For more information, see Creating Cluster.

#### **Cluster Scale-out**

In the split policy, both the primary and secondary AZ can support on-demand scale-out of Task, Core, and Router node types.

#### **Cluster Scale-in**

When you perform scale-in under the split policy, the total number of Core node types in both the **primary and secondary AZs can be reduced** to a minimum of 3, with no restrictions on other node types.

#### **Balanced Policy**

The following document explains how to create and manage clusters based on the balanced policy.

#### **Creating Clusters**

Log in to the EMR Console, and on the EMR on CVM cluster list page, click Create Cluster.

1. Software configuration step: **Select Hadoop Cluster Type**, and select the required cluster scene. For other information, see the single cluster purchase page, and see the Creating Cluster.

2. Region and hardware configuration step: Select the Cross-AZ option for the cross-AZ information item, and **choose the balanced policy for the deployment policy item.** You will need to select **three AZs as per your requirements. The first row represents the primary AZ, the second row the secondary AZ, and the third row the balanced AZ**. Choose the subnets accordingly based on the selected AZs. For node configuration items, set the node specifications and quantities for each selected AZ. In the balanced policy, the high availability option is **enabled by default and cannot be disabled.** 

3. For basic configuration and to confirm the configuration details, see the single-cluster purchase page. For more information, see Creating Cluster.

#### **Viewing Cluster**

After the cluster is successfully created, you can view the primary and secondary AZ information in the cluster list and on the cluster instance information page.

Log in to the EMR Console to access the cluster list page. Click the cluster **ID/Name** to enter the cluster instance information page:

#### **Cluster Scale-out**

The balanced policy allows for on-demand scale-out of Task, Core, and Router type nodes in the primary AZ, secondary AZ, and balanced AZ.

The balanced policy scale-out is solely for resource expansion. If the user's cluster services are deployed with a primary/replica scheme, it is recommended to scale out resources equally between the primary and secondary AZs.

#### **Cluster Scale-in**

When you perform scale-in under the split policy, the total number of Core node types in both the primary and secondary AZs can be reduced to a minimum of 3, with no restrictions on other node types.

In the balanced AZ, full-scale scale-in is supported for Core, Task, and Router node types during the scale-in process.



## Cross-AZ High Availability

Last updated : 2024-10-30 10:34:18

### **Overview of Rack Awareness**

Rack awareness in a Hadoop cluster refers to the technique where Hadoop organizes nodes according to the network topology and prioritizes task scheduling and data storage between nodes within the same rack. This improves cluster performance and reliability.

It is supported by two components: HDFS and YARN. HDFS achieves high reliability and availability by distributing replicas of data blocks across different racks. YARN improves task execution efficiency and performance by assigning tasks to nodes or containers that are physically closer.

Since Hadoop cannot automatically detect the network topology of nodes, it provides the following methods to enable rack awareness:

Custom Java Class implements the DNSToSwitchMapping API method, and the class name is specified in the coresite.xml configuration file using the net.topology.node.switch.mapping.impl parameter.

Topology mapping is based on a script, and the net.topology.script.file.name parameter in the core-site.xml configuration file is used to specify the script file.

Below is an example of configuring rack awareness policies based on a script. The basic method involves mapping AZ subnets to rack information.

#### Note

The setup of rack awareness should be based on a cross-AZ deployment architecture for the cluster (for cluster creation, see Cross-AZ Cluster Deployment). It is not applicable to single-AZ clusters.

### Configuring Rack Awareness Policy Based on Scripts

1. Prepare a cross-AZ EMR cluster. Log in to the EMR Console, click the cluster **ID/Name** to enter the cluster details page, and under **Instance Information > Deployment Information**, confirm the VPC network information of the cluster and the subnets corresponding to different AZs.

Then, in **VPC > Subnet**, obtain the CIDR and AZ mapping information for each subnet.

#### Note

Both the VPC name and subnet name may be duplicated, so you need to further verify the information in Instance Information under Cluster Resources.

2. Prepare the rack awareness script RackAware.py based on the subnet CIDR and AZ mapping information. **Note:** 

This script uses the /usr/bin/python path with the Python 2 version as an example. Replace #CIDR# with the actual subnet CIDR in the script.

```
#!/usr/bin/python
import sys
import IPy
import re
DEFAULT_RACK="/default-rack"
cidrToRack = {
 ' #CIDR#' : 'rack-1',
 ' #CIDR#' : 'rack-2',
 ' #CIDR#' : 'rack-3'
 }
for name in sys.argv[1:]:
 rack = DEFAULT_RACK
 ips = re.findall(r'[0-9]+(?:\backslash .[0-9]+){3}', name)
 if len(name) > 0 and len(ips) > 0:
     ip = ips[0]
     for cidr in cidrToRack.keys():
       if ip in IPy.IP(cidr):
             rack = cidrToRack[cidr]
             break
 print "/{0}".format(rack)
```

3. In the Cluster Service > HDFS > Configuration Management, add the RackAware.py file and update the NameNode node's core-site.xml file with the configuration item

```
net.topology.script.file.name=/usr/local/service/hadoop/etc/hadoop/RackAware.py .
```

4. In the console, restart NameNode and ResourceManager.

### Viewing Cluster Rack Information

HDFS service: Log in to the NameNode, and execute the command hdfs dfsadmin -printTopology as the hadoop user, as shown below:

[hadoo Rack:	op@10 ~]\$ /rack-1	hdfs dfsadmin	-printTo	opology	
	100.10.01	:4001 (	)		
		1001 (	, )		
		:4001 (	2		
		:4001 (	)		
Rack:	/rack-3				
		:4001 (			
		• 4001 (		)	
		- 4001 (	/		
		:4001 (	)		

YARN service: You can log in to the WebUI to view the information:

									Nodes	of t	he cl	ust	er						Logged in as
• Cluster	Cluster Me	etrics																	
About Nodes Node Labels	Apps Submitte	d Pe	Apps ending	Apps Running	Apps Complete	Apps d Failed	Ap Kil	led	Containers Running	Memory Used	Memor Total	y I F	Vemory Reserved	VCores Used	VCores Total	VCores Reserve	Phys d U	ical Mem sed %	Physical V Used
<u>Applications</u> <u>NEW</u>	Cluster No	udes Me	trics	0		0	0	0		0 B	43.20 Gt	5 UB		0	24	0	20		0
NEW SAVING SUBMITTED	Active	Nodes		Decommiss	sioning No	des		Decom	missioned Node	s	Lost No	odes	Unh	ealthy Node	25	Rebooted I	lodes	Shu	tdown Node:
ACCEPTED RUNNING	<u>6</u>	110000	Q	00001111100	Johning Ho	400	0	booom			0		<u>0</u>	ountry riou	<u>0</u>	1100001001		<u>0</u>	ao minino ao m
FINISHED	Ci	ontainers	AssignedPe	rSec	pending container allocated container re						rese	erved con	tainer	R	MEventQueueS	ize	Sch	edulerEvent	QueueSize
KILLED	0				0			0			0			0			)		
Scheduler	User Metri	cs for h	nadoop																
→ Tools	Apps Submitte	d	Apps Pending	Apps Bunning	Ap	ops oleted	Contain Bunnir	iers na	Containers	Ci	ontainers Reserved	Me	mory sed	Memory Pending	Memo	y V ad U	Cores Ised	VCores	VCor
1	0	0	r ontaning	0	0	0	- turnin	.9	0	0	10001100	0 B	0	B	0 B	0	0	1 onung	0
	Scheduler Metrics																		
,	Sche	duler Typ	e	Schedu	uling Reso	Resource Type Minimum Allocation					Maximum Allocation Maximum Cluster Application					Application I	Priority		
	Fair Sched	uler	[M	EMORY, CPU	]	<memory:16, vcores:1=""> <me< td=""><td colspan="6">1emory:7373, vCores:4&gt; 0</td><td></td></me<></memory:16,>						1emory:7373, vCores:4> 0							
	Show 20	entries										Search:							
	Node Labels *	Rack ≎	Node State \$	Node Add	Iress 🗘	Node HT Address	TP \$	Last h	nealth-update \$	Health-	report ≎	Contain	ers Men	n Mem ≎ Avail ≎	Phys Mem Used %	<ul><li>VCore</li><li>↓ Used</li></ul>	<ul><li>VCore</li><li>Avail</li></ul>	s Phys \$ Use	VCores ∖ d%≎
		/rack- 3	RUNNING	1.75.41	I:5006	1.01.4.14	:5008	611	23 ) 2023			0	0 B	7.20 GB	20.940504	0	4	0.9996	3675 2
		/rack- 1	RUNNING		5006	1.01.1.00	<u>5008</u>	21:41:0	23 0 +0800 2023			0	0 B	7.20 GB	20.901533	0	4	5.5	2
		/rack- 1	RUNNING	1.26.1.3	5006	1.00.0.00	<u>5008</u>	21:41:0	23 0 +0800 2023			0	0 B	7.20 GB	20.200052	0	4	5.08163	193 2
		/rack- 1	RUNNING		5006	1.01.1.01	<u>5008</u>	21:41:0	] 23 0 +0800 2023			0	0 B	7.20 GB	20.914524	0	4	5.5	2
		/rack- 3	RUNNING	1.22.41	3:5006	1.01.6.10	:5008	21:41:0	23 0 +0800 2023			0	0 B	7.20 GB	20.667706	0	4	5.16494	45 2
		/rack- 3	RUNNING	1.75.41	:5006	12142	:5008	21:41:0	23 0 +0800 2023			0	0 B	7.20 GB	20.888542	0	4	5.16494	45 2
	Showing 1	:o 6 of 6	entries														F	irst Previo	us 1 Next

## Configuring Cluster Creating Cluster

Last updated : 2024-10-30 10:32:07

### Overview

This document introduces the directions and configurations for creating an EMR on CVM cluster through the EMR console.

### Directions

Log in to the EMR console, click **Create cluster** on the EMR on CVM cluster list page. When creating a cluster, you need to go through four steps: Software configuration, AZ and hardware configuration, basic configuration, and Confirmation configuration.

#### Software Configuration

Configuration Item	Configuration Description
Region	The physical data center where the cluster is deployed. Each region represents an independent physical data center, and private networks of CVMs between different regions are not interconnected. Note: Once the cluster is created, the region cannot be changed, so choose carefully.
Cluster type	EMR on CVM supports multiple cluster types. Select the appropriate cluster type based on your business needs. The default is the Hadoop cluster type. For an introduction to each cluster type, see Cluster Type.
Use cases	The Hadoop cluster type supports five use cases: Default Scene, Zookeeper, HBase, Presto, and Kudu. Select the appropriate use case for deployment based on your business needs.
Product version	The components and their versions bundled with different product versions vary.
Components to deploy	Optional components that can be customized and combined based on your needs.
Kerberos mode	Disabled by default. When it is enabled, the open-source components in the cluster will be started in Kerberos security mode. For more details, see Kerberos Overview.



Component dependency	Disabled by default. When it is enabled, components from an existing cluster are shared for use in the current cluster. For more information, see Component Configuration Sharing. Note: Since the selected components are deployed on an existing cluster, the cluster providing the dependency components cannot be directly terminated. You should first terminate all clusters that rely on its components.
Software configuration	Optional configuration. Before starting the cluster, you can specify a JSON file to modify component configuration parameters or access external clusters. For more details, see Software Configuration.

### AZ and Hardware Configuration

Configuration Item	Configuration Description
Billing mode	The billing mode supports monthly billing as well as pay-as-you-go billing. Monthly subscription: Prepay for N months of product fees, which is more cost-effective compared to pay-as-you-go pricing. Pay-as-you-go: Pay based on usage duration. Account identity verification is required, and a 2-hour fee will be frozen at the time of activation (vouchers cannot be used as a freezing deposit). The frozen resource fee will be refunded upon termination.
(Cross) AZ	<ol> <li>You can choose whether to deploy across regions as needed. Cross-AZ deployment distributes service and management roles across multiple AZs, providing varying levels of high availability. By default, a single AZ is used, and cross-AZ deployment is available through allowlist access.</li> <li>Different AZs within the same region may support different model specifications; it is recommended to choose the latest AZ. Cloud products in different regions cannot communicate through private networks and cannot be changed after purchase. It is recommended to choose a region and AZ close to your business data to reduce access delay and improve download speed.</li> </ol>
Deployment policy	<ul> <li>When Cross-AZ deployment is selected, both the equalization policy and balancing policy are supported. Equalization policy: A deployment scheme using Primary AZ + Secondary AZ. If the secondary AZ encounters an issue, the primary AZ can continue to provide services normally.</li> <li>Balancing policy: A deployment scheme using Primary AZ + Secondary AZ + Balanced AZ. The primary and secondary AZs are mutually redundant. If any AZ has an issue, the other AZs are not affected and continue to provide services.</li> </ul>
Cluster network	To ensure the security of the EMR cluster, all cluster nodes will be placed in a VPC. You need to set up a VPC to ensure the correct creation of the EMR cluster.
Cluster public network	It can be used for SSH log-in and component webui access from the public network, and the network can be adjusted in the console after the cluster is successfully created. The public network is enabled by default for the master1 node.



Security group	The security group functions as a firewall to configure network access control for CVM. If no security group is available, EMR will automatically create one for you. If there is an existing security group in use, you can select it directly. If the number of security groups has reached the upper limit and new ones cannot be created, you can delete some unused security groups. View the security groups currently in use. Create a security group: EMR helps users create a security group, enabling ports 22 and 30001, as well as the necessary private network IP ranges. Existing EMR security group: Select an already created EMR security group as the security group for the current instance, enabling ports 22 and 30001, as well as the necessary private network IP ranges.
Remote log-in	Port 22 is commonly used for remote login. The created security groups will have this port enabled by default. You can close this port based on your business needs, but it is enabled by default.
High availability (HA)	HA is enabled by default. Different cluster types and use cases have varying numbers of nodes deployed in HA or non-HA environments. For more details, see Cluster Type.
Node type	<ul> <li>Select the appropriate model configuration for different node types based on business requirements. For more details, see Business Assessment.</li> <li>Note: <ol> <li>Currently, Core nodes, Task nodes, and Router nodes support mounting multiple types of cloud disks (each type can only be selected once) and multiple disks (up to 20).</li> <li>Local disk models are not supported for deployment on Master and Common nodes.</li> <li>Select a non-local disk model.</li> </ol> </li> </ul>
Placement group	Optional configuration. A placement group is a policy for distributing CVM instances across underlying hardware. For more details, see Placement Group.
Hive metadatabase	If the Hive component is selected, Hive Metastore offers two storage options: The first is the default cluster option, where Hive metadata is stored in a separately purchased MetaDB for the cluster. The second option is to associate an external Hive Metastore, where you can choose to link to EMR-MetaDB or a self-built MySQL database, with metadata stored in the associated database, which will not be destroyed when the cluster is terminated. For more details, see Hive Metadata Management. Note: When you select one or more of the following components such as Hue, Ranger, Oozie, Druid, and Superset, the system will automatically purchase a MetaDB for storing metadata of components other than Hive.

### **Basic Configuration**

Configuration Item	
-----------------------	--



Project	Assign the current cluster to different project groups. Note: The project cannot be modified once the cluster is created.
Cluster name	Set a cluster name to distinguish between different EMR clusters. The system generates a name randomly, which can be modified. The cluster name needs to be between 6 and 36 characters and can only contain Chinese characters, letters, digits, hyphens (-), and underscores (_).
Login method	Currently, EMR provides two methods for logging in to cluster services, nodes, and MetaDB: Custom password setup and key association. SSH keys are used only for quick access through the EMR-UI. The default username is root, while the username for the Superset component's WebUI quick access is admin.
Bootstrap action	Optional configuration. Bootstrap scripts allow you to run custom scripts during cluster creation, enabling you to modify the cluster environment, install third-party software, and use your own data. For more settings, see Bootstrap Actions.
Tag	Optional configuration. You can add tags to cluster or node resources during creation to facilitate resource management. A maximum of 5 tags can be added, and tag keys should not be duplicated.

#### **Configuration confirmation**

Configuration Item	Configuration Description
Configuration list	Confirm if there is any error in the deployment information.
Auto-renewal	Automatic renewal Optional. Seven days before the cluster expires, the system will check daily if the user account has sufficient available balance to renew cluster resources set for automatic renewal. Monthly subscription clusters are automatically set for renewal by default, but users can manually deselect this option.
Agreement	Agree to the EMR Service Level Agreement and Refund Agreement.

After the above configurations are completed, click Purchase to proceed with payment. Once the payment is successful, the EMR cluster will start the creation process. In approximately 10 minutes, you can find the created cluster in the EMR console.

#### Note

Pay-as-you-go cluster: Creation starts immediately. Once the cluster is created, its status changes to Running. Monthly subscription cluster: An order is generated first, and the cluster creation starts after the payment is completed. You can view the instance information for each node in the CVM console. To ensure the proper functioning of the EMR cluster, do not modify the configuration of these instances in the CVM console.

### Next Steps

After the cluster is successfully created, you can log in and further configure the cluster as needed. See the following documentation for detailed operations:

Configure the cluster: Software Configuration, Mounting CHDFS, and Unified Management of Hive Metadata. Managing the cluster: Set Tags, Setting Bootstrap Actions, and Cluster Termination.

## Administrative rights CAM Overview

Last updated : 2023-12-27 09:59:20

### CAM Overview

When using Tencent Cloud EMR, different departments and roles need different permissions in order to avoid security risks such as leakages and maloperations. To this end, you can assign different permissions to different users through sub-accounts. By default, a sub-account does not have the permission to use EMR or related resources. Therefore, you need to create a policy to grant the required permission to the sub-account first.

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

When using EMR, you can associate a policy with a user or user group to allow or forbid them to use specified resources to complete specified tasks. For more information on CAM policies, see Element Reference. For more information on how to use CAM policies, see Policy.

### CAM Use Cases

#### 1. Resource management permissions need to be granted to sub-users

You can create users or roles in CAM and assign them separate security credentials (console login passwords, TencentCloud API keys, etc.) or request temporary credentials for them to access Tencent Cloud resources. You can also manage permissions to control the operations users and roles can perform and the resources they can access.

#### 2. Users with external Tencent Cloud roles can access Tencent Cloud resources

You can use your existing authentication system through CAM to grant your employees and services the access permissions for Tencent Cloud services and resources.

#### 3. Strengthen your security with an additional layer of protection

Currently, three authentication methods are supported: WeChat QR code, hardware/virtual MFA, and mobile verification code.

### Supported CAM Granularities

EMR supports **resource-level authentication** and **API-level authentication**. For API-level authentication, API input parameters include the cluster string ID, and the TencentCloud API directly forwards the authentication request to CAM for authentication. Resource-level authentication involves fine-granularity authentication on the EMR backend. EMR supports three authorization methods: resource-level authorization, API-level authorization, and authorization by tag.

**Resource-level and API-level authorization**: You can use policy syntax to grant sub-accounts permissions to manage individual resources. For more information, see Authentication Granularity Scheme.

**Authorization by tag**: You can tag resources and grant sub-accounts permissions to manage resources with particular tags. For more information, see Authentication Granularity Scheme.

## **Role Authorization**

Last updated : 2024-10-30 15:51:14

When using the EMR service, users need to grant the service account the default system role EMR\_QCSRole. Once the role is successfully granted, EMR can call related services (such as TKE and COS) to create clusters and save logs.

#### Note

When enabling EMR for the first time, you need to complete the role authorization process using the root account; otherwise, neither sub-accounts nor the root account can use EMR.

### **Role Authorization Process**

1. When a user creates a cluster or creates an on-demand execution plan, if the EMR\_QCSRole role authorization for the service account fails, the user will be redirected to a page notifying the permission limitations. Then click **Go to CAM** to proceed with role authorization.

2. Click **Grant** to authorize the default role EMR\_QCSRole to the EMR service account.

3. After authorization is completed, the user needs to refresh the EMR console or purchase page, after which normal operations can proceed. For more detailed information on EMR\_QCSRole policies, you can log in to the CAM console. The permissions included in EMR\_QCSRole can be found in Collaborator/Sub-account Permissions.

# Special Instructions for Service Role Authorization Related to EMR on TKE Clusters

When you create or use an EMR on TKE cluster, data needs to be directly written to or calculated in Cloud Object Storage (COS). To ensure data security, EMR should be granted temporary keys to read and write COS resources. Therefore, the relevant EMR service-related role EMR\_QCSLinkedRoleInApplicationDataAccess should be authorized and bound to the QcloudAccessForEMRLinkedRoleInApplicationDataAccess preset policy. 1. When viewing the EMR on TKE cluster list, you need to check if the service-related role

EMR\_QCSLinkedRoleInApplicationDataAccess is bound to the EMR service.

2. If the EMR service-related role EMR\_QCSLinkedRoleInApplicationDataAccess does not exist, authorization and binding need to be performed.

#### Note

If you need to specify cluster access permissions for the corresponding COS resources in a more refined manner, see Custom Service Roles for settings.

### EMR on TKE Cluster Authentication Description

The permission settings for sub-accounts and collaborators are consistent with that of the EMR on CVM version. For details, see Collaborator/Sub-account Permissions.

Tag authentication and API authentication settings are consistent with that of the EMR on CVM version. For details, see Collaborator/Sub-account Permissions.

## Collaborator/Sub-user Permissions

Last updated : 2025-02-18 16:59:52

Tencent Cloud EMR (hereinafter referred to as EMR) requires access to or operations on other related cloud products during operation. To ensure normal operation and use of EMR by Collaborator or Sub-user, this document provides guidance on granting relevant permissions.

### Permission Policies Overview

Policy Name	Description	Required or Optional	Description
QcloudCamRoleFullAccess	Full read-write access for CAM Users and Roles	No	Used as a Custom service role, it provides granular permission control for inter- product data access. Refer to Cluster COS Service Role for details
QcloudEMRFullAccess	Full read-write access for EMR (EMR)	No	Full feature operation permissions for EMR products. For details, see Purchase and Management of EMR Clusters
QcloudEMRReadOnlyAccess	Read-only access for EMR (EMR)	No	View permissions for all features of EMR products
QcloudEMRPurchaseAccess	Financial permissions for EMR products	No	For details, see Purchase and Management of EMR Clusters. If purchase or change of configuration is not needed, this permission can be disabled.

#### Caution

The pre-defined QcloudEMRPurchaseAccess policy allows you to manage the financial permissions for purchasing EMR products for all users. When this policy is granted to a user, it includes the financial permissions for CVM, CDB,
and EMR. If you need to restrict a user's ability to purchase CVM and CDB, do not grant the corresponding product ordering permissions.

# Application scenario

### Authorize EMR to access other cloud product permissions

1. **Service Role (mandatory)**: When using EMR services, access to cloud services like CVM, CBS, and TencentDB is required. During the initial product purchase, the service role EMR\_QCSRole for EMR must be assigned and granted the policy QcloudAccessForEMRRole (for requesting basic resources such as CVM, CBS, TencentDB, and read permissions for COS) and EMR's access permissions to cloud resources.

2. Service-related role (optional): If EMR directly writes to or processes data in COS, to ensure data security, EMR needs the corresponding service role EMR\_QCSLinkedRoleInApplicationDataAccess and must bind to the QcloudAccessForEMRLinkedRoleInApplicationDataAccess pre-defined policy to read and write COS resources using temporary keys.

### Special Note on COS Bucket Access Authorization:

1. Since August 20, 2023, for new users or existing users modifying their authorization policies, the default servicerelated role EMR\_QCSLinkedRoleInApplicationDataAccess is granted.

2. The current authorization policy for existing users is to bind the

QcloudAccessForEMRRoleInApplicationDataAccess policy in the service role EMR\_QCSRole.

3. When both the service-related role and service role are authorized, the service-related role is used by default. In the cluster instance information authorization policy, COS will show as authorized, and the cluster COS service role will display the EMR\_QCSLinkedRoleInApplicationDataAccess role.

### Purchase and manage EMR clusters

For scenarios involving resource purchasing, such as creating clusters, adding components, modifying configurations, or expanding collaborators/sub-users, it is necessary to grant QcloudEMRFullAccess. This is according to the Definition TencentDB purchasing policy. If there are no resource purchasing scenarios, such as for service configuration management or restarts, only the QcloudEMRFullAccess policy needs to be granted.

### Caution

For the annual/monthly subscription purchase method, if financial permissions are not granted, a pending order will be generated and linked to an account with financial permissions for approval. The pay-as-you-go purchase method does not support order approval; financial permissions must be granted.

Policy Category	Policy Name	Policy Description
EMR Preset Policy	QcloudEMRFullAccess	Full read-write access for EMR (select one)
EMR Preset Policy	QcloudEMRReadOnlyAccess	Read-only access for EMR (choose



		one of the two)
EMR Preset Policy	QcloudEMRPurchaseAccess	Financial permissions for EMR products

The root account grants the above permissions to the Sub-user or Collaborator. The steps are as follows:

1. Log in to CAM Console, find the corresponding Sub-user or collaborator in **Users > User List**, then click **Authorize**.

User List						
	How to view more information?     CAM protects the security of your sen     subscription. You can also click the us	sitive information. You can click ername to view or edit the user i	the drop-down button (>) on the left side of information.	the list to view more information about t	the user, such as identity security statu	s, groups the user has joined, and mes
	Create User More V				Search by username/I	D/SecretId/mobile/email/remarks (se
	Username	User Type	Account ID	Creation Date	Associated Info	Action
	> Transistante	Root Account	1000103806	2019-06-24 13:21:02		Authorize More 🔻
	+ 🗌 desperandet	Sub-user	2000160209	2020-10-10 16:06:11	L 🔽	Authorize More 🔻
	> 🗆 alaafaa	Sub-user	100015595	2020-09-04 17:27:30	G	Authorize More 🔻
	> 🗆 putagai	Sub-user	10001557 <mark>.</mark> 808	2020-09-03 15:31:26		Authorize More 🔻
	> Distant	Sub-user	1000153 6639	2020-08-14 23:07:35		Authorize More 🔻
	•	Sub-user	100014   54939	2020-07-13 17:50:43	Lø	Authorize More 🔻
	) 🗋 10,000	Sub-user	1000147	2020-06-22 15:21:36		Authorize More 🔻

2. Search for the policy listed in the table above (the following image takes QcloudEMRFullAccess policy as an

example) within Associated Policies. Once the policy is selected, click  $\mathbf{OK}$  to confirm.

3. Grant the EMR financial policy QcloudEMRPurchaseAccess, similar to step 2.

#### **Associate Policy**

Select Policies (21 Total)				1 selected		
Support search by policy name/description/remark	(S	Q		Policy Name	Policy Type	
Policy Name Cross-service access of YouMall to Cloud	Policy Type T	•		QcloudCamSubaccountsAuthorizeRoleFu CAM sub-account authorization service r	Preset Policy	(
CAM sub-account authorization service ro	Preset Policy					
QcloudCollApiKeyManageAccess Provides the ability for an CAM users to	Preset Policy		÷			
QcloudCollMFAManageAccess Provides the ability for an CAM user to m	Preset Policy	ł				
QcloudCollPasswordManageAccess Provides the ability for an CAM sub-user t	Preset Policy	•				
Press Shift to select multiple items						
		Confirm		Cancel		

# **Cluster COS Service Role**

EMR root account or Collaborator and Sub-user with QcloudCamRoleFullAccess can precisely control COS bucket permissions and other cloud resource permissions. For details, see Cluster COS Service Role. Root account grants QcloudCamRoleFullAccess to Sub-user or Collaborator, the specific steps are as follows: 1. Log in to CAM Console, find the corresponding Sub-user or collaborator in **Users > User List**, then click **Authorize**.

User List						
	How to view more information?     CAM protects the security of your sensit     subscription. You can also click the users	ive information. You can click the dre name to view or edit the user informa	op-down button [▶] on the left side ation.	of the list to view more information about the	user, such as identity security stat	us, groups the user has joined, and mes
	Create User More 💌				Search by username,	ID/SecretId/mobile/email/remarks (se
	Username	User Type	Account ID	Creation Date	Associated Info	Action
	> Transistan	Root Account	1000103806	2019-06-24 13:21:02	]. 🖸	Authorize More 🔻
	+ 🗌 desparadoti	Sub-user	20001602091	2020-10-10 16:06:11	], 🖸	Authorize More 🔻
		Sub-user	10001559560	2020-09-04 17:27:30	G	Authorize More 🔻
		Sub-user	10001557_B08	2020-09-03 15:31:26	G	Authorize More 🔻
		Sub-user	10001530639	2020-08-14 23:07:35		Authorize More 🔻
		Sub-user	100014 1054939	2020-07-13 17:50:43	G	Authorize More 🔻
	• 🗌 10_000	Sub-user	1000147	2020-06-22 15:21:36	L 🖸	Authorize More 🔻

2. Search for QcloudCamRoleFullAccess policy within Associated Policies. Once the policy is selected, click **OK** to confirm.

Associate Policy					
Select Policies (2 Total)			1 selected		
qcloudcamr	<b>8</b> Q		Policy Name	Policy Type	
Policy Name	Policy Type <b>T</b>		OcloudCamRoleFullAccess		
QcloudCamReadOnlyAccess Read-only access to Cloud Access Managem	Preset Policy		Full read-write access to Role	Preset Policy	8
QcloudCamRoleFullAccess Full read-write access to Role	Preset Policy	÷			
Press Shift to select multiple items					
	Confirm		Cancel		

# CAM-Enabled EMR API Authorization Granularity Details

Last updated : 2023-12-27 10:00:46

## List of APIs supporting resource-level authorization

EMR supports resource-level authorization. You can grant a specified sub-account the API permission of a specified resource.

### Caution

A permission error may occur when new APIs are subsequently added. In case of such permission error, you can add the missing API permission in the policy based on the error message.

APIs supporting resource-level authorization include:

API	Description	Resource Type/Six-Segment Example Resource
DescribeKeyTabFile	Exports keytab file (user management)	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeShowUserManagerTab	Specifies whether to show the user management tab (user management)	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeResourceSchedule	Gets data from the YARN resource scheduling page	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeCamUserList	Queries CAM user list	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeClusterServiceInfo	Queries service information	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeUserManagerUserList	Queries user list (user management)	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
ModifyUserManagerPwd	Changes user	emr-



	password (user management)	instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyResourcePools	Refreshes dynamic resource pool	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DeleteUserManagerUserList	Deletes user (user management)	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyResourceScheduler	Modifies the YARN resource scheduler (the change will take effect after you click **Apply**)	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyResourceScheduleSwitch	After this switch is turned on, the configuration file of resource scheduling needs to be synced to the resource scheduler first before the resource scheduler page can be viewed and operations can be performed on the page	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
AddUserManagerUserList	Adds user (user management)	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyUserManagerInfo	Modifies user information (user management)	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyResourceScheduleConfig	Modifies the resource configuration of YARN resource scheduling	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyResourceScheduleConfigForRollback	Cancels saving the resource configuration of YARN resource scheduling	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
InquirePriceRefundEmr	Queries the refund	emr-



	amount of terminated node	instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifySecurityGroup	Modifies cluster security group	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeInstanceRenewNodes	Queries renewed nodes of EMR cluster	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
InquirePriceRenewEmr	Queries cluster price for renewal	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeInstancesList	Queries the list of EMR cluster instances	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyFlowStatus	Changes workflow status	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyMasterIp	Updates EMR cluster IP	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
CheckFlowCanBeCancelled	Checks whether the workflow can be canceled	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeModifySpec	Queries target specification	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
AddShellScriptTask	Generates cluster script task	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeShellScriptTaskList	Queries the list of cluster script tasks	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeShellScriptNodes	Queries the list of nodes in a single cluster script task	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeNodeList	Views node	emr-



information	instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Views node information	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Queries the detailed execution result of cluster script task on a single node	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Deletes the records of cluster script tasks	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Gets EMR cluster instance IP	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Queries the API information of TKE cluster	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Adds custom configuration file	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Deletes custom configuration file	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Modifies bootstrap script	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Gets alias	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Gets bootstrap script	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Describes EMR subtask workflow	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
	informationViews node informationQueries the detailed execution result of cluster script task on a single nodeDeletes the records of cluster script tasksGets EMR cluster information of TKE clusterQueries the API information of TKE clusterDeletes custom configuration fileDeletes custom configuration fileGets aliasGets aliasGets bootstrap scriptDescribes EMR subtask workflow



ListConfLogs	Gets configuration delivery logs	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
GenerateScaleoutGoodsDetail	Generates scale-out order	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
GenerateRenewGoodsDetail	Generates renewal order	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
GenerateModifyGoodsDetail	Generates configuration adjustment order	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
ModifyAutoScaleGlobalConf	Updates the global configuration of auto scaling	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeFlowStatusDetail	Queries EMR task execution details	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeFileIps	Queries the list of IPs of the specified configuration file	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeExportConfsList	Queries the list of exportable configuration files	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeFlowStatus	Queries the EMR instance workflow status	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeFlowNum	Queries the number of workflows in EMR cluster	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeClusterNodes	Queries the hardware node information of EMR cluster	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
ClearMetadata	Clears the metadata information of EMR	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId



	cluster and terminates cluster	
DescribeAutoScaleGlobalConf	Gets the global configuration of auto scaling	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeAutoScaleSpecs	Gets auto scaling specification	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyAutoScaleSpecs	Modifies auto scaling specification	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
AddMetricScaleStrategy	Adds metric load- based scaling rule	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DeleteAutoScaleSpec	Deletes auto scaling specification	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeAutoScaleRecords	Gets auto scaling history	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DeleteAutoScaleStrategy	Deletes auto scaling rule	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyStrategyPriority	Changes rule priority	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeAutoScaleMetaRange	Gets auto scaling metadata	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeAutoScaleStrategies	Gets auto scaling rules	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyAutoScaleStrategy	Modifies auto scaling rule	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
AddAutoScaleSpec	Adds auto scaling	emr-



	specification	instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
TerminateAutoScaleNodes	Terminates all automatically scaled nodes	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeCbsEncrypt	Describes CBS encryption	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyConfigGroup	Modifies configuration group	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DeleteConfigGroup	Deletes configuration group	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
AddConfigGroup	Adds configuration group	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeConfigGroup	Describes configuration group	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
SynchronizeGroupConfCheck	Syncs configuration check	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
UnbindInstanceAndNodesTags	Unbinds cluster from tags	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeNodeResourceConfigFast	Quickly gets node specification configuration	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
InstallCos	Enables COS upon installation	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeScaleoutableService	Describes services that can be scaled out	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeNodeResourceConfig	Gets node	emr-



specification configuration	instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Deletes node specification configuration	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Adds node specification configuration	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Sets the default attributes of node specification configuration	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Gets HBase table-level monitoring data overview	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Pulls node resources from the Tag console	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Updates the password of proxy component	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Terminates task nodes	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Terminates nodes	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Terminates EMR instance	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Stops component service	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
Stops component monitoring	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
	specification configurationDeletes node specification configurationAdds node specification configurationSets the default attributes of node specification configurationGets HBase table-level monitoring data overviewPulls node resources from the Tag consoleUpdates the password of proxy componentTerminates task nodesTerminates task nodesStops component serviceStops component monitoring

🔗 Tencent Cloud

StartService	Starts component service	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
StartMonitor	Starts component monitoring	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ScaleOutRouter	Adds router nodes	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeDestroyInfo	Queries the termination information of EMR cluster or node	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ScaleOutInstance	Scales out instance	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
RollBackConf	Rolls back configuration	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
RestartService	Restarts component service	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyServiceParams	Modifies service component parameters	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyResource	Adjusts instance configuration	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
AssignInstancesProject	Assigns EMR cluster to specified project	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyOptionalSpecStatus	Changes the status of optional specification	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyOptionalSpec	Modifies optional specification	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId



DescribeSelectedOptionalSpec	Queries selected optional specification	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyInstanceBasic	Renames cluster	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
InstallSoftware	Installs component	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
InquiryPriceAddRouter	Queries the price of router adding	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
InquiryPriceRenewInstance	Queries the price for renewal	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
InquiryPriceUpdateInstance	Queries the price for configuration adjustment	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
InquiryPriceScaleOutInstance	Queries the price for scale-out	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribleClusterNodes	Queries the hardware node information of EMR cluster	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeServiceNodeInfos	Queries the service process information of EMR cluster	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeServiceGroups	Queries the service group information of EMR cluster	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeServiceConfs	Queries all configuration information of the EMR cluster service	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeScaleoutGoodsDetail	Queries scale-out order	emr- instance qcs::emr:\${region}:uin/\${uin}:e



		instance/\$emrInstanceId
DescribeRouterGoodsDetail	Queries router node order	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeRenewGoodsDetail	Queries renewal order	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeOptionalSpec	Queries the optional specifications of EMR instance node	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeModifyGoodsDetail	Queries configuration adjustment order	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeMetricsDimension	Queries the dimension values of monitoring at different levels	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeMetricMeta	Queries the monitoring metadata of EMR cluster	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeInstanceOplog	Queries the operation log of EMR instance	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeExecCustomScript	Queries the information of custom script	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeInstanceAlerts	Queries the alarm information of EMR cluster	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeInstances	Queries the information of EMR instances	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeInstallSoftwareInfo	Queries the information of components installed and uninstalled for EMR cluster	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId

DescribeExportConfs	Queries export configuration	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
CheckDiskInfo	Checks whether the disk metadata in the console is updated	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
SyncDiskInfo	Updates disk metadata in console	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeServiceConfsNew	Gets the configuration information of component (configuration management page)	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeConfFileList	Gets the list of configuration files (configuration management page)	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeServiceConfCategories	Gets the configuration type of component (configuration management page)	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeServiceConfDiff	Compares configurations (configuration management page)	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeConfigGroupList	Queries the configuration group information of node type	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeHbaseStatus	Displays the information of `DescribeHbaseStatus`	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyHbaseRit	Fixes HBase RIT issue	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
TerminateSparkApp	Terminates Spark	emr-



	job	instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeAccessKey	Gets the `AccessKey` of cluster	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeAttachableDisks	Queries the cloud disks that can be mounted to the node	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
InquirePriceRenewDisks	Queries the price for cloud disk renewal	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
AttachDisks	Mounts cloud disk	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeNodeDataDisks	Queries the data disk information of node	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
InquirePriceResizeDisks	Queries the price for cloud disk expansion	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ResizeDataDisks	Expands cloud disk	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyAutoRenewFlag	Automatic Renewal	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ScaleOutCluster	Adds cluster nodes	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
StartStopServiceOrMonitor	Starts or stops monitoring or service	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeClusterClients	Views the client information on the page	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeDelayedServiceConfig	Gets expired	emr-



	configuration (configuration management page)	instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeFailedServiceConfig	Gets failed configuration (configuration management page)	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeServiceConfDimensionDiff	Compares the configurations at different levels (configuration management page)	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyEmrManagerAgent	Updates EMR Manager Agent	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyServiceConfDiff	Overwrites different configuration (configuration management)	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyYarnLabelState	Delivers an instruction to add, delete, or bind a label	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeYarnScheduleHistory	Views the YARN resource scheduling history	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeYarnDeployMessage	Gets the prompt message	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyYarnDeploy	Applies deployment	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeParentLabels	Gets the list of labels of the parent queue	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DescribeYarnLastestLabels	Gets the latest label information	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId

🔗 Tencent Cloud

ModifyYarnLabels	Syncs YARN node labels	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
ModifyYarnQueue	Modifies the resource pool in resource scheduling	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
ModifyOldLabelConfig	Cancels saving the edited content in YARN label management	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeNodeLabelWebUrl	Gets the web URL of YARN node labels	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeNodeManagerHosts	Gets the list of NodeManager IPs	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
ModifyGlobalScaleConf	Modifies the scaling configuration of cluster, including whether to enable scaling and the scaling type	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeGlobalScaleConf	Gets the scaling configuration of cluster, including whether to enable scaling and the scaling type	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
ModifyManagedScaleGlobalConf	Updates the global configuration of managed scaling	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeManagedScaleGlobalConf	Gets the global configuration of managed scaling	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeServiceComponentInfos	Describes the role information of container cluster	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId
DescribeRssClusterList	Gets the RSS clusters in the same EKS	emr- instance qcs::emr:\${region}:uin/\${uin}:¢

	cluster as the current Spark cluster	instance/\$emrInstanceId
CreateSparkLinkRss	Associates RSS cluster with Spark cluster	emr- instance qcs::emr:\${region}:uin/\${uin}:e instance/\$emrInstanceId
DeleteSparkLinkRss	Disassociates RSS cluster from Spark cluster	emr- instance qcs::emr:\${region}:uin/\${uin}:¢ instance/\$emrInstanceId

# List of APIs supporting API-level authorization

API	Description
RunJobFlow	Creates and runs job
DescribeJobFlow	Queries running jobs
DescribeK8sResourcePrice	Queries the price of K8s resource
DescribePodSpecs	Describes Pod specification
GenerateCreateGoodsDetail	Generates cluster creation order
DescribeLogSearchFileNames	Gets the list of log search files
DescribeInstanceServiceRoleNames	Gets service role names
DescribeCompareMetricsList	Gets the list of comparison metrics
DescribeHeatMapMetricList	Returns the list of metrics in the aggregated cluster server dimension
DescribeHBaseRegionList	Gets the list of HBase regions
DescribeClusterOverview	Queries cluster overview
DescribeEMRNodeOverview	Queries the deployment status of various processes on node
DescribeNodeHardwareInfo	Queries the basic configuration of the host
DescribeTopNMeta	Gets the top N metadata information items

DescribeMetricDataAutoGranularity	Gets monitoring data and automatically sets time granularity
DescribeLogSearchTabs	Gets log search tabs
DescribeEMRHostOverview	Queries the process status on the host details page
DescribeLogSearchRecords	Gets log search content
DescribeTopNByProcess	Gets the top N processes
DescribeDiskInfo	Returns specific disk information
DescribeTopNByHost	Queries the top N hosts on the overview page
DescribeHeatMapDistribution	Returns the heat map data of cluster
DescribeInstanceServiceRoleTable	Gets the data of service role table
DescribeNodeAlias	Gets the alias of EMR node
DescribeInstanceNodes	Gets the host information of cluster
DescribeKeyPairs	Queries the information of key pairs
DescribeHbaseTableMetricData	Queries HBase table-level monitoring data
DescribeKeyPairs	Queries the information of key pairs
DescribeEmrMetaDB	Gets the unified metadatabase information of Hive
ModifyEmrRole	Updates the information of EMR role
DescribeDisasterRecoverGroup	Gets the information of spread placement group
DescribeTags	Pulls all tags of cluster
InquiryPriceCreateInstance	Queries the price of created instance
DescribleAccountBalance	Queries account balance
DescribeSecurityGroup	Gets the information of EMR security group
DescribeCvmSpec	Queries CVM instance specification
DescribeCdbPrice	Queries the price of TencentDB instance
CreateInstance	Creates EMR instance
GetMetricDataForMcController	Gets the monitoring information on the details page in the console

DescribeVpcList	Queries the list of VPCs
DescribeSceneProductInfo	Gets the cluster use case, type, product, and component information from the purchase page
DescribeRegionAndZoneSaleInfo	Gets the region and AZ information from the console and the purchase page
DescribeCgwProjects	Gets the list of CGW projects
DescribeServiceUpgradeVersion	View the service minor version that can be upgraded
ModifyServiceVersion	Upgrades service minor version
DescribeServiceRoleInstanceConstraints	Gets the constraint information of role instances removed from the service
CheckSupportServiceRoleInstance	Checks whether a role instance can be added to the service
DescribeServiceRoleNames	Gets the drop-down list of role names (or role types) of the service
DescribeServiceRoleInstanceDeployableNodes	Queries the list of nodes where the service role can be deployed
DescribeServiceDeletableRoleInstances	Queries the list of role instances that can be deleted from the service
DeleteServiceRoleInstance	Removes role instance
AddServiceRoleInstance	Adds role instance
ModifyResourcesTags	Forcibly modifies tag
CreateCluster	Creates cluster
RepairDisk	Repairs the business after the customer's corrupted disk is replaced at the CVM side
CreateOutsideClient	Provides component packages and deployment scripts for servers outside the EMR cluster to support their access to EMR and job submission
ModifyCvmReplaceFlag	Modifies the EMR cluster to support abnormal node replacement



For more information on resource-level and API-level authorization, see Authorization Granularity Scheme.

# Authentication Granularity Scheme

Last updated : 2023-12-27 10:01:09

# Six-Segment Resource Description

qcs: Describes the abbreviation of qcloud service, indicating that a resource is a Tencent Cloud resource. This field is required.

project\_id: Describes the project information, which is used to enable compatibility with legacy CAM logic and can be left empty.

service\_type: Describes the product abbreviation, such as CVM .

region: Describes the region information, such as bj (see the CAM documentation).

account: Describes the root account of the resource owner, such as uin/164256472 .

resource: Describes the detailed resource information of each product, such as instance/instance\_id1 or

instance/\* .

## Tag Authentication Operation Guide

This section describes how to authenticate EMR tags. Before using tags to manage the authentication of EMR resources, you need to plan tags based on your departments or organizations; for example, you can set usernames, clusters, and tags for departments A, B, and C respectively. For more information on tags, see Tag.



After completing the planning, to enable department A to manage cluster A1, you need to perform the following operations in sequence.

### 1. Create a tag and tag the cluster

Enter the Tag List page and click Create Tag.

### Note

The tag\_A tag is used as an example here and below.

Tag	Tag List			
Resource Tag	Create Tag Delete			
🖃 lag List	Tag Key T	Tag Value 🔻	Resource Count	Operation
	tke-clusterid	cls-2dqpztmb	0	Bind Resources Delete
	tke-name	1307265578-emr-66pdbnbo/emr-66pdbnbo-hs-service	0	Bind Resources Delete
	tke-created	yes	34	Bind Resources Delete
	tke-kind	service	0	Bind Resources Delete
	tke-clusterid	cls-41n2ew8l	0	Bind Resources Delete
	test01	03	0	Bind Resources Delete
	test02	02	0	Bind Resources Delete
	test01	01	0	Bind Resources Delete
	tsg_A	Dept_A	0	Bind Resources Delete
	test	test01	1	Bind Resources Delete
	Total items: 10			10 🔻 / page 🛛 H 🔄 1 🔢 / 1 page

Enter the tag key and tag value and click **OK**.

Create Tag		>
<ul> <li>Enter a new tag key and value</li> <li>Each tag key can have a magental</li> </ul>	alue, or select a tag key and add a new value to it. aximum of 1,000 values. You can add 10 values at a time.	
Tag Key tag_A	Tag Value : Dept_B 😮 Dept_A	Delete
Add Tag Key	Dept_A OK Cancel	

On the Cluster List page, select the target cluster, click More at the top, and click Edit Tag.

Cluster list	🔇 Guangzhou 1 🔻								Cluster management d
Create clust	ter More 🔻 Se	eparate keywords with " "; press Ent	er to separate filter tags		Q				
V ID/Name	Edit tag	Monitor	Project	AZ	Billing mode	Network	Master node public IP	Tag	Action
Ei.	Cluster is runnin	9 ili	DEFAULT PROJECT	Guangzhou Zone 7	Pay-as-you-go	Roy-001 ceshi001		©1	Service Resource More
Total 1 item								Lines per page 10 🔻	1 /1 page

On the Edit Tag page, select the created tag\_A tag and click Complete.

Tags are used to m	anage resource	es hy category in differ	ent dimensions	f the existing tags
don't meet your re	quirements, yo	ou can manage tags 🗹 .		in the existing tags
1 resource(s) select	ed			
tag_A	Ŧ	Dept_A		×
+ Add				

### 2. Create a custom policy

Currently, certain EMR APIs don't support authentication by tag; therefore, you need to set two policies:

Policy\_Dept\_A1 (for APIs that support authentication by tag) and Policy\_Dept\_A2 (for APIs that don't).

### Create custom policy Policy\_Dept\_A1

Go to the Policies page in the CAM console and click **Create Custom Policy**.

#### Select Authorize by Tag.

In **Tag Policy Generator**, select **Dept\_A** for **Authorized Users**, leave **User Groups** empty, select the tag key and tag value of the planned tag for **Tag Key** and **Tag Value** respectively, add services and operations as needed, and click **Next**.

Create the Policy\_Dept\_A1 policy, replace the content of the action and resource fields with the following in **Policy**, and click **Complete**.

```
{
 "effect": "allow",
 "action": [
     "emr:DescribeClusterNodes",
     "emr:DescribeInstancesList"
],
 "resource": [
     " * "
],
 "condition": {
     "for_any_value:string_equal": {
              "qcs:resource_tag": [
                      "tag_A&Dept_A",
                      "tag_B&Dept_B",
                      "tag_C&Dept_C",
                      "tag_D&Dept_D"
              ]
     }
 }
}
```

### Caution

The content of the resource field needs to be entered based on the actual resource. For more information on how to specify it, see Resource Description Method.

### Create custom policy Policy\_Dept\_A2

Go to the Policies page in the CAM console and click Create Custom Policy.

Select Create by Policy Syntax.

### In Select Policy Template, select Blank Template and click Next.

Create the Policy\_Dept\_A2 policy, clear the existing content in **Policy**, paste the following content, and click **Complete**.

```
{
    "version": "2.0",
    "statement": [
        {
            "effect": "allow",
                "action": [
               "emr:RunJobFlow","emr:DescribeJobFlow","emr:FindClusterListForMC","emr:EmrScale
        outClusterWithAuth","emr:CreateClusterForMC","emr:CheckSoftRelation","emr:Descri
beK8sResourcePrice","emr:DescribePodSpecs","emr:DescribePodSpec","emr:Describe
TkeWhiteList","emr:GenerateCreateGoodsDetail","emr:DescribeLogSearchFileNames",
        "emr:DescribeEMRInstancesExtra","emr:DescribeClusterHardwareInfo","emr:Describe
```

InstanceServiceRoleNames", "emr:DescribeServiceComponents", "emr:DescribeCompareM etricsList", "emr:DescribeHeatMapMetricList", "emr:DescribeHBaseRegionList", "emr: DescribeAutoScaleWhiteList", "emr:DescribeOptionalSpecWhiteList", "emr:DescribeSe rviceOverview", "emr:DescribeEMRInstances", "emr:DescribeEMRNodeType", "emr:Descri beEMRNodes", "emr:DescribeStepByTimeRange", "emr:DescribeClusterOverview", "emr:De scribeEMRNodeOverview", "emr:DescribeNodeHardwareInfo", "emr:DescribeTopNMeta", "e mr:DescribeMetricDataAutoGranularity", "emr:DescribeLogSearchTabs", "emr:Describe EMRHostOverview", "emr:DescribeLogSearchRecords", "emr:DescribeTopNByProcess", "em r:DescribeDiskInfo", "emr:DescribeMetricInfo", "emr:DescribeInstanceServiceAbstra ct", "emr:ModifyMetricMetaPerInstance", "emr:DescribeOverviewData", "emr:DescribeC lusterHardwareType", "emr:DescribeClusterEmergencyCallback", "emr:DescribeTopNByH ost","emr:DescribeHeatMapDistribution","emr:DescribeInstanceServiceRoleTable"," emr:DescribeEmrSaleInfoExtend", "emr:DescribeNodeAlias", "emr:DescribeInstanceNod es","emr:DescribeKeyPairs","emr:DescribeHbaseTableMetricData","emr:DescribeHbas eMetricMeta","emr:DescribeEmrMetaDB","emr:CheckMetaDBNet","emr:ModifyEmrRole"," emr:DescribeEmrRole", "emr:DescribeDisasterRecoverGroup", "emr:DescribeTags", "emr :InquiryPriceCreateInstance", "emr:GetCreateGoodsDetail", "emr:DescribleAccountBa lance", "emr:DescribeSecurityGroup", "emr:DescribeMetricData", "emr:DescribeEmrSal eInfo","emr:DescribeCvmSpec","emr:DescribeCdbPrice","emr:CreateInstance","emr:C heckCustomConfig", "emr:CheckCosInfo", "emr:GetCreateGoodsDetailList", "emr:GetSca leoutGoodsDetailList","emr:GetRenewGoodsDetailList","emr:UpdateWebProxyPassForM cController", "emr:SubmitServiceParamsForMC", "emr:ReleaseClusterForMC", "emr:Rest artServiceForMC", "emr:GetNodeListForMcController", "emr:GetMetricDataForMcContro ller", "emr:GetServiceGroupForMcController", "emr:EmrDestroyTaskNodeWithAuth", "emr:DescribeModifySpec", "emr:DescribeMasterIp", "emr:DescribeMetricsDimension" ],

```
}

]
```

### 3. Authorize and authenticate the sub-user

On the **User List** page, find the target sub-user and click **Authorize** on the right. For more information on how to create a sub-user, see Creating Sub-user.

Select Policy\_Dept\_A1 and Policy\_Dept\_A2 and click Complete.

Log in as a sub-user for authentication.

Tencent Cloud

At this point, when logging in to EMR, the sub-user can see and manage the test\_A cluster only through the tag\_A tag.

#### Caution

}

As the APIs in the Policy\_Dept\_A2 policy can only be fully listed, a permission error may occur when new APIs are subsequently added. In case of such permission error, submit a ticket or contact us for assistance.

## Resource-Level and API-Level Authentication Operation Guide

This section describes how to authorize EMR resources. When using resource-level authentication to authorize an EMR feature, you need to plan resources for your departments or organizations first; for example, you can plan usernames, APIs, and clusters for departments A, B, and C as shown below.

After completing the planning, to grant department A only access to the DescribeRegionAndZoneSaleInfo API of the cluster, you need to perform the following operations in sequence:

### 1. Create a custom policy

Go to the Policies page in the CAM console and click Create Custom Policy.

### Select Create by Policy Syntax.

### In Select Policy Template, select Blank Template and click Next.

Create a custom policy named DescribeRegionAndZoneSaleInfo-test , clear the existing content in Policy, configure the access deny policy for EMR's DescribeRegionAndZoneSaleInfo API, and save the configuration.

```
{
    "version": "2.0",
    "statement": [
        {
            "effect": "allow",
            "action": [
"emr:DescribeRegionAndZoneSaleInfo"
                ],
                "resource": "*"
        }
    ]
}
```

### Caution

For APIs that support resource-level authentication, see **List of APIs supporting resource-level authorization**. As the APIs can only be fully listed, a permission error may occur when new APIs are subsequently added. In case of such permission error, you can add the missing API permission in the policy based on the error message.

### 2. Authorize and authenticate the sub-account

On the **User List** page, find the target sub-user and click **Authorize** on the right. For how to create a sub-user, see Creating Sub-user.

Select the DescribeRegionAndZoneSaleInfo-test custom policy and click Complete.

After you log in with the sub-account and call the above API, if a pop-up window similar to the following requesting authorization is displayed, the authentication has been successfully configured.

# **Cluster COS Service Role**

Last updated : 2024-10-30 09:51:23

The custom service role feature enables you to create a CAM service role for granting access to COS resources. Please select **Tencent Cloud Product Service** as the service type and **Elastic MapReduce** as the service supporting the role. If no role is set, the system defaults to using the

EMR\_QCSLinkedRoleInApplicationDataAccess service-related role for accessing COS resources.

The following steps are for creating a new custom Definition service role and associating it with the cluster COS service role.

#### Note:

### Note:

1. The cluster COS service role by default displays the currently bound role identity, and the cluster uses this role identity for reading and writing COS resources.

2. If both the EMR service-related role and the service role are authorized simultaneously, the new cluster will default to using the service-related role to access COS. The instance information authorization policy will show COS as authorized, and the cluster COS service role will display the EMR\_QCSLinkedRoleInApplicationDataAccess service-related role.

3. When both the service-related role and service role for EMR are authorized, and the service-related role is not authorized via the EMR on CVM Console cluster instance information cluster COS service role binding, the default service role EMR\_QCSRole and identity QcloudAccessForEMRRoleInApplicationDataAccess will be used for COS access. If you need to use the service-related role, please manually bind the

EMR\_QCSLinkedRoleInApplicationDataAccess service-related role in the cluster instance information cluster COS service role.

## Step 1. Customize a Permission Policy

1. Log in to the CAM console, click Create Custom Policy, and select Create by Policy Syntax in the Select Policy Creation Method pop-up window.

2. On the Create by Policy Syntax page, select Blank Template as the template type.

3. Set the syntax policy as follows:

```
{
    "version": "2.0",
    "statement": [
        {
            "action": "cos:*",
            "effect": "allow",
            "resource": "qcs::cos::uid/appId:bucketName/*"
```

}

Where appId is the AppID of the root account, and bucketName is the name of the target bucket. Generate a policy named TestPolicy . In practical cases, you can customize the name.

## Step 2. Create a Custom Role

1. In the CAM console, click Create Role to open the Select role entity pop-up window, select Tencent Cloud Product Service to go to the Create Custom Role page, and select Elasticsearch MapReduce (emr) in Product Service.

2. Associate the TestPolicy policy generated in step 1. In practical cases, you can associate a policy as needed.

3. Set the role tag keys and values and click Next.

4. Generate a custom role named EMRCOSROLE . In practical cases, you can customize the name.

### Step 3. Bind the Role with an EMR Cluster

In the EMR console, click a cluster ID/name to go to the instance information page, select **Instance info > Basic** configuration > Custom Service Role and then click Set. Set the custom service role to the EMRCosRole role generated in step 2. In practical cases, you can customize the name.

÷	Instance info		⊟ Help docu
Cluster overview	Basic configuration		
Instance info	Instance ID e	Region info Guangzhou—Guangzhou Zone	.7 Network info Roy-001
Cluster services	Security group sg-Oru72wdp	Creation time 2023-05-17 17:44:46	Master public IP Check for updates
Cluster resource v	Billing mode Pav-as-you-on	Auto replenishment (i)	COSID Unauthorized Authorize
Cluster monitor	Custom service role () Not cat Sat	Host login(1) Password Set	
* DashBoard BETA			security group -
Cluster events			
Log search	Software info		Deployment info
Cluster inspection	Cluster type Hadoop		High availability cluster Enabled
Alarm records	Use cases Hadoop-Default		Cloud disk encryption Not enabled
<ul> <li>Java analysis</li> </ul>	Product version EMR-V2.7.0		Guangzhou Zone 7
Auto scaling	Components to deploy hdfs-2.8.5,yam-2.8.5,zookeeper-3.6.3,openIdap-2.4.44,knox	-1.6.1, hive-2.3.9, hbase-2.4.5, kudu-1.15.0	Subnet ceshi001
Scrints v	MetaDB cdb-axtt2mkp		Node Common x3 Master x2 Core x3 Task x1 Router x0
licere	Hive metadatabase cdb-axtt2mkp		
Operation logs	Kerberos mode Not enabled		
operation logs			

Note:

1. Make sure that the custom service role associated with the cluster's COS service role includes specific access permissions for COS. If the associated role does not include the specific access permissions for COS, it will not be able to access COS.

2. Query Custom Service Role or Modify Custom Service Role with Predefined Permission Policy CAM.

# Setting Tags

Last updated : 2023-12-27 10:01:45

## Overview

Tags are key-value pairs provided by Tencent Cloud EMR to identify cluster types or node resources. You can manage clusters or node resources in a categorized manner by using tags in various dimensions such as business, purpose, and responsible person, or use tags to easily identify them. Tag key-value pairs have no semantic meaning for EMR and are parsed and matched strictly based on strings.

## Limits

A tag is a key-value pair. You can set tags for EMR clusters or node resources to manage them in a categorized manner. Then, you can easily view and identify clusters and node resources under the corresponding tags. You can edit tags in the EMR console.

There are several limits on tags:

Quantity limit: Each cluster or node can have up to 50 tags. (**Up to five tags can be added at a time**.) Tag key limits:

You cannot use "qcloud", "tencent", or "project" at the beginning of a tag key because they are reserved by the system. A tag key can contain up to 255 characters. Only digits, letters, and the + = . @ - symbols are supported. Tag value limit: A tag value can contain up to 127 characters. Only digits, letters, symbols + = . @ -, and a null string are supported.

## Directions

### **Editing cluster tags**

1. Log in to the EMR console, select the target cluster on the Cluster list page and click More > Edit tag.

Cluster list 🔇 Guar	ngzhou 1 💌								Cluster management (
Create cluster M	lore 🔻 Separate k	eywords with " "; press Er	iter to separate filter tags		Q				
D/Name	dit tag	Monitor	Project	AZ	Billing mode	Network	Master node public IP	Tag	Action
emr-edqx53yb	Cluster is running	di	DEFAULT PROJECT	Guangzhou Zone 7	Pay-as-you-go	Roy-001 ceshi001	111.230.90.160	©1	Service Resource More
Total 1 item								Lines per page 10 🔻	H - 1 /1 page

2. Add, modify or delete tags as needed in the pop-up dialog box.

### Note

You can edit tags for up to 20 clusters at a time.

Edit Tag				×
Tags are used to mana don't meet your requi	age resources rements, you	by category in differer can manage tags 🗹 .	it dimensions. If	the existing tags
1 resource(s) selected				
tke-created	•	yes	Ŧ	×
Tag Key		Tag Value	Ψ.	×
+ Add				
	_		_	
		OK Cancel		

### **Editing node tags**

- 1. Log in to the EMR console and click a cluster ID/name to go to the cluster details page. Select Cluster resources
- > Resources, select the target node, and click More > Edit tag.

Resources												🖽 Help documer
i The co page :	urrent cluster res shall prevail. It is	ources were not allowed	e purchased base d to adjust the fe	ed on the off e policy on t	ficial EMR billing the CVM Conso	g rules and the r le.	enewal status	es of CVM resources are managed by E	MR in a unified manner. For mo	onthly-subscription CVM r	esources, the bill	ing statuses shown on the current
All nodes	Master	Core	Common	Task	Router	Metadb	Recyle	Renew				
Scale out	Scale in	More	<b>*</b>									
Separate keyv		Check	disk update									
- Resour	rce ID	Set no	de specification					Configuration	Creation time 1	Expiration time	Tag	Operation
	w2ukf7z	Chang Set au Chang	e from pay-as-yc to-renewal e configuration			0 (privat 0.160 (j	te) public)	EMR StandardSA2 CPU: 4-core; memory: 16GB System disk: Cloud SSD 50G x 1	2023-05-17 17:46:00	Pay-as-you-go	©1	Change N configuration
								Data disk: Cloud SSD 200G x 1				

2. Add, modify or delete tags as needed in the pop-up dialog box.

### Note

You can edit tags for up to 20 clusters at a time.

Tags are used to man don't meet your requi	age resources irements, you	by category in differer can manage tags 🗹 .	nt dimensions.	If the existing tags
1 resource(s) selected				
tke-created	Ŧ	yes	•	×
Tag Key		Tag Value	Ŧ	×
+ Add				
		OK Cancel		
## **Bootstrap Actions**

Last updated : 2023-12-27 10:02:08

## Overview

A bootstrap action involves executing custom scripts during cluster creation so that you can modify the cluster environment, install third-party software, and utilize proprietary data. This action runs the bootstrap script during cluster creation (including cluster scale-out) and cluster termination (including cluster scale-in), except for router nodes.

Currently, you can specify a bootstrap action only for cluster creation and termination in the console. You can also specify a bootstrap action for the scale-out or scale-in process via API. By default, if no bootstrap action is specified via API, the bootstrap action specified for cluster creation is executed during the scale-out process, and the bootstrap action specified for cluster termination is executed during the scale-in process.

1. The bootstrap action specified for cluster creation (including cluster scale-out) can be executed in any of the following circumstances:

a. After server initialization: After server resource initialization and before EMR cluster software installation.

b. Before cluster start: Before cluster service start.

c. After cluster start: After cluster service start.

2. The bootstrap action specified for cluster termination (including cluster scale-in) can be executed in the following circumstance:

a. Before service deactivation: Before cluster service deactivation.

Bootstrap actions run scripts during cluster creation and scale-out in the order in which the scripts are added. The number of bootstrap actions cannot exceed 16.

#### Note:

We recommend that you create a small pay-as-you-go cluster first to test whether the bootstrap action works properly, and if yes, create a cluster for actual business.

### Directions

#### Option 1: Add a bootstrap action when creating a cluster on the purchase page

1. Choose **Basic Configuration** > **Advanced Settings** > **Add Bootstrap Action** to add a bootstrap action.



Advanced settings					
Bootstrap Actions 🚯	Run	Name	Script Location	Parameter	Operation
			No data yet		
			+ Add Bootstrap Actio	n	
Tag 🚯			+ Add		

2. Edit or delete the added bootstrap action as needed.

Add Bootstrap Action							
Run	After node initialization	Before cluster start	After cluster start				
	Before service stop						
Name	Enter						
	1-64 characters; supports Chine	ese characters, letters, digits	:, -, and _				
Script Location	Enter						
	Use the COS STANDARD stora beijing.myqcloud.com/data/test	age class. Script example: ht .sh	tps://beijing-111111.cos.ap-				
Parameter	Enter						
	Follow standard Shell specificat	tions.					
	Cancel	Subr	nit				

Select a time for running the script and enter relevant parameters.

Name: We recommend you enter your object name.

Script Location: We recommend you copy the location information from the COS details page. Go to the COS console, click **Bucket List**, select the target script, and choose **Operation** > **Details**.

On the **Details** page, you can see the object name and object address.

Parameter: This refers to the parameters for running the script. Separate multiple parameters by spaces, and do not add spaces in individual parameters. The total length of **Parameter** and **Name** cannot exceed 240 characters.

#### Option 2: Add a bootstrap action on the Basic information page of the cluster

1. Log in to the EMR console and click the ID/Name of the target cluster in the cluster list to enter the cluster details page. Then, choose Basic information > Bootstrap Actions and click Add Bootstrap Action.

2. Edit or delete the added bootstrap action as needed. Select a time for running the script and enter relevant parameters.

Name: We recommend you enter your object name.

Script Location: We recommend you copy the location information from the COS details page. Go to the COS console, click **Bucket List**, select the target script, and choose **Operation** > **Details**.

Parameter: This refers to the parameters for running the script. Separate multiple parameters by spaces, and do not add spaces in individual parameters. The total length of **Parameter** and **Name** cannot exceed 240 characters.

## Viewing Bootstrap Result

Currently, you can specify a bootstrap action for scale-out via API but not in the console. If a bootstrap action is specified, it will be executed during scale-out; otherwise, the one specified for cluster creation will be executed.

1. View the bootstrap result in the script's system log.

Logs and script files to be executed are stored in the /usr/local/service/scripts/ directory. The script system log is script\_syslog . Naming convention 1: "Execution order" + "" + "Time" + "Script name" + "" + stderr.

Naming convention 2: "Execution order" + "" + "Time" + "Script name" + " + stdout.

#### Note:

The scripts will be executed on all types of nodes, and the script files and log files output by script execution will be stored on each node.

Bootstrap scripts need to be encoded in UTF-8.

2. View the bootstrap result in the Task Center.

Log in to the EMR console, click **Task Center** on the left sidebar, or enter a cluster, click **Task** in the top-right corner, and select the target process (creating cluster, scaling out cluster, or initializing node). Then, you can click **Run Details** in the service initialization step of **Task details** to view the bootstrap result.

## Software Configuration

Last updated : 2023-12-27 10:02:29

## Feature

Software configuration enables you to customize configurations of components such as HDFS, YARN, and Hive when creating a cluster.

## Custom Software Configuration

Software programs such as Hadoop and Hive have many configuration items. By using the software configuration feature, you can customize component parameters when creating a cluster. During the configuration, you need to provide the corresponding JSON files as required. You can customize the files or generate them by exporting software configuration parameters of an existing cluster for quick cluster creation. For more information on how to export software configuration parameters, please see Exporting Software Configuration.

Currently, only parameters in the following files can be customized: HDFS: core-site.xml, hdfs-site.xml, hadoop-env.sh, log4j.properties YARN: yarn-site.xml, mapred-site.xml, fair-scheduler.xml, capacity-scheduler.xml, yarn-env.sh, mapred-env.sh Hive: hive-site.xml, hive-env.sh, hive-log4j2.properties

#### Sample JSON file and description:

```
{
    {
        "serviceName": "HDFS",
        "classification": "hdfs-site.xml",
        "serviceVersion": "2.8.4",
        "properties": {
            "dfs.blocksize": "67108864",
            "dfs.client.slow.io.warning.threshold.ms": "900000",
            "output.replace-datanode-on-failure": "false"
        }
    },
    {
        "serviceName": "YARN",
        "classification": "yarn-site.xml",
    }
}
```

```
"serviceVersion": "2.8.4",
                      "properties": {
                                 "yarn.app.mapreduce.am.staging-dir": "/emr/hadoop-yarn/staging",
                                 "yarn.log-aggregation.retain-check-interval-seconds": "604800",
                                 "yarn.scheduler.minimum-allocation-vcores": "1"
                      }
           },
                      "serviceName": "YARN",
                      "classification": "capacity-scheduler.xml",
                      "serviceVersion": "2.8.4",
                      "properties": {
                                 "content": "<?xml version=\\"1.0\\" encoding=\\"UTF-8\\"?>\\n<?xml-</pre>
stylesheet type=\\"text/xsl\\" href=\\"configuration.xsl\\"?>\\n<configuration>
<property>\\n
                                                        <name>yarn.scheduler.capacity.maximum-am-resource-
percent</name>\\n
                                                                    <value>0.8</value>\\n</property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<prime<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property>\\n<property<\\n<property>\\n<property>\\n<property>\\n<property<\\n<pr
<name>yarn.scheduler.capacity.maximum-applications</name>\\n
<value>1000</value>\\n</property>\\n<property>\\n
<name>yarn.scheduler.capacity.root.default.capacity</name>\\n
<value>100</value>\\n</property>\\n<property>\\n
<name>yarn.scheduler.capacity.root.default.maximum-capacity</name>\\n
<value>100</value>\\n</property>\\n<property>\\n</property>\\n
<name>yarn.scheduler.capacity.root.default.user-limit-factor</name>\\n
<value>1</value>\\n</property>\\n<property>\\n
<name>yarn.scheduler.capacity.root.queues</name>\\n
<value>default</value>\\n</property>\\n</configuration>"
                      ł
           }
]
```

#### Configuration parameter descriptions:

serviceName: component name, which must be in uppercase.

classification: filename, which must be a full name with file extension.

serviceVersion: component version, which must be the same as the corresponding component version in the EMR product version.

properties: parameters that need to be customized.

If you want to modify configuration parameters in capacity-scheduler.xml or fair-scheduler.xml, set key in properties to content, and set value to the content of the entire file.

If you want to adjust the component configuration of an existing cluster, you can configure component parameters.

### Accessing External Clusters

After configuring the HDFS access address information of an external cluster, you can read data in it.

#### **Configuration during purchase**

EMR allows you to configure access to external clusters when you create an EMR cluster. On the purchase page, you only need to import a JSON file that meets the requirements in the software configuration section. Below is an example based on assumption:

#### Assumption

Assume that the nameservice of the external cluster to be accessed is HDFS8088 , and the access method is as follows:

```
<property>
    <name>dfs.ha.namenodes.HDFS8088</name>
    <value>nn1,nn2</value>
</property>
<property>
    <name>dfs.namenode.http-address.HDFS8088.nn1</name>
    <value>172.21.16.11:4008</value>
</property>
<property>
    <name>dfs.namenode.https-address.HDFS8088.nn1</name>
    <value>172.21.16.11:4009</value>
</property>
    <name>dfs.namenode.rpc-address.HDFS8088.nn1</name>
    <value>172.21.16.11:4007</value>
<property>
    <name>dfs.namenode.http-address.HDFS8088.nn2</name>
    <value>172.21.16.40:4008</value>
</property>
<property>
    <name>dfs.namenode.https-address.HDFS8088.nn2</name>
    <value>172.21.16.40:4009</value>
</property>
<property>
    <name>dfs.namenode.rpc-address.HDFS8088.nn2</name>
    <value>172.21.16.40:4007</value>
<property>
```

#### **JSON file and description:**

Taking the assumption as an example, import the JSON file (the requirements for its content are the same as those for custom software configuration) in the box.

```
[
{
    "serviceName": "HDFS",
    "classification": "hdfs-site.xml",
    "serviceVersion": "2.7.3",
    "properties": {
```



```
"newNameServiceName": "newEmrCluster",
    "dfs.ha.namenodes.HDFS8088": "nn1,nn2",
    "dfs.namenode.http-address.HDFS8088.nn1": "172.21.16.11:4008",
    "dfs.namenode.https-address.HDFS8088.nn1": "172.21.16.11:4009",
    "dfs.namenode.rpc-address.HDFS8088.nn1": "172.21.16.11:4007",
    "dfs.namenode.http-address.HDFS8088.nn2": "172.21.16.40:4008",
    "dfs.namenode.https-address.HDFS8088.nn2": "172.21.16.40:4009",
    "dfs.namenode.https-address.HDFS8088.nn2": "172.21.16.40:4009",
    "dfs.namenode.https-address.HDFS8088.nn2": "172.21.16.40:4008",
    "dfs.namenode.https-address.HDFS8088.nn2": "172.21.16.40:4009",
    "dfs.namenode.https-address.HDFS8088.nn2": "172.21.16.40:4009",
    "dfs.namenode.https-address.HDFS8088.nn2": "172.21.16.40:4009",
    "dfs.namenode.rpc-address.HDFS8088.nn2": "172.21.16.40:4009",
    "dfs.namenode.rpc-address.HDFS8088.nn2": "172.21.16.40:4009",
    "dfs.namenode.rpc-address.HDFS8088.nn2": "172.21.16.40:4007"
}
```

#### Configuration parameter description

serviceName: component name, which must be "HDFS".

classification: filename, which must be "hdfs-site.xml".

serviceVersion: component version, which must be the same as the corresponding component version in the EMR product version.

properties: the content must be the same as that in the assumption.

newNameServiceName: it indicates the nameservice of the newly created cluster, which is optional. If this parameter is left empty, its value will be generated by the system; if it is not empty, its value must consist of a string, digits, and hyphen.

Access to external clusters is supported only for high-availability clusters.

Access to external clusters is supported only for clusters with Kerberos disabled.

#### **Configuration after purchase**

EMR allows you to use the configuration distribution feature to access external clusters after creating an EMR cluster. **Below is the assumption:** 

Assume that the nameservice of the cluster is HDFS80238 (if it is not a high-availability cluster, the

nameservice will usually be masterIp:rpcport , such as 172.21.0.11:4007).

The nameservice of the external cluster to be accessed is HDFS8088 , and the access method is as follows:

```
<property>
<name>dfs.ha.namenodes.HDFS8088</name>
<value>nn1,nn2</value>
</property>
<property>
<name>dfs.namenode.http-address.HDFS8088.nn1</name>
<value>172.21.16.11:4008</value>
</property>
```

<property></property>
<name>dfs.namenode.https-address.HDFS8088.nn1</name>
<value>172.21.16.11:4009</value>
<name>dfs.namenode.rpc-address.HDFS8088.nn1</name>
<value>172.21.16.11:4007</value>
<property></property>
<name>dfs.namenode.http-address.HDFS8088.nn2</name>
<value>172.21.16.40:4008</value>
<property></property>
<name>dfs.namenode.https-address.HDFS8088.nn2</name>
<value>172.21.16.40:4009</value>
<property></property>
<name>dfs.namenode.rpc-address.HDFS8088.nn2</name>
<value>172.21.16.40:4007</value>
<property></property>

If the information above is in the EMR cluster, you can view it on the management page of configuration distribution or log in to the server to view the /usr/local/service/hadoop/etc/hadoop/hdfs-site.xml file.

1. Enter the Configuration Distribution page and select the hdfs-site.xml file of the HDFS component.

2. Change the value of the configuration item dfs.nameservices to HDFS80238, HDFS8088 .

3. Add configuration items and their values.

Configuration Item	Value
dfs.ha.namenodes.HDFS8088	nn1, nn2
fs.namenode.http-address.HDFS8088.nn1	172.21.16.11:4008
dfs.namenode.https- address.HDFS8088.nn1	172.21.16.11:4009
dfs.namenode.rpc-address.HDFS8088.nn1	172.21.16.11:4007
fs.namenode.http-address.HDFS8088.nn2	172.21.16.40:4008
dfs.namenode.https- address.HDFS8088.nn2	172.21.16.40:4009
dfs.namenode.rpc- address.HDFS8088.nn2	172.21.16.40:4007
dfs.client.failover.proxy.provider.HDFS8088	org.apache.hadoop.hdfs.server.namenode.ha.ConfiguredFailoverPr



dfs.internal.nameservices

HDFS80238

dfs.internal.nameservice needs to be added; otherwise, if the cluster is scaled out, the "datanode" may report an error and be marked as "dead" by the namenode .

4. Distribute the configuration by using the Configuration Distribution feature.

For more information on relevant configurations and principles, please see the community documentation.

## Mounting CHDFS Instance

Last updated : 2023-12-27 10:02:48

## Overview

CHDFS is a high-performance distributed file system with standard HDFS access protocols and hierarchical namespaces. EMR supports read/write of data in CHDFS. This document describes how to mount a CHDFS instance to an EMR cluster.

## Directions

#### Scenario 1. Mounting a CHDFS instance to a new cluster

#### Note:

New cluster: This refers to clusters created on or after December 31, 2019. For new clusters, the default CHDFS mounting address of EMR is /data/emr/hdfs/tmp/chdfs .

An EMR cluster is automatically adaptive to CHDFS. Create a CHDFS instance and set permissions reasonably to interconnect the CHDFS instance and EMR cluster. The configuration steps are as follows:

- 1. Create a CHDFS instance in the same region as the EMR cluster as instructed in Creating CHDFS Instance.
- 2. Create a permission group as needed as instructed in Creating Permission Group.
- 3. Create permission rules as needed as instructed in Creating Permission Rule.
- 4. Create a mount point on the same network as the EMR cluster as instructed in Creating Mount Point.

5. Check the connectivity between the CHDFS instance and the EMR cluster. Use the hadoop fs command line tool to run the hadoop fs -ls ofs://\${mountpoint}/ command. Here, mountpoint is the mounting address. If the file list is output properly, the CHDFS has been successfully mounted.

#### Scenario 2. Mounting a CHDFS instance to a legacy cluster

#### Note:

Legacy cluster: This refers to clusters created before December 31, 2019.

For more information on how to mount a CHDFS instance to a legacy EMR cluster, see Mounting CHDFS Instance.

## Unified Management of Hive Metadata

Last updated : 2023-12-27 10:03:06

## Overview

When you choose to deploy the Hive component, the system provides two storage methods for Hive metadata: you can store the metadata in a MetaDB instance separately purchased for the cluster (default method) or associate the metadata with EMR-MetaDB or a self-built MySQL database. In the latter case, metadata will be stored in the associated database and will not be deleted when the cluster is terminated. This way, Hive metadata can be managed in a unified manner.

## Prerequisites

**Cluster default**: A MetaDB instance is automatically purchased to store Hive metadata, alongside metadata of other components. The MetaDB instance will also be terminated when the cluster is terminated. If you need to retain the metadata, please go to the TencentDB console and save it manually in advance.

Hive metadata is stored together with the metadata of Hue, Ranger, Oozie, Presto, Druid, and Superset. A MetaDB instance needs to be purchased separately for the cluster as the metadatabase.

The MetaDB instance will be terminated when the cluster is terminated, i.e., the metadata stored in it will be deleted. **Associate EMR-MetaDB**: When a cluster is created, the system will pull a MetaDB instance available in the cloud as the Hive metadatabase, so you don't need to purchase a MetaDB instance separately, which helps reduce your storage costs. Moreover, Hive metadata will not be deleted when the cluster is terminated.

The ID of the available MetaDB instance is the ID of an existing MetaDB instance in the EMR cluster under the same account.

If one or more components among Hue, Ranger, Oozie, Druid, and Superset are selected, a MetaDB instance will be automatically purchased to store the metadata of such components other than Hive.

To terminate the associated EMR-MetaDB instance, you need to go to the TencentDB console to do so. Hive metadatabases cannot be recovered once terminated.

Make sure that the associated EMR-MetaDB instance is in the same network as the newly created cluster.

**Associate self-built MySQL**: You can associate with your local self-built MySQL database to store Hive metadata, so you don't need to purchase a MetaDB instance separately, which helps reduce your storage costs. Please enter the address beginning with "jdbc:mysql://", username, and login password of the database correctly and make sure that the database can be connected with the cluster.

Make sure that the self-built database is in the same network as the EMR cluster.

Enter the username and password of the database correctly.

If one or more components among Hue, Ranger, Oozie, Druid, and Superset are selected, a MetaDB instance will be automatically purchased to store the metadata of such components other than Hive.

Make sure that the Hive metadata version in the self-built database is not below the Hive version in the new cluster.

## Directions

#### Installing Hive when creating a cluster

1. Log in with your Tencent Cloud account and click Buy Now. Select the Hive component in **Components to Deploy** on the **Software Configuration** page.

Availability Zone	e and Software Co	onfiguration	2.Hard	lware Config	uration	3.Basic Con	figuration	
Billing Mode	Pay-as-you-go	ast China N	lath China D	Youth Asia Davifia				
Region	Guangzhou	Shanghai	Beijing I	Mumbai				
AZ	Guangzhou Zone	3 Guangzho	ou Zone 4					
Cluster Type	HADOOP	DRUID CI	LICKHOUSE					
Product Version	EMD V2.2.0 flipuy							
	EIVIR-V2.2.0.UIIIUX	~						
Required Components	zookeeper 3.5.5	knox 1.2.0	hadoop 2.8	3.5				
Required Components Optional	zookeeper 3.5.5 presto 0.228	knox 1.2.0	hadoop 2.8 spark_hadoo	9.5	sqoop 1.4.7	storm 1.2.3	superset 0.35.2	tez 0.9.2
Required Components Optional Components	zookeeper 3.5.5 presto 0.228 zeppelin 0.8.2	knox 1.2.0 ranger 1.2.0 oozie 5.1.0	hadoop 2.8 spark_hadoo	p2.8 2.4.3 kylin 2.5.2	sqoop 1.4.7 flink 1.9.2	storm 1.2.3 flume 1.9.0	superset 0.35.2 ganglia 3.7.2	tez 0.9.2 hbase 1.4.9
Required Components Optional Components	zookeeper 3.5.5 presto 0.228 zeppelin 0.8.2 hive 2.3.5	xnox 1.2.0 ranger 1.2.0 oozie 5.1.0 udi 0.5.1	spark_hadoop livy 0.7.0 ue 4.6.0	0.5 0p2.8 2.4.3 kylin 2.5.2 1pala 2.10.0	sqoop 1.4.7 flink 1.9.2 alluxio 1.8.1	storm 1.2.3 flume 1.9.0	superset 0.35.2 ganglia 3.7.2	tez 0.9.2 hbase 1.4.9
Required Components Optional Components	zookeeper 3.5.5 presto 0.228 zeppelin 0.8.2 hive 2.3.5 h	knox 1.2.0 ranger 1.2.0 oozie 5.1.0 udi 0.5.1	spark_hadoo livy 0.7.0 ue 4.6.0	1.5 pp2.8 2.4.3 kylin 2.5.2 npala 2.10.0	sqoop 1.4.7 flink 1.9.2 alluxio 1.8.1	storm 1.2.3 flume 1.9.0	superset 0.35.2 ganglia 3.7.2	tez 0.9.2 hbase 1.4.9
Required Components Optional Components Hive Metadatabase	zookeeper 3.5.5 presto 0.228 zeppelin 0.8.2 hive 2.3.5 h Default A	ranger 1.2.0 oozie 5.1.0 udi 0.5.1 ht ssociate EMR-M	spark_hadoon livy 0.7.0 ue 4.6.0 im etaDB Assoc	1.5 pp2.8 2.4.3 kylin 2.5.2 npala 2.10.0 ciate self-built N	sqoop 1.4.7 flink 1.9.2 alluxio 1.8.1	storm 1.2.3 flume 1.9.0	superset 0.35.2 ganglia 3.7.2	tez 0.9.2 hbase 1.4.9
Required Components Optional Components Hive Metadatabase	zookeeper 3.5.5 presto 0.228 zeppelin 0.8.2 hive 2.3.5 h Default A	knox 1.2.0 ranger 1.2.0 oozie 5.1.0 udi 0.5.1 ht	spark_hadoop livy 0.7.0 ue 4.6.0 im etaDB Assoc	1.5 pp2.8 2.4.3 kylin 2.5.2 ppala 2.10.0 ciate self-built N	sqoop 1.4.7 flink 1.9.2 alluxio 1.8.1 MYSQL	storm 1.2.3 flume 1.9.0	superset 0.35.2 ganglia 3.7.2	tez 0.9.2 hbase 1.4.9

2. Select one of the following values for **Hive metadatabase** as needed: **Cluster default**, **Associate EMR-MetaDB**, and **Associate self-built MySQL**.

3. Restrictions vary depending on the selected value. For more information, see **Prerequisites**.

#### Installing Hive after a cluster is created

1. Log in to the EMR console and click **Cluster List** on the left sidebar. On the page that appears, click the **ID/Name** of the target cluster.

2. Choose Cluster Service > Add Component and select the Hive component.

Add Component							
The current cluster has no metadatabase. If you want to add a component such as Hive, Sqoop, Hue, Ranger, Oozie, and Presto, you need to purchase a new TencentDB instance to store metadata.							
Product Version	EMR-V2.2.0.tlinux						
Optional Components	Zookeeper-3.5.5	hadoop-2.8.5	knox-1.2.0	zeppelin-0.8.2			
	livy-0.7.0	hbase-1.4.9	✓ hive-2.3.5	hue-4.6.0			
	✓ oozie-5.1.0	presto-0.228	ranger-1.2.0	spark_hadoop2.8-2.4			
	storm-1.2.3	sqoop-1.4.7	tez-0.9.2	ganglia-3.7.2			
	flume-1.9.0	impala-2.10.0	alluxio-1.8.1	flink-1.9.2			
	kylin-2.5.2	superset-0.35.2	hudi-0.5.1				
Hive Metadatabase	Default Associat	e EMR-MetaDB Associ	ate self-built MYSQL				
Instance ID	Select data 💌						
MetaDB (?)	Pay-as-you-go 1 high IO Te	ncentDB with 4,000 MB memo	ry and 100 GB disk				
Cost							
		Confirm	ncel				

3. Select one of the following values for **Hive metadatabase** as needed: **Cluster default**, **Associate EMR-MetaDB**, and **Associate self-built MySQL**.

4. Restrictions vary depending on the selected value. For more information, see **Prerequisites**.

## Setting Security Groups

Last updated : 2024-01-10 10:02:25

EMR uses Tencent Cloud VPC as the underlying network. Security groups in EMR are used as virtual firewalls to control the access between the internal nodes in a cluster and access from external nodes to internal nodes. This document describes the best practices of using security groups in EMR to help you choose security group policies.

## Security Groups

A security group is a virtual firewall for stateful data packet filtering. As an important network isolation approach provided by Tencent Cloud, it is used to control network access to CVM instances (nodes). When creating an EMR cluster, you can select an existing security group. If there is no existing security group, EMR will automatically create one for you. If the number of security groups has reached the upper limit and you want to create a new one, delete those you no longer use. You can view existing security groups in the VPC console.

### Use Limits and Rules

For use limits and quotas of security groups, see the **Security Group Limits** section in Use Limits Overview. A security group rule consists of: Source: IP address of the source data (inbound) or target data (outbound) Protocol type and protocol port: protocol type such as TCP and UDP Policy: allow or reject access requests

## Rules for Selecting a Security Group

By default, **Select an existing security group** is selected and an EMR security group is selected. You can create a new EMR security group or select a non-EMR security group.

1. For a new EMR security group, port 22, port 30001, and necessary private network IP ranges are enabled. The security group is named in the format of "emr-xxxxxxx\_yyyyMMdd". Do not modify its name.

2. Select an existing security group available in the current region for the current instance. A security group starting with "emr" is recommended, as EMR service has been activated and necessary policies are running properly for such security groups. Security groups not starting with "emr" may lack necessary inbound and outbound rules. This may cause cluster creation failure or cluster unavailability.

3. If you add nodes in a cluster, by default, the nodes inherit the security group policy selected for the cluster when it is created.

## Details of EMR Security Group Policies

If you select a non-EMR security group when creating an EMR cluster, the following inbound and outbound rules must be included. Otherwise, the cluster creation will fail.

#### Inbound rules

Source	Protocol Port	Policy	Note
10.0.0/8	ALL	ACCEPT	Opens IP range A.
172.16.0.0/12	ALL	ACCEPT	Opens IP range B.
192.168.0.0/16	ALL	ACCEPT	Opens IP range C.
0.0.0.0/0	ICMP	ACCEPT	Opens local ICMP.

#### **Outbound rules**

Source	Protocol Port	Policy	Note
0.0.0/0	ALL	ACCEPT	Opens all outbound ports.

#### Inbound rules for accessing webUI

To access the cluster service WebUI normally when using a non-EMR security group, the inbound rules must include the following policies:

Source	Protocol Port	Policy	Note
0.0.0/0	TCP:13000	ACCEPT	Opens port 13000 and port hue.
0.0.0/0	TCP:30001	ACCEPT	Opens port 30001.
0.0.0/0	TCP:30002	ACCEPT	Opens port 30002.
0.0.0/0	TCP:22	ACCEPT	Opens the remote login port.

For more information, see Security Group.

## **Component Configuration Sharing**

Last updated : 2023-12-27 10:14:00

## Overview

You can share the components of an existing cluster with other clusters without deploying such components again. This makes it easier to manage multiple clusters with the same component configurations. Currently supported components include Ranger, ZooKeeper, and Kerberos. Currently supported versions include all EMR versions containing the above components. Dependent: The current cluster is using components of other existing clusters. Shared: One or more components of the current cluster are added to other newly created clusters. Caution

Shared components such as ZooKeeper, Kerberos, and Ranger are no longer deployed in dependent clusters separately, and dependencies cannot be canceled once set. When a cluster with shared components deployed is terminated, all its dependent clusters need to be terminated first. Note that if your account has any overdue payment, dependent clusters won't work properly.

## Limits

1. Your account must have the management permissions of one existing cluster.

2. When Kerberos is enabled, other clusters under the current account must have Kerberos deployed before you can enable the dependency.

3. If a dependent component involves user identity management, you need to manually sync the users in the shared cluster to the new cluster after creating it or adding a user. Note that as the current cluster doesn't deploy Kerberos, krb5 users cannot be created.

## Service Architecture

## Directions

#### **Purchasing a cluster**

- 1. Go to the purchase page and select multiple components under Hadoop.
- 2. Click **Enable** and select the shared cluster.



3. After the new cluster is successfully purchased, relevant components will automatically depend on the selected cluster.

advanced settings	
Kerberos mode 🤅	Enable Kerberos authentication
Dependent Component Pat	E Cable dependent components
Dependent Components 🧃	zookeeper-3.6.3
	zookeeper-3.6.3 dependent c emr-edqx53yb ~

#### Adding a component

- 1. Select Cluster Service > Add Component on the cluster details page in the console
- 2. Select the components to be added.
- 3. Click Enable and select the dependent components and cluster.
- 4. Click **OK** and wait for the components to be successfully added to the cluster.

### **Cluster Information**

If a cluster has dependent or shared components, the component information is displayed in Instance info.

#### Caution

The management features of dependent components can be used only on the **Components** page of the dependedon cluster.

## Account Sync

User accounts need to be created in both the dependent and depended-on clusters.

1. Log in to the depended-on cluster and click Create user to create an account.

2. Log in to the dependent cluster and click **Create user** to create an account with the same username, user group, and password.

3. The dependent cluster will automatically sync the permission and usage information of the account.

emr-edqx53yb EMR-jhaoy7jm	Users				🖾 Help docur
Cluster overview	i The user list shows only	r the user information added via the Users page.			
Instance info					
Cluster services	Create user Delete				Separate keywords with " "; press Enter to separate filter tags
Cluster resources	Username	User group	Remarks	Creation time	Operation
Cluster monitor	root	root	- /	2023-05-17 17:46:42	Reset password Download Keytab
* DashBoard BETA	hadoop	hadoop	-1	2023-05-17 17:46:42	Reset password Download Keytab
<ul> <li>Cluster events</li> <li>Logs</li> </ul>	Total 2 items				Lines per page 10 v H 4 1 /1 page
Cluster inspection					
Alarm records					
<ul> <li>Java analysis</li> </ul>					
Auto-scaling					
Scripts ~					
Users					
Operation logs					

## Managing Cluster Instance Information

Last updated : 2023-12-27 10:26:23

## Overview

Instance information records the basic information of your EMR cluster. You can view the basic configuration, software information, and hardware information of the cluster on the **Instance Information** page.

The **Basic Configuration** section displays basic information of the cluster, such as network information, creation time, security group, high-availability switch status, master public IP, billing mode, host login method, automatic service role definition switch status, and COS switch status.

The **Software Info** section displays the cluster type, use case, product version, deployed component, MetaDB, Hive metadatabase, and Kerberos mode switch status.

The Hardware Configuration section displays the quantity of nodes in each type.

This document describes how to view the cluster instance information in the console.

## Directions

1. After successfully creating a cluster, log in to the EMR console and click the ID/Name of the cluster.

2. If COS is not enabled for the cluster, click Authorize.



	Landar			
	Instance Information			Help Documentation
Basic information  Instance Information Bootstrap Actions Cluster Service	Basic Configuration Instance ID Network information Creation time 2022-04-24 15:21:12	Region Info Guangzhou(Guangzhou Zone 7) Security Group Configure Security Master Public IP	High Av urity Groups Cloud d Jpdates Billing m	ailability Cluster Enabled isk encryption Not enabled sode Pay-as-you-go
Cluster Resource	COS (1) Authorized	Custom Service Role () Not set Set	Host Lo	gin Dassword Set
<ul> <li>Node Status</li> </ul>				
Resource Management	Software Info	Hard	ware Info	
Cluster Monitoring	Cluster Type Hadoop	Maste	er Node 2	
Auto Scaling	Use Cases Hadoop-Default	Core M	Node 4	
Cluster Scripts	Product Version EMR-V2.2.0.tlinux	Comm	non Node 3	
Operation Logs	Components to Deploy hdfs-2.8.5,yarn-2.8.5,zookeeper-3.5.5,kno	x-1.2.0,hive-2.3.5		
	MetaDB cdb-kcx3dajr			
	Hive Metadatabase cdb-kcx3dajr			
	Kerberos mode Not enabled			

÷	Role Manag	ement
	Service Authoriza	tion
	After you agree to gra	Int permissions to EMR, a preset role will be created and relevant permissions will be granted to EMR
	Role Name	EMR_QCSRole
	Role Type Description	Service Role Current role is a EMR service role, which will access your other cloud service resources within the permissions of the associated policies.
	Authorized Policies	Preset policy QcloudAccessForEMRRole①
	Grant Ca	ncel

If you need refined authorization, you can set a custom service role for accessing Tencent Cloud resources during the execution of big data jobs. You should select **Tencent Cloud product service** as the service role type and **Elastic MapReduce** as the service supporting the role.





To change the host login password or method of cluster nodes after scale-out, click **Change**.



<ol> <li>1. Modifica</li> <li>2. You can</li> <li>see <u>Shuttin</u></li> </ol>	tion of the default host login me go to <u>CVM console</u> 🔀 to chang <u>g Down Instances</u> 🔀 to unders	thod applies only to new nodes of the cluster. Ie host passwords/keys of existing nodes of the cluster. Before change, please tand the potential shutdown risks.
Set Password	Associate SSH Key Pair	
Username	root	
New Password *	Enter a password	
	Enter the peepword again	

#### Note

The login method set for initializing nodes in an EMR cluster. After the cluster is created, modification of the setting here applies only to new nodes.

You can go to the CVM console to change host passwords/keys of existing nodes of the cluster. Before change, see **Shutting Down Instances** to understand the potential shutdown risks.

## Node Specification Management

Last updated : 2023-12-27 10:26:39

## Feature Overview

Node specification management is used to set hardware specifications when adding more core, task, or router nodes. Each node type can have up to 3 specifications and will automatically select the current default specification during scaling. If the current default specification is sold out or unavailable, you can set a new default specification. You are recommended to keep the configuration of resources such as CPU cores, memory size, and disk capacity in the new specification the same as that of the existing nodes for optimal compatibility with current component parameters.

## Directions

1. Log in to the EMR Console, select **Cluster List**, and click the ID/Name of the target cluster to enter the cluster details page.

2. Select Cluster Resource > Resource Management > More > Set Node Specification on the cluster details page

← emr-ldjnqr4a v_pjianpan_官网截图集	Resource Management									Help Documentation
Basic Information   Instance Information	The current cluster resources shall prevail. It is not allowed	s were purchased bas I to adjust the fee polic	ed on the official EMR b cy on the CVM Console.	illing rules and the rene	ewal statuses of CVM re	sources are managed by EMR in a unified m	anner. For monthly-sub	scription CVM resources, the	billing statuses :	shown on the current page
Bootstrap Actions	All Nodes Master	Core Comm	on Task R	outer Metadb	Recyle Rener	И				
Cluster Service	Scale Out Scale In	More v							Search IP	Q
Cluster Resource	Node Type	Re Set Node S	Resource Name/	IP	Configuration	Deployment Process	Expiration Time	Creation Time	lag	Operation
Node Status     Resource Management	Master	HOST	en estateata		EMR StandardSA2 CPU: 2-core; memory: 8GB	knox, HiveMetaStore, HiveServer2, HiveW ebHcat, Impala-Catlog, Impala- StateStore, JobHistoryServer NameNode	Pay-as-you-go	2021-06-15 20:01	>	Change Configurati
Cluster Monitoring				(public)	Premium cloud disk: 100G x 1	,Zookeeper,ResourceManager				
Auto Scaling Cluster Scripts Operation Logs	Core	HOST	100-00-020080 No-1402000	N. MILLON (Months)	EMR StandardSA2 CPU: 2-core; memory: 8GB Premium cloud disk: 100G x 1	DataNode,Impala-Server,NodeManager	Pay-as-you-go	2021-06-15 20:01 U	>	Change Configuration Start/Stop Services
	Core	HOST	atte och divergion Nordenskat	NUMBER BOOK	EMR StandardSA2 CPU: 2-core; memory: 8GB Premium cloud disk: 100G x 1	DataNode,Impala-Server,NodeManager	Pay-as-you-go	2021-06-15 20:01	>	Change Configurati Start/Stop Services
	Task	HOST	an a gatta	N. TOL DO LANSING	EMR StandardSA2 CPU: 2-core; memory: 8GB Premium cloud disk: 100G x 1	Impala-Server, NodeManager	Pay-as-you-go	2021-06-15 20:13	5	Change Configurati Start/Stop Services More ▼
	Total 4 items							Lines per page	20 -	< 1 /1 page

3. The default initial specification of the core nodes is the specification selected when the cluster is created. You can set or delete the default one for normal scaling.

- Stencent Cloud
- 4. The task and router nodes have no default specification. Please add the specification as prompted.

Node specification managed default specifications for each of the comparison of the comparameters are compatible parameters are compatible of the comparison	gement is used to set the specific each node type are automatically PU, memory, disk, and other reso le.	ations for core, task, and router node scale selected for scale-out. If there is no quota urces of the new specifications consistent	e-out. Up to 3 node specifications of default specifications, you car with the existing nodes so that th	can be set for each node type. T n set a new default specification. e component configuration
Pay-as-you-go Specific	ation			
Core				
Туре	Model	Disk	Add Time	Operation
EMR StandardSA2	2-core 8 GB	Premium cloud disk 100G * 1	2021-06-15 20:00:53	Set as default Delete
+ New Specification				
Task				
Туре	Model	Disk	Add Time	Operation
EMR StandardSA2	2-core 8 GB	Premium cloud disk 100G * 1	2021-06-15 20:12:47	Set as default Delete
+ New Specification				
Router				
Туре	Model	Disk	Add Time	Operation
+ New Specification				

## Checking and Updating Public IP

Last updated : 2023-12-27 10:26:58

### Feature Overview

The EMR console supports checking and updating the public IP on master1. After you update or add the master1 public IP, the master public IP in the cluster configuration information and the WebUI access address will be updated too.

## Directions

1. Log in to the EMR console and click a **cluster ID** in the cluster list to go to the instance information page. Click **Check for Updates** next to **Master Public IP** in the **Basic Configuration** section.

-			
emr-6dmpunsz Roy-emr-01	Instance Information		Help Documentation
Basic Information	Basic Configuration		
Instance Information	Instance ID ellipse	Region Info Guangzhou(Guangzhou Zone 3)	High Availability Cluster Disabled
<ul> <li>Bootstrap Actions</li> </ul>	Network Info	Security Group	Cloud disk encryption Disabled
Cluster Service	Creation Time 2021-02-23 10:14:38	Master Node Public IP .229 Check for Updates	Billing Mode Pay-as-you-go
Cluster Resource	COS () Authorized	Custom Service Role ① Not set Set	
Cluster Monitoring			
Auto-scaling	Software Info	Hardware Info	
Cluster Scripts			
Operation Logs	Cluster Type HADOOP	Core Node 2	
	Product Version EMR-V2.5.0	Master Node 1	
	Component Info zookeeper-3.6.1,hadoop-2.8.5,knox-1.2.0		
	MetaDB None		
	Hive Metadatabase None		
	Kerberos mode Disabled		

2. If the actual master public IP is different from the public IP saved in EMR, you can update the public IP saved in EMR to the actual master public IP. This page will refresh automatically after synchronization finishes.

Check	and Update Master Public IP	×
0	If the master public IP configured in EMR is not the act IP, you can update the master public IP to the actual p	tual public ublic IP.
Curren	t Public IP	
Actual	Public IP Consistent	
	No need to update Cancel	

## Logging in to Clusters

Last updated : 2023-12-27 09:58:34

## Logging in Through a Remote Login Tool (on Windows)

This section uses Xshell as an example to describe how to log in to an EMR cluster with a password by using a remote login tool on Windows.

#### Applicable OS

Windows

#### Logging in with a password

1. Download PuTTY (a remote login tool on Windows) here and install it.

2. Launch the PuTTY client, enter the following in the PuTTY Configuration window, and click **Open** to create a session as shown below:

**Host Name**: The public IP address of the EMR cluster, which can be viewed on the list page or details page in the EMR Console.

**Port**: The port number of the CVM instance, which has to be 22. Please make sure that the port 22 in the CVM instance is opened. For more information, please see Security Group and Network ACL. **Connection type**: Select **SSH**.

🕵 PuTTY Configurati	on	? ×
Category:		
Session     Logging     Terminal     Keyboard     Bell     Features     Window     Appearance     Behaviour     Translation	Basic options for your PuTTY set     Specify the destination you want to connection type:     O Raw O Telnet O Rlogin O Saved Session     Saved Sessions	ession ect to <u>P</u> ort 22 H O Se <u>r</u> ial
Colours Connection Data Proxy Telnet Rlogin SSH	Test Default Settings	Load Sa <u>v</u> e Delete
Hout Host Keys	Close window on exit: Always Never Only on or elp	clean exit <u>C</u> ancel

3. In the PuTTY session window, enter the obtained admin account and press Enter.

4. Enter the obtained login password and press Enter to log in as shown below:

login as: root root	

## Log in Through SSH (on Linux or macOS)

This section describes how to log in to an EMR cluster through SSH on Linux or macOS.

#### Applicable OS

Linux or Mac OS

#### Logging in with a password

1. On macOS, launch Terminal and run the following command. On Linux, run the following commands directly:

ssh <username>@<hostname or IP address>

username: The admin account, such as root.

hostname or IP address: The public IP address or custom domain name of your EMR cluster.

2. Enter the obtained password (only the input but not the output is displayed here) and press Enter to log in.

## **Cluster Scale-Out**

Last updated : 2025-06-11 16:04:16

## Feature overview

When the compute and storage resources of your EMR cluster are insufficient, you can add more core and task nodes in the console. When the master node is overloaded or lacks sufficient resources, you can scale out the router nodes in the console. Router nodes can share the load of the master node or to serve as task submitters in the cluster. You can scale the cluster at any time.

#### Note:

By default, the current instance specification for scale-out is the instance specification selected when the cluster was created. If the current default specification is unavailable, or if you want to adjust the scale-out configuration, you need to set it in **Node spec** as instructed in Node Specification Management.

By default, the selected component will inherit the cluster-level configuration and fall into the default configuration group for that node type. You can also set the configuration of the target component through the **Specify configuration** parameter.

You cannot specify a configuration group for a ClickHouse cluster to be scaled out.

For non-high availability StarRocks clusters, Master (FE) can be scaled out to 3 nodes, but scale-in is not supported after the scale-out. Please perform a careful evaluation before proceeding.

## Prerequisites

Add router nodes: A router node can be used as a submitter, through which you can submit YARN, Hive, and Spark computing tasks to the cluster. We recommend you select a model with larger memory, preferably not lower than the master node specification.

When the cluster uses monthly subscription billing, the billing mode for scale-out nodes (task nodes and router nodes) supports mixed deployment, that is, both pay-as-you-go billing and monthly subscription billing are supported. Core nodes do not support mixed deployment and only support monthly subscription billing.

For a pay-as-you-go cluster, all newly added nodes support the pay-as-you-go billing mode.

### Directions

1. Log in to the EMR console and click the **ID/Name** of the target cluster in the cluster list to enter the cluster details page.



2. On the cluster details page, select **Cluster Resource** > **Resource Management** > **Scale Out** and select the type of nodes to be added (core, task, or router), billing mode, optional services, and scale-out quantity as needed.

Scale-out Service: The clients of selected services will be deployed to the new node by default.

Specify configuration: Find the target component and select the level from which to inherit the configuration.

If you choose to inherit the configuration of the cluster, an added node will inherit the cluster-level configuration and fall into the default configuration group for that node type.

If you choose to inherit the configuration of the configuration group, an added node will inherit the configuration grouplevel configuration and fall into the selected configuration group.

Deployment Process: Displays the information of service deployment processes after components are selected for the nodes to be added. You can edit the deployment process as needed.

Do not start services after scaling: If this option is selected during scaling, added nodes will not start the service. When needed, you can start the service as instructed in Service Start/Stop.

Tag: Used to identify the added node resources.

Current Specification: The default specification.

You can set the default specification for scale-out as instructed in Node Specification Management.

The default node specification is used for scale-out. You can adjust it by clicking Set Node Specification.

Placement Group: Optional configuration. A placement group is a policy for distributing cloud server instances on the underlying hardware. For details, see Placement Group.

After a placement group is enabled, if the underlying hardware resources are insufficient to achieve node distribution, creation will fail. Operate with caution.

The node specification is configured with a local data disk. If a placement group is not enabled, when multiple local data disks are distributed on the same host, there will be a risk of data loss.

3. After selecting the desired components and number of nodes to be added, click **Confirm** and make the payment. Then, the cluster will start scaling out, which usually takes 10 to 20 minutes.

4. For ClickHouse cluster scale-out, you can add an even number of nodes in a high-availability (HA) instance, and unlimited nodes in a non-HA instance. You can choose an existing cluster or a new one for the nodes to be added. **Note:** 

# There will be no data in the newly added nodes after the virtual ClickHouse cluster is scaled out. You need to migrate the data manually as the system won't automatically do it for you. In order to achieve data balance and improve resource utilization, migrate your data in time after successful cluster scaling.

## Cluster Scale-in

Last updated : 2023-12-27 10:27:35

## Overview

When the resources of your EMR cluster are excessive, you can remove task nodes in the console. When the cluster has router nodes that do not share the load of the master node or that serve as task submitters in the cluster, you can remove nodes to achieve cluster scale-in.

## Preparations

Since core nodes store data, we recommend you not terminate them to avoid data security issues. After a node is terminated, the data stored in it will be deleted. By terminating a node, you confirm that the data in it can be deleted. Pay-as-you-go cluster: Once terminated, the cluster will not be retained in the recycle bin but will be permanently terminated. Please proceed with caution.

#### Caution

Before terminating a cluster, please make sure that your data has been backed up as it cannot be recovered after cluster termination.

## Directions

1. Log in to the EMR console and click the ID/Name of the target cluster in the cluster list.

2. On the **Cluster details** page, select **Cluster resources** > **Resources**, select cloud resources (task or router nodes) that can scale in from the node list. If two or more types of nodes are selected, the **Scale in** button will be unavailable.

1 Termination options	Confirm		
1. If there are EIPs (including	the IPs on the secondary ENI), the server wil	l be retained after returned, and the	e idle IPs will continue to incur fees. If you
don't need to retain it, please	e release it on the corresponding resource m	anagement page.	
2. Please make sure that you	back up the data before terminating the ins	tance. After the instance is terminat	ted, the data cannot be restored.
3. Please ensure that the data terminated. Please (migrate /	a in the terminated nodes has been migrated data] first.	d. If the data hasn't been migrated,	it cannot be recovered after the node is
ode selected.Learn more 💌			
ode selected.Learn more 👻	Resource ID	Node type	Status
ode selected.Learn more 👻	Resource ID	Node type	Status
Node selected.Learn more 👻	Resource ID ins-5lsds90k	Node type Task	Status Terminable
No. Node IP	Resource ID ins-5lsds90k	Node type Task	Status Terminable
No. Node IP 10.0.0.110	Resource ID ins-5lsds90k	Node type Task	Status Terminable
node selected.Learn more ▼ No. Node IP 1 10.0.0.110 aceful scale-in ⑦ ①	Resource ID ins-5lsds90k	Node type Task	Status Terminable
No. Node IP 1 10.0.0.110 aceful scale-in (?)	Resource ID ins-51sds90k Termination documentation 🛂	Node type Task	Status Terminable
No. Node IP 1 10.0.0.110 aceful scale-in (?)	Resource ID ins-5lsds90k Termination documentation	Node type Task	Status Terminable

3. In the pop-up window, confirm the instances to be terminated (removed). If the cluster metaDB is a shared metadatabase, the disk will be retained in this operation and needs to be manually removed in the CDB console.

## **Releasing Now**

Last updated : 2024-10-30 10:37:21

## Feature Introduction

When the cluster is in an isolated status or a node is in a pending-recycle status, the EMR console supports one-click release to prevent ongoing fees for associated products, helping to reduce costs.

#### Note

After immediate release, the cluster or node will be permanently released, and all data will be deleted and cannot be recovered. Proceed with caution.

Associated products include, but are not limited to: cloud data disks, cloud data warehouse MySQL, and public network traffic.

### Directions

#### **Immediate Cluster Release**

#### Note

The cluster is in an isolated status.

1. Log in to the EMR console, and click the corresponding cluster **ID/Name** in the cluster list to enter the cluster details page.

In the cluster list, select the corresponding cluster, choose Management > Release Now, and click Confirm Release in the pop-up window.

#### Immediate Node Release

#### Note

Nodes that are in a pending-recycle status in the recycle bin.

1. Log in to the EMR console, and click the corresponding cluster **ID/Name** in the cluster list to enter the cluster details page.

2. On the cluster details page, select **Cluster Resources > Resource Management > Recycle Bin**, and check the corresponding node.

3. In Operations, select **More > Release Now**, and click Confirm Release in the **Confirm Release** the Selected Node Immediately pop-up window.

## Auto Scaling Overview

Last updated : 2023-12-27 10:27:53

## Overview

You can configure scaling rules in EMR to automatically increase or reduce the computing resources of task nodes as your cluster load changes. This saves costs while quickly responding to changes in the computing needs. Automatic scaling supports two scaling policies: load-based scaling (for YARN-enabled clusters only) and time-based scaling.

### Notes

1. Automatic scaling is disabled by default. It can be either custom scaling or managed scaling. Currently, managed scaling is made available through an allowlist. To use it, submit a ticket for application.

2. Custom scaling can be either load-based or time-based. If you switch the scaling policy, the original scaling rules will be retained. However, they will be in an invalid state and will not be triggered or executed. The added nodes will be retained unless the scale-in rule is triggered.

3. In managed scaling, only **HOST** is supported as the resource type; in custom scaling, **HOST** and **POD** are supported and cannot be used at the same time. If you switch the resource type, the resource specification and instance deployment methods set for the original resource type will be retained. However, they will be in an invalid state and will not be triggered or executed. The added nodes will be retained unless a scale-in rule is triggered. Pod resources are currently made available through an allowlist. To use them, submit a ticket for application.

4. **Pay-as-you-go** and **Spot instances preferred** are supported as the instance deployment policies. However, Pod resources can be deployed only on a pay-as-you-go basis.
# **Custom Scaling Configuration**

Last updated : 2023-12-27 10:28:14

# **Basic Settings**

The **Basic settings** section allows you to set the node quantity range of the custom scaling feature, configure the elastic resource type, and configure whether to support graceful scale-in. It also displays the number of elastic node resources in the current cluster and supports quick release of elastic instances.

to Scaling 🛛 🚺				🖂 Help Docum
Custom scaling Managed	scaling			
<b>plicy Management</b> Sc	aling History			
<ul> <li>Cluster load varies dependi whether the remaining reso and 50 EKS instance(s).</li> </ul>	ing on compute tasks. Auto Scaling a ource quota is enough. If the quota is	llows you to set policies to automatically increase or de not enough, apply for a larger one in advance. In the o	crease task nodes based on cluster load to meet the changing bus uurrent account, for Guangzhou Zone 7, the remaining quota incluc	iness needs. When Auto Scaling is enabled, pay attention t les 500 pay-as-you-go CVM instance(s), 100 spot instance(
Basic Settings				
Elastic node count: 1 Release All				
Min node count 1				
Max node count 20				
Resource type HOST				
Global graceful scale-in 🛈 🛛 Dis	able			
Scaling Specification Mana	gement(i)			Add Specific
Node selection policy() P	'ay-as-you-go <mark>Modify</mark>			
Model	AZ	Hardware Configuration	Priority	Operation
EMR StandardS5	Guangzhou Zone 7	4-core 8 GB System disk:1 x 50GB高效云盘 Data Disk:1 x 200GB高效云盘	1	Delete

Min node count: The minimum number of elastic task nodes to be retained in the cluster after the auto scale-in policy is triggered.

Max node count: The maximum number of elastic task nodes to be retained in the cluster after the auto scale-out policy is triggered. The total number of nodes to be added based on one or more types of specifications cannot exceed the max node count.

Edit basic settings		×
Min node count	— 1 + (0 - 9999)	
Max node count	- 20 + (1 - 9999)	
Resource type	O HOST	
	Switch to POD if resources are in shortage	
Global graceful scale-in 🚯		
	Confirm	

Release all: Clears all nodes added in an auto-scaling action with one click.

Release spot instances: Clears all spot instance nodes added in an auto-scaling action with one click.

Release pay-as-you-go instances: Clears all pay-as-you-go nodes added in an auto-scaling action with one click.

Release elastic node	×
Total elastic node count: 1	
O Pay-as-you-go node count: 1	
Release 0 spot instances	
C Release All	
Confirm Cancel	

Global graceful scale-in: This feature is disabled by default. It applies to all scale-in rules and can be disabled for individual scale-in rules.

Resource type: HOST resources are billed on a pay-as-you-go or spot basis, whereas POD resources are billed on a pay-as-you-go basis and can only be used to assume the NodeManager role of YARN.

Caution

When you switch the resource type, the scaling specifications and node selection policy of the new resource type apply.

## Scaling Specifications Management

Scaling specifications refer to node specifications for custom scaling. You can configure each cluster with up to five scaling specifications. When a scale-out rule is triggered, the scaling will be carried out according to the specification priority. When the high-priority specification has fewer resources than the number of resources to be added, the specification of the next priority will be used for supplement. In order to maintain the linear change of the cluster load, we recommend you keep the CPU and memory of the scaling specifications consistent.

Node selection policy: It can be **Pay-as-you-go** or **Spot instances preferred**. **Pay-as-you-go**: Only pay-as-you-go nodes are added to provide the required computing power when a scale-out rule is triggered.

**Spot instances preferred**: Spot instances will be preferred if they are available to provide the required computing power when a scale-out rule is triggered. **Min proportion of pay-as-you-go nodes**: The minimum proportion of pay-as-you-go nodes to the scale-out quantity.

#### Example:

If 10 nodes are to be added and the **Min proportion of pay-as-you-go nodes** is 20%, at least two pay-as-you-go nodes will be added when a scale-out rule is triggered, with the remaining eight being spot instance nodes. When spot instance nodes are insufficient, additional pay-as-you-go nodes will be added to make up the difference. You can add, delete, change, and query the nodes in the scaling specifications, and adjust the priority of scaling specifications as needed.

The sequence of the five specifications for scale-out is as follows (same for pay-as-you-go and spot instances): When the resources are sufficient: 1 > 2 > 3 > 4 > 5.

#### Example:

If five preset specifications are available and the resources are sufficient, when a scale-out rule is triggered to add 10 nodes, 10 nodes with specification 1 will be added according to the sequence, and other preset specifications will not be selected.

When the resources are insufficient: 1 + 2 > 1 + 2 + 3 > 1 + 2 + 3 + 4 > 1 + 2 + 3 + 4 + 5.

#### Example:

If there are eight nodes with specification 1, four with specification 2, and three with specification 3, when a scale-out rule is triggered to add 13 nodes, eight nodes with specification 1, four with specification 2, and one with specification 3 will be added following the sequence.

When a resource specification is out of stock (take specification 2 as an example): 1 + 3 > 1 + 3 + 4 > 1 + 3 + 4 + 5. **Example:** 

If there are eight nodes with specification 1, no nodes with specification 2 (out of stock), and three nodes with specification 3, when a scale-out rule is triggered to add 10 nodes, eight nodes with specification 1 and two nodes with specification 3 will be added according to the sequence, while specification 2 will not be selected.

If there are eight nodes with specification 1 and nodes with other preset specifications are all out of stock, when a scale-out rule is triggered to add 10 nodes, eight nodes with specification 1 will be added, with the scale-out being partially successful.

Scaling Specification Management()					
Node selection policy ()	Pay-as-you-go Modify				
Model	AZ	Hardware Configuration	Priority	Operation	
EMR StandardS5	Guangzhou Zone 7	4-core 8 GB System disk:1 x 50GB高效云盘 Data Disk:1 x 200GB高效云盘	1	Delete	

## Scaling Rule Management

Scaling rules are business policies used to configure the trigger conditions of scaling actions and the number of nodes to be added or removed. Automatic scaling supports two types of policies: load-based scaling and time-based scaling. You can choose to use either type, but not both at the same time. When you switch the policies, the original scaling rules will be retained in an invalid state and will not be triggered. Those already added nodes will also be retained unless a scale-in rule is triggered. Each policy can be configured with up to 10 scaling rules. When two rules are triggered at the same time, the rule of the higher priority will be executed first.



To set a scaling rule, select Load-based scaling or Time-based scaling for Auto-scaling type in the Scaling rule management pane and click Add.

#### Load-based scaling

When it is impossible to accurately estimate the peaks and valleys of cluster computing, you can configure the scaling policy based on the load to ensure that important jobs are completed on time. The load is mainly based on the preset YARN metric statistics rules and the task nodes are automatically adjusted when the preset conditions are triggered.

Scaling Rule Manage	Scaling Rule Management①						
Auto Scaling Type 🛈	O Load-based scaling O Tim	e-based scaling					Add Ru
Rule Type	Rule Name	Rule Details		Rule Status	Priority	Operation	
			No data				

To add a load-based scaling rule, select **Load-based scaling** as the **Auto-scaling type**, click **Add**, and configure the following fields:

Rule status: Used to mark whether the rule is enabled. The default status of a rule is "Enabled". When you don't want a rule to be executed but still want to retain it, you can set the rule status to "Disabled".

Rule name: The name of the scaling rule. The scaling rule names in the same cluster must be unique (including scaleout and scale-in rules).

Validity: The load-based scaling rule can be triggered only within the validity. **Unlimited** is selected by default, and you can customize this field.

Target service: By default, the selected component will inherit the cluster-level configuration and fall into the default configuration group for that node type. You can also set the configuration of the target component through the **Specify configuration** parameter.

Node label: This field is empty by default, and the added nodes will be assigned the default label. If you have set a label, resources will be added to the node with this label.

Load metric: YARN load metric. You can set the condition rule for triggering the threshold based on the load metric selected here.

Category	Dimension	EMR Auto-Scaling Metric	Description
AvailableVCores	root	AvailableVCores#root	Number of virtual cores available in the root queue
	root.default	AvailableVCores#root.default	Number of virtual cores available in the root.default queue
	Custom subqueue	Example: AvailableVCores#root.test	Number of virtual cores available in the root.test queue



PendingVCores	root	PendingVCores#root	Number of virtual cores waiting to be available in the root queue
	root.default	PendingVCores#root.default	Number of virtual cores waiting to be available in the root.default queue
	Custom subqueue	Example: PendingVCores#root.test	Number of virtual cores waiting to be available in the root.test queue
	root	AvailableMB#root	Amount of memory available in the root queue, in MB
AvailableMB	root.default	AvailableMB#root.default	Amount of memory available in the root.default queue, in MB
	Custom subqueue	Example: AvailableMB#root.test	Amount of memory available in the root.test queue, in MB
	root	PendingMB#root	Amount of memory waiting to be available in the root queue, in MB
PendingMB	root.default	PendingMB#root.default	Amount of memory waiting to be available in the root.default queue, in MB
	Custom subqueue	Example: PendingMB#root.test	Amount of memory waiting to be available in the root.test queue, in MB
AvailableMemPercentage	Clusters	AvailableMemPercentage	Available memory in percentages
ContainerPendingRatio	Clusters	ContainerPendingRatio	Ratio of the number of containers to be allocated to the number of allocated containers
AppsRunning	root	AppsRunning#root	Number of running tasks in the root queue
	root.default	AppsRunning#root.default	Number of running tasks in

🔗 Tencent Cloud

			the root.default queue
	Custom subqueue	Example: AppsRunning#root.test	Number of running tasks in the root.test queue
	root	AppsPending#root	Number of pending tasks in the root queue
AppsPending	root.default	AppsPending#root.default	Number of pending tasks in the root.default queue
	Custom subqueue	Example: AppsPending#root.test	Number of pending tasks in the root.test queue
	root	PendingContainers#root	Number of containers to be allocated in the root queue
PendingContainers	root.default	PendingContainers#root.default	Number of containers to be allocated in the root.default queue
	Custom subqueue	Example: PendingContainers#root.test	Number of containers to be allocated in the root.test queue
	root	AllocatedMB#root	Amount of memory allocated in the root queue
AllocatedMB	root.default	AllocatedMB#root.default	Amount of memory allocated in the root.default queue
	Custom subqueue	Example: AllocatedMB#root.test	Amount of memory allocated in the root.test queue
	root	AllocatedVCores#root	Number of virtual cores allocated in the root queue
AllocatedMB	root.default	AllocatedVCores#root.default	Number of virtual cores allocated in the root.default queue
	Custom subqueue	Example: AllocatedVCores#root.test	Number of virtual cores allocated in the root.test queue

🔗 Tencent Cloud

ReservedVCores	root	ReservedVCores#root	Number of virtual cores reserved in the root queue
	root.default	ReservedVCores#root.default	Number of virtual cores reserved in the root.default queue
	Custom subqueue	Example: ReservedVCores#root.test	Number of virtual cores reserved in the root.test queue
	root	AllocatedContainers#root	Number of containers allocated in the root queue
AllocatedContainers	root.default	AllocatedContainers#root.default	Number of containers allocated in the root.default queue
	Custom subqueue	Example: AllocatedContainers#root.test	Number of containers allocated in the root.test queue
	root	ReservedMB#root	Amount of memory reserved in the root queue
ReservedMB	root.default	ReservedMB#root.default	Amount of memory reserved in the root.default queue
	Custom subqueue	Example: ReservedMB#root.test	Amount of memory reserved in the root.test queue
	root	AppsKilled#root	Number of tasks terminated in the root queue
AppsKilled	root.default	AppsKilled#root.default	Number of tasks terminated in the root.default queue
	Custom subqueue	Example: AppsKilled#root.test	Number of tasks terminated in the root.test queue
AppsFailed	root	AppsFailed#root	Number of tasks failed in the root queue



	root.default	AppsFailed#root.default	Number of tasks failed in the root.default queue
	Custom subqueue	Example: AppsFailed#root.test	Number of tasks failed in the root.test queue
	root	AppsCompleted#root	Number of tasks completed in the root queue
AppsCompleted	root.default	AppsCompleted#root.default	Number of tasks completed in the root.default queue
	Custom subqueue	Example: AppsCompleted#root.test	Number of tasks completed in the root.test queue
	root	AppsSubmitted#root	Number of tasks submitted in the root queue
AppsSubmitted	root.default	AppsSubmitted#root.default	Number of tasks submitted in the root.default queue
	Custom subqueue	Example: AppsSubmitted#root.test	Number of tasks submitted in the root.test queue
AvailableVCoresPercentage	Clusters	AvailableVCoresPercentage	Percentage of virtual cores available in the cluster
	root	MemPendingRatio#root	Percentage of memory waiting to be available in the root queue
MemPendingRatio	root.default	MemPendingRatio#root.default	Percentage of memory waiting to be available in the root.default queue
	Custom subqueue	Example: MemPendingRatio#root.test	Percentage of memory waiting to be available in the root.default queue

Statistical rule: The cluster load metric selected by the user is triggered once when the threshold is reached according to the selected aggregate dimension (average value) within a specific statistical period.

Statistical period: The statistical duration of the metric. Currently, three statistical periods are supported: 300s, 600s, and 900s.

Repeat count: The number of times that the threshold of the aggregated load metric is reached. When the repeat count is reached, the cluster will be automatically scaled.

Scale-out mode: You can select **Node**, **Memory**, or **Core**, and their values must be a positive integer.

In case of scale-out in the core or memory mode, the quantity of nodes guaranteeing the maximum computing power will be added.

#### Example:

1. In the core mode, if 10 cores are to be added, but the specification for auto-scaling is to add 8-core nodes in order of priority, then **two 8-core nodes** will be added when the rule is triggered.

2. In the memory mode, if 20 GB memory is to be added, but the specification for auto-scaling is to add 16 GB memory nodes in order of priority, then **two 16 GB MEM nodes** will be added when the rule is triggered.

Rule Type	O Scale Out ○ Scale In
Rule Name (j)	
	1-64 characters; supports Chinese characters, letters, digits, -, and _
Valid Time 🚯	Unlimited 💌
Scale-out Service	HDFS-2.8.5 YARN-2.8.5
	Specify configuration The component inherits the cluster-level configuration by default. To adjust the configuration, you can specify a configuration group.
Deployment Process	NodeManager
	A deployment process refers to the process information of node deployment.
Node Label (j)	
	The name of node label can only contain numbers, letters, - or It cannot start with - or _ and up to 64 characters are supported.
Load Metric 🛈	AvailableVCores#root 💌
Statistical Rule 🛈	Statistical period 300Second 🔻 > 💌
Repeat Count 🚯	1 •
Scale-out mode 🛈	Node 🔻
Resource supplement retry 🕄	Node Core
Cooldown Period 🛈	Memory
	Cool-down time range: 0-43200 seconds

Scale-in mode: You can select **Node**, **Memory**, or **Core**, and their values must be a positive integer.

In case of scale-in in the core or memory mode, a minimum quantity of nodes will be released in reverse chronological order while keeping the business running properly, with at least one node released.

Example:



1. In the core mode, if 20 cores are to be released, and the cluster has three 8-core 16 GB MEM and two 4-core 8 GB MEM elastic nodes in reverse chronological order when a scale-in rule is triggered, then **two 8-core 16 GB MEM nodes** will be released.

2. In the memory mode, if 30 GB memory is to be released, and the cluster has three 8-core 16 GB MEM and two 4core 8 GB MEM elastic nodes in reverse chronological order when a scale-in rule is triggered, then **one 8-core 16 GB MEM node** will be released.

Rule Type	O Scale Out ○ Scale In
Rule Name 🚯	
	1-64 characters; supports Chinese characters, letters, digits, -, and _
Valid Time (j	Unlimited 💌
Scale-out Service	HDFS-2.8.5 YARN-2.8.5
	Specify configuration The component inherits the cluster-level configuration by default. To adjust the configuration, you can specify a configuration group.
Deployment Process	NodeManager A deployment process refers to the process information of node deployment.
Node Label 访	
	The name of node label can only contain numbers, letters, - or It cannot start with - or _ and up to 64 characters are supported.
Load Metric 🚯	AvailableVCores#root *
Statistical Rule	Statistical period 300Second 🔻 > 💌
Repeat Count 🚯	1 *
Scale-out mode 🛈	Node 🔻
Resource supplement retry 🛈	Node Core
Cooldown Period	Memory
	Cool-down time range: 0-43200 seconds

Resource supplement retry: When auto scale-out is performed during peak hours, the number of nodes added may be less than the target number due to the lack of resources. In such a case, if the resource supplement retry policy is enabled, the system will automatically retry to request resources when there are sufficient resources of the set scaling specifications until the target number is reached or approximated. If frequently encountering this situation, you can

enable this feature. Please note that retry may extend the time of auto scale-out and this policy change may affect your business.

Cooldown period: The interval (60s to 43200s) before carrying out the next auto-scaling action after the rule is successfully executed.

Graceful scale-in: After this feature is enabled, if the scale-in action is triggered when a node is executing tasks, the node will not be released immediately. Instead, the scale-in action will be executed after the tasks are completed. However, the scale-in action will still be executed if the tasks are not completed within the specified time.

Create Rule	
<ol> <li>1. If there are consistent ar due to the in 2. When the</li> </ol>	e multiple preset resource specifications, we recommend you keep CPUs and memory sizes for different specifications ad match them with the resource occupation configuration of the service, so as to avoid the service process start failure consistency between node hardware configuration and service configuration. scale-in rule is triggered, nodes will be removed in reverse chronological order of their creation time.
Rule Type	◯ Scale Out   ◯ Scale In
Rule Name 🛈	
	1-64 characters; supports Chinese characters, letters, digits, -, and _
Valid Time (	Unlimited 💌
Load Metric	AvailableVCores#root *
Statistical Rule 🧊	Statistical period 300Second 🔻 > 🔻
Repeat Count	1 *
Scale-in mode 🛈	Node 💌
Cooldown Period (	
	Cool-down time range: 0-43200 seconds
Graceful Scale-In 访	Value range: 60-1800 seconds
	Confirm Cancel

#### **Time-based scaling**

If there are obvious peaks and valleys in your cluster computing workload during a certain period, you can configure time-based scaling policies to ensure that important jobs are completed on time. Time-based scaling policies allows you to automatically add or remove task nodes daily, weekly, or monthly during a set time period.

Scaling Rule Manage	ement(j)					
Auto Scaling Type 🛈	O Load-based scaling	O Time-based scaling				Add Ru
Rule Type	Rule Name	Rule Details	Rule Status	Priority	Operation	
Scale Out	test-11	Quantity: 1 Cooldown Period : 0 seconds Execution Type : Once Execution Time : 2023-04-25 17:30:00 Retry Time after Expiration : 0 seconds Scheduled termination : Unlimited		1	Edit Delete	

To add a time-based scaling rule, select **Time-based scaling** as the **Auto-scaling type**, click **Add**, and configure the following fields:

Rule name: The name of the scaling rule. The scaling rule names in the same cluster must be unique (including scaleout and scale-in rules).

Target service: By default, the selected component will inherit the cluster-level configuration and fall into the default configuration group for that node type. You can also set the configuration of the target component through the **Specify configuration** parameter.

Once: Triggers a scaling action once at a specific time, accurate to the minute.

Recurring: Triggers a scaling action daily, weekly, or monthly at a specific time or time period.

Retry period: Auto-scaling may not be executed for various reasons at the specified time. After you set the retry period, the system will try to execute the scaling at a certain interval within the period until it is executed when the conditions are met.

Rule expiration: The time-based scaling rule is valid until the end of the date specified here.

Scale-out mode: You can select Node, Memory, or Core, and their values must be an integer other than 0.

In case of scale-out in the core or memory mode, the quantity of nodes guaranteeing the maximum computing power will be added.

#### **Example:**

1. In the core mode, if 10 cores are to be added, but the specification for auto-scaling is to add 8-core nodes in order of priority, then **two 8-core nodes** will be added when the rule is triggered.

2. In the memory mode, if 20 GB memory is to be added, but the specification for auto-scaling is to add 16 GB memory nodes in order of priority, then **two 16 GB MEM nodes** will be added when the rule is triggered.

Rule Type	O Scale Out ○ Scale In
Rule Name 🛈	
	1-64 characters; supports Chinese characters, letters, digits, -, and _
Scale-out Service	HDFS-2.8.5 YARN-2.8.5
	Specify configuration The component inherits the cluster-level configuration by default. To adjust the
	configuration, you can specify a configuration group.
Deployment Process	NodeManager
	A deployment process refers to the process information of node deployment.
Node Label (i)	
	The name of node label can only contain numbers, letters, - or It cannot start with - or _ and up to 64 character are supported.
Execution Type	Once Recurring
Execution Time	2023-04-25 🗊 17:43
Retry Time after Expiration 🚯	
	The retry time after expiration should be between 0 and 21600 seconds
Scale-out mode	Node 💌
Resource supplement retry (j)	Node Enable
Cooldown Period 访	Memory
	Cool-down time range: 0-43200 seconds

Scale-in mode: You can select **Node**, **Memory**, or **Core**, and their values must be an integer other than 0.

In case of scale-in in the core or memory mode, a minimum quantity of nodes will be released in reverse chronological order while keeping the business running properly, with at least one node released.

Example:



1. In the core mode, if 20 cores are to be released, and the cluster has three 8-core 16 GB MEM and two 4-core 8 GB MEM elastic nodes in reverse chronological order when a scale-in rule is triggered, then **two 8-core 16 GB MEM nodes** will be released.

2. In the memory mode, if 30 GB memory is to be released, and the cluster has three 8-core 16 GB MEM and two 4core 8 GB MEM elastic nodes in reverse chronological order when a scale-in rule is triggered, then **one 8-core 16 GB MEM node** will be released.



consistent and match t due to the inconsisten 2. When the scale-in ru	them with the resource occupation configuration of the service, so as to avoid the service process start failure cy between node hardware configuration and service configuration. Jule is triggered, nodes will be removed in reverse chronological order of their creation time.
Rule Type	🔾 Scale Out 🛛 O Scale In
Rule Name 🛈	
	1-64 characters; supports Chinese characters, letters, digits, -, and _
Execution Type	Once Recurring
Execution Time	
	2023-04-25 🖽 17:43
Retry Time after Expiration (	
	The retry time after expiration should be between 0 and 21600 seconds
Scale-in mode 🛈	Node 💌
Cooldown Period 🛈	Node
Ŭ	Core nge: 0-43200 seconds
	Memory

Resource supplement retry: When auto scale-out is performed during peak hours, the number of nodes added may be less than the target number due to the lack of resources. In such a case, if the resource supplement retry policy is enabled, the system will automatically retry to request resources when there are sufficient resources of the set scaling specifications until the target number is reached or approximated. If frequently encountering this situation, you can enable this feature. Please note that retry may extend the time of auto scale-out and this policy change may affect your business.



Cooldown period: The interval (60s to 43200s) before carrying out the next auto-scaling action after the rule is successfully executed.

Scheduled termination: You can specify the usage duration for the added nodes, so that the nodes will not be affected when a scale-in rule is triggered. **Unlimited** is selected by default. You can enter an integer between 1–24 (in hours) here.

#### Use cases:

This feature is suitable when you need more computing resources for a specified period of time (within one day) and the resources should not be affected by other scale-in rules.

Resource supplement retry (j)	O Disable O Enable
Cooldown Period (j)	
	Cool-d It must be an integer between 1 and 24
Scheduled termination (i)	Custom 🔻 hour
Tag ③	+ New Tag
	Up to <mark>5</mark> tags can be bound
	Confirm

Rule status: Used to mark whether the rule is enabled. The default status of a rule is "Enabled". When you don't want a rule to be executed but still want to retain it, you can set the rule status to "Disabled".

Graceful scale-in: After this feature is enabled, if the scale-in action is triggered when a node is executing tasks, the node will not be released immediately. Instead, the scale-in action will be executed after the tasks are completed. However, the scale-in action will still be executed if the tasks are not completed within the specified time.

Create Rule	e preset resource specifications, we recommend you keep CPUs and memory sizes for different specifications
consistent and match due to the inconsister 2. When the scale-in r	them with the resource occupation configuration of the service, so as to avoid the service process start failure ncy between node hardware configuration and service configuration. ule is triggered, nodes will be removed in reverse chronological order of their creation time.
Rule Type	◯ Scale Out O Scale In
Rule Name	1-64 characters; supports Chinese characters, letters, digits, -, and _
Execution Type	Once Recurring
Execution Time	2023-04-25 🗊 17:43
Retry Time after Expiration (i)	The retry time after expiration should be between 0 and 21600 seconds
Scale-in mode 🛈	Node 🔻
Cooldown Period	Cool-down time range: 0-43200 seconds
Graceful Scale-In	Value range: 60-1800 seconds
	Confirm

# Managed Scaling Configuration

Last updated : 2023-12-27 10:28:49

After the managed scaling feature is enabled, the system will continuously monitor the load of the YARN cluster and calculate the peak fluctuations in the last 10 minutes, so as to automatically add or remove task nodes. Managed scaling applies only to clusters containing the YARN component.

# **Basic Settings**

The **Basic settings** section allows you to set the node quantity range of the automatic scaling feature and the minimum number of pay-as-you-go nodes.

Min number of nodes: The minimum number of task nodes a cluster can have when the managed scale-in policy is triggered.

Max number of nodes: The maximum number of task nodes a cluster can have when the managed scale-out policy is triggered. The accumulated number of task nodes under one or more rules cannot exceed the maximum number of nodes.

Min number of pay-as-you-go nodes: The minimum number of pay-as-you-go nodes to be added after a scale-out is triggered, which is used to set the proportion of pay-as-you-go nodes and spot instances. By default, the value is the maximum number of nodes.

#### For example:

If the minimum number of nodes is set to 0, the maximum number of nodes to 100, and the minimum number of pay-as-you-go nodes to 10, at least 10 pay-as-you-go nodes will be added when a scale-out is triggered, with the remaining ones being spot instances. When spot instances are insufficient, additional pay-as-you-go nodes will be added to make up the difference.

### Scaling Specifications Management

Scaling specifications refer to node specifications for managed scaling. You can configure each cluster with up to five scaling specifications. When a scale-out rule is triggered, the scaling will be carried out according to the specification priority. When the high-priority specification has fewer resources than the number of resources to be added, the specification of the next priority will be used for supplement. In order to maintain the linear change of the cluster load, we recommend you keep the CPU and memory of the scaling specifications consistent. Managed scaling is supported only for host resources.

You can add, delete, change, and query the nodes in the scaling specifications, and adjust the priority of scaling specifications as needed.

The sequence of the five specifications for scale-out is as follows (same for pay-as-you-go and spot instances):

When the resources are sufficient: 1 > 2 > 3 > 4 > 5

#### For example:

If five preset specifications are available and the resources are sufficient, when a scale-out rule is triggered to add 10 nodes, 10 nodes with specification 1 will be added according to the sequence, and other preset specifications will not be selected.

When the resources are insufficient: 1 + 2 > 1 + 2 + 3 > 1 + 2 + 3 + 4 > 1 + 2 + 3 + 4 + 5.

#### For example:

If there are eight nodes with specification 1, four with specification 2, and three with specification 3, when a scale-out rule is triggered to add 13 nodes, eight nodes with specification 1, four with specification 2, and one with specification 3 will be added following the sequence.

When a resource specification is out of stock (take specification 2 as an example): 1 + 3 > 1 + 3 + 4 > 1 + 3 + 4 + 5.

#### For example:

If there are eight nodes with specification 1, no nodes with specification 2 (out of stock), and three nodes with specification 3, when a scale-out rule is triggered to add 10 nodes, eight nodes with specification 1 and two nodes with specification 3 will be added according to the sequence, while specification 2 will not be selected.

If there are eight nodes with specification 1 and nodes with other preset specifications are all out of stock, when a scale-out rule is triggered to add 10 nodes, eight nodes with specification 1 will be added, with the scale-out being partially successful.

## Managed scaling monitoring metrics

Managed scaling monitors many metrics and calculates the suggested number of nodes under each metric. Then, it determines on a scale-out or scale-in based on the number of nodes.

Managed Scaling Monitoring Metric	Description
AvailableMemPercentage	Available memory in percentages
AvailableVCoresPercentage	Proportion of virtual cores available in YARN in percentages

Statistical Rule: The cluster load metric is set to process the peak loads in the last 10 minutes based on the specified aggregation dimension within a statistical period.

Statistical Period: The statistical duration of the metric, which is 1 minute.

By default, managed scaling applies the principle of quick scale-out and cautious and graceful scale-in.

# **Viewing Scaling Records**

Last updated : 2025-03-05 11:00:39

The records of auto scaling actions can be viewed in **Scaling history**. Grading of auto scaling events is supported, and event alarm policies are set based on the event level. For more information about event levels, see Cluster Events. For more information about event alarm configuration, see Alarm Configurations.

# **Custom Scaling Records**

Supports filtering the scaling records by execution time range and searching by policy name.

Sorts by **Execution time**, and displays **Execution time**, **Policy name**, **Scaling type**, and **Execution status**. You can click **Details** in the **Operation** column to view details.

There are four auto scaling execution statuses:

Executing: The automatic scaling is being executed.

Successful: All nodes are added to or removed from the cluster.

Partially successful: Some nodes are successfully added or removed from the cluster, while others fail due to disk quota management or CVM inventory.

Failed: No node is added to or removed from the cluster.

**Resource supplement retry** displays whether this feature is enabled. If it is enabled, the retries will be displayed. **Elastic node count** displays the details of the execution result. If the execution fails, it will display the cause of the failure as well as the solution.

**Scaling Specification** displays the instance specification and the number of added or removed nodes after the rule is triggered.

## Managed scaling records

Supports filtering the scaling records by **Execution Time** or **Scaling Type**.

Supports sorting by **Execution Time**, **Scaling Type**, **Model Specification**, **Quantity**, **Execution Status**, or **Reason**.

There are two execution statuses of managed scaling:

Successful: The target node is added to or removed from the cluster based on the cluster load.

Failed: The target node failed to be added to the cluster based on the cluster load due to resource insufficiency. We recommend you modify the preset resource specification.

Model Specification: The target model specification and type after the rule is triggered for a scale-out or scale-in. Quantity: The number of added or removed nodes in various specifications after the operation.

Custom scaling Manag Policy Management S	ed scaling				
Today Yesterday Execution Time	Last 7 days Last 30 days Scaling Type	2023-05-16 19:51:20 to 2023-05-23 19:51:20	Quantity	Execution Status	Reason
			No data vet		
			no data jer		

# Graceful Scale-In

Last updated : 2025-07-03 11:35:04

## Overview

After graceful scale-in is enabled, if the scale-in action is triggered when a node is executing tasks, the node will not be released immediately. Instead, it will be released after completing the tasks. Graceful scale-in is available for both automatic scaling and manual scale-in.

#### Note:

YARN, HBase, and Presto (renamed Trino in EMR v2.7.0 and v3.40 or later) support graceful scale-in. However, Presto cannot be scaled in gracefully in Ranger, Kerberos, and OpenLDAP integration scenarios.

### Directions

#### Auto-scaling

Auto-scaling allows you to enable or disable graceful scale-in for all scale-in rules. This feature is called global graceful scale-in and is disabled by default. When you add or edit a single scale-in rule, graceful scale-in is enabled by default, and the default duration is 60 seconds (valid range: 60–86400 seconds).

#### Note:

When both graceful scale-in and a single scale-in rule are enabled, the graceful scale-in takes effect.

1. Log in to the EMR console, click the **ID**/name of the target cluster to enter the cluster details page, and click **Auto-scaling**.

2. In the Scaling rule management section on the Auto-scaling page, click Add rule and add a scale-in rule.

#### Manual scale-in

When you try to manually remove a node, graceful scale-in is disabled by default. When you enable this feature, the default duration is 60 seconds (valid range: 60–86400 seconds).

1. Click the **ID**/name of the target cluster to enter the cluster details page. Then, select **Cluster resources** >

#### Resources.

2. Select the node to remove and click **Scale in**. In this case, graceful scale-in is disabled by default. You can enable it and set a duration.

3. After finishing your settings, click Next, confirm the node information, and click Start termination.

# **Changing Configurations**

Last updated : 2023-12-27 11:06:56

# Overview

In actual use, you may find that the configurations of the nodes in your cluster need to be upgraded, especially when the CPU or memory resources of the master nodes are insufficient. This document describes how to change the instance configuration in the EMR console.

#### Caution

The node will be shut down during the change. Note that the shutdown may affect the normal use of the cluster and even interrupt your business. Proceed with caution.

The size of data disks and system disks cannot be changed.

## Prerequisites

1. After the configurations of pay-as-you-go nodes are changed, the billing tier will restart from the first tier. For monthly subscribed clusters, you need to make up the difference.

2. The configurations of local disk models, Pod resources, and spot instance models cannot be changed.

3. If you batch adjust configurations, the system will automatically deduct fees one by one. Make sure your account has sufficient balance.

4. The refund will be credited to your Tencent Cloud account at the ratio of cash to trial credit paid upon purchase, but the discount amount or voucher (if any) will not be refunded.

### Directions

1. Log in to the EMR console and click the **ID/Name** of the target cluster in the cluster list to enter the cluster details page.

2. Select **Cluster Resource** on the cluster details page to enter the **Resource Management** page. Select the target node and **change the configuration**. Batch change is supported, but you can only change the configurations of nodes in the same billing mode to the same configuration.

All nodes	Master	Core	Common	Task	Router	Metadb	Recyle	Renew				
Scale Out	Scale In ords with " ") pre	More Cneck	т овк update ode Specification	n		•						
Node ty	/pe	R Chang	ge from pay-as-y			tion		(	Configuration	Creation time 1	Expiration time	Operation
✓ Master		Edit ta Set au H Chang	ag ito-renewal ge Configuration	_		Ţ			:MR StandardSA2 :PU: 4-core; memory: 16GB iystem disk: Cloud SSD 50G x 1 Data Disk: Cloud SSD 200G x 1	2022-10-25 15:58:23	Pay-as-you-go	Change Configuration
✓ Core		HOST					)		:MR StandardSA2 IPU: 4-core; memory: 8GB iystem disk: Cloud SSD 50G x 1 Data Disk: Cloud SSD 200G x 1	2022-10-25 15:58:25	Pay-as-you-go	Change Configuration
✔ Core		HOST						8 ( 5	:MR StandardSA2 IPU: 4-core; memory: 8GB iystem disk: Cloud SSD 50G x 1 Data Disk: Cloud SSD 200G x 1	2022-10-25 15:58:27	Pay-as-you-go	Change Configuration
Total 3 items											Lines per page 20 🔻	H

3. On the configuration adjustment page, confirm the relevant information, read carefully the notes, and select **I have** read and understood the notes and agree to the operation.

4. On the **Select Target Configuration** tab, select the model, instance type, model list, and other configuration items. After confirming the cost, click **Start Adjustment**.

5. The fees incurred by adjusting the configurations of different nodes to the same configuration will be displayed on the billing details page.

hange Conf	iguration				
Select Config	Target uration	> (2	Cost details > (	3 Shutdown notes	
(i) • The	configuration char	ige does not appl	y to local models and pod resour	ces.	
• The	configuration char	ige does not appl	y to data and system disks.		
<ul> <li>If you</li> <li>"Contraction</li> </ul>	ou want to change nfiguration manage	resources of the c ement" and restar	orresponding component after th t the service.	ne configuration change, rede	liver the configurations in
odes selected:	1 See Less 👻				
Resource ID		Node type	Current configuration	Billing Period	Operation
		master	EMR StandardSA2 CPU: 4-core; memory: 16GB	-	Adjustable configuration
lect a target c	configuration				
odel	Standard	MEM Optimiz	ced Compute		
pe	Standard SA2	2 Standard	S5		
odel List 访	Specification: S	A2.2XLARGE16; C	PU: 8-core; m 💌		
otal cost to be	frozen				
			Novt	ancel	

6. (Optional) To adjust the resources of the component after configuration adjustment, you need to deliver the configuration again in **Configuration Management** and restart the service.

Select Target Configuration	>	2 Cost details	> 3 Shutdown not	25	
<ul> <li>Costs in a batch</li> <li>Discounts or vo</li> <li>If the configura</li> </ul>	h configuration chan buchers (if any) used tion is changed, the	ige will be deducted item by it in the purchase will not be re pay-as-you-go nodes will be	tem. Please ensure your account turned. charged again from the first tier.	balance is sufficient. Please proceed with caution	. <u>View details</u>
Resource ID	Node type	Current configuration	Target configuration	Billing Period	Cost
		EMR StandardSA2	Standard SA2		

#### Caution

The YARN resources will be automatically adjusted according to the model and specification by default. After configuration adjustment, the size of the resources will change as the specification changes and does not need to be adjusted manually.

If you have manually adjusted the configuration of YARN resources, then after configuration adjustment, you need to modify the parameter values of the yarn.nodemanager.resource.cpu-vcores and

yarn.nodemanager.resource.memory-mb configuration items in Configuration Management, click Save configuration to deliver the configuration, and restart the NodeManager service for the configuration of YARN resources to be updated.

# Automatic Replacement

Last updated : 2023-12-27 11:07:18

## Overview

The system continuously monitors task nodes and router nodes in a cluster for any abnormal running status. If any exceptions are detected, the system automatically purchases nodes of the same model to replace the affected nodes. Meanwhile, alarms are sent to notify users of the replacement result.

#### Caution

When "Instance Termination Protection" is enabled for nodes, this action cannot be triggered.

## **Replacement Principle**

1. Only pay-as-you-go and spot instance nodes support automatic replacement.

2. The system purchases only nodes of the same model in the same billing mode as abnormal nodes for automatic replacement.

3. When automatic replacement is not enabled, you can still receive alarms in the event of any node abnormalities.

4. Automatic replacement and exception monitoring are not supported for Pod nodes.

### Directions

1. Log in to the EMR console and click **Cluster List** in the left sidebar. On the page that appears, click the **ID/Name** of the cluster for which you want to enable automatic replacement.

2. Choose Instance Information > Basic Configuration and enable Automatic Replacement.

Basic configuration	The system can monitor the status of task and router nodes in the cluster. When detecting node exceptions, it automatically removes abnormal pay- as-you-go and spot instance nodes and repurchase nodes with the corresponding configurations for	Network information
Security Group sg-Or curity group	replenishment.	Master Public IP 11 Check for Updates
Billing mode Pay-as-you-go	Auto replenishment 🕥 🔹	COS 🛈 Not aut rize
Custom Service Role 🕄 Not set Set	Host Login  Password Set	Security Group -

# Exporting Software Configuration

Last updated : 2023-12-27 11:07:38

# Overview

You can export the software configuration parameters of existing clusters in the EMR console, so that you can reuse these parameters for software configuration to quickly create a cluster in the future.

#### Note

You can export configuration files at the cluster but not configuration group or node level.

To reuse the configuration of an existing cluster for a new cluster, we recommend you export only the modified configuration files of the existing cluster but not all the configuration files.

### Directions

- 1. Log in to the EMR console and go to the cluster list page.
- 2. In the Action column of the target cluster, select More > Export Software Configuration.

Elastic MapReduce	Cluster South China (Guar	ngzhou) 🔻					Cluster Mar	nagement Do	cum
Cluster	Create Cluster More	Ŧ							¢
⑦ Monitor	D/Name \$	Instance St Pro	ject Region/Zone	Billing Mode	Network	Master Nod	_ Creation _ \$	Action	
Components		Cluster runni DEF	AULT PR South China Guangzhou	Pay-as-you-go		η	<u>a</u> 2020-02-27	Details More ▼	
	Total 1 item					Lines per page 10 💌	Restart Component Terminate	t	•
							Export Software Co	nfiguration	

3. Select the files to be exported and click **Export Configuration** to download the software configuration files.

Export Software Configuration	×
Components Involved in Cluster Software Configuration Parameters: HDFS, YARN Select configuration files:	
Select All	
HDFS	
hadoop-env.sh hdfs-site.xml	I
log4j.properties	
Export Configuration Cancel	

# **Cluster Scripts**

Last updated : 2023-12-27 11:07:55

### Feature Overview

The cluster script feature allows you to run a script on multiple nodes at a time for higher efficiency. Only one cluster script can be run in a cluster at a time. If a cluster script is running (such as installing a third-party software or modifying the running environment of the cluster), you cannot submit and execute new cluster scripts.

#### Caution

Only files in COS can be selected as the script to run.

Only nodes in the current cluster can be selected to run the script on.

You can set custom parameters based on your business requirements.

The EMR console does not verify the scripts you run. Therefore, running a custom script is risky, so please operate with caution.

## Prerequisites

The cluster script feature is only available for clusters that are running. The script to be run must be a shell script file in the COS STANDARD storage class. To use the cluster script feature, you must grant EMR the permission to access COS.

## Directions

1. Log in to the EMR console and click a cluster ID in the cluster list to go the instance information page. Click **Cluster Scripts** > **Set Task**.

← mer delengement Hay serviced	Cluster Scripts							Help Documentatio
Basic Information   Instance Information	1. You can run a 2. The EMR con	script on multiple nodes at a sole does not verify the script	time as needed, for purposes s you run. Therefore, running :	such as installing a thi a custom script is risky	rd-party software or modifyin , so please operate with cau	ng the running environment ttion.	of the cluster.	
Bootstrap Actions Cluster Service	Set Task	Today Yesterday	Last 7 days Last 30	days 2021-02-23	o 2021-02-23 📋			Enter a task or si C
Cluster Resource	Task ID	Task Name	Script Address	Start Time	End Time	Parameter	Status	Operation
Cluster Monitoring					No data yet			
Auto-scaling	Total items: 0						10 👻 / page	✓ 1 / 1 page >
Cluster Scripts								
Operation Logs								



2. Set the **Task Name**, **Script**, **Nodes**, and **Custom Parameter** fields. After that, click **Run** to generate a task in the task list.

Set Task					×
Task Name *	The length should b	e 6-36 characte	ers. Chinese, letters, numb	ers, -, and _ are supported.	
Script *	D Select a shell script	file in the COS	STANDARD storage class		
Nodes *	Separate keyword	ds with " "; press	Enter to separate filter tag	js	Q
	Node Type	Node IP	Deploy Task	Deployment Process	
	core		HDFS, FILEBEAT, YARN	DataNode、Filebeat、 NodeManager	•
	master		YARN、HDFS、 ZOOKEEPER、 FILEBEAT、KNOX	Ldap, NameNode, QuorumPeerMain, ResourceManager, Filebeat, Gateway, JobHistoryServer	Ţ
	Selected nodes (0):				
Custom Parameter	Up to 200 charact	iers			
		Run	Cancel		

3. There are various task statuses depending on the execution results, including successful, failed, partially failed, etc. You can click **Details** to view details.

emr-6dmpunsz	Cluster Scripts Help Documen	tatior
Basic Information	1. You can run a script on multiple nodes at a time as needed, for purposes such as installing a third-party software or modifying the running environment of the cluster. 2. The EMR console does not verify the scripts you run. Therefore, running a custom script is risky, so please operate with caution.	
Bootstrap Actions Cluster Service Cluster Resource	Set Task         Today         Yesterday         Last 7 days         Last 30 days         2021-02-23 to 2021-02-23 to         Enter a task or size         Enter a task or size           Task ID         Task Name         Script Address         Start Time         End Time         Parameter         Status         Operation	Q
Cluster Monitoring   Auto-scaling	191 Roy-001 2021-02-23 10:19:52 2021-02-23 10:19:53 Running Details Delete	
Cluster Scripts Operation Logs	interningite internitional internet int	
	Effective for the former of th	
	Total items: 1 10 ¥ / page × × 1 / 1 page	P

Copy Suco	cessful Node IPs	Copy Failed Nod	le IPs		
Node IP	Node Y	Start Time *	End Time *	Execu Y	Running Result
	core	2021-02-23 10:19	2021-02-23 10:19	Running	≜ stdout: stderr:
					stderr:

4. A cluster script may run successfully on some nodes and fail on other nodes. You can batch copy the failed nodes to run the script again.
# **Cluster Termination**

Last updated : 2023-12-27 11:08:14

## Feature Overview

When you no longer need an EMR cluster, you can terminate it in the EMR console.

# Prerequisites

Pay-as-you-go cluster: once terminated, the cluster will not be retained in the recycle bin but will be completely terminated and cannot be recovered. Please do so with caution.

Before terminating a cluster, please make sure that your data has been backed up as it cannot be recovered after termination.

If there is an EIP (an IP with a secondary ENI), the instance will be retained once returned, and the idle IP will continue to incur fees. If you don't need to retain it, please release it on the corresponding resource management page.

## Directions

### Caution

If MetaDB in the cluster to be terminated is associated with an external cluster as a Hive metadatabase, it will be retained when the cluster is terminated; if you no longer need it, please go to the TencentDB console to terminate it. Hive metadatabases cannot be recovered once terminated. Please do so with caution.

Log in to the EMR console, select Action > More > Terminate to enter the cluster termination page. Confirm the information of the cluster that needs to be terminated. After confirming that everything is correct, check I have read and agree to and click Next.

<ul> <li>emr-cygigid b Cluster is ru wendanli-using * Cluster is ru * * Cluster is ru</li> <li>Check Info of the Item to Terminate (2) Confirm Termination</li> <li>Check Info of the Item to Terminate (2) Confirm Termination</li> <li>Cluster is ru * * Cluster is ru * * Cluster is ru * * Cluster is ru</li> <li>Check Info of the Item to Terminate (2) Confirm Termination</li> <li>Cluster is ru * * Cluster is ru * * * Cluster is ru</li> <li>Check Info of the Item to Terminate (2) Confirm Termination</li> <li>Cluster is ru * * * * * * * * * * * * * * * * * * *</li></ul>		ID/Name *	Instance Sta	Termina	te Cluster			,	× 1	ime Created ‡	Action
emr-8dgmahbj @       Cluster is ru         if there are EIPs (including the IPs on the secondary ENI), the instance will be retained after returned, and the idle IPs will continue to incur fees. If you don't need to retain it, please release it on the corresponding resource management page. data in advance.       2019-07-22 11:3         emr-95bccwill @       cluster is ru       I node selected. View Details *       2019-07-22 11:3         Total 3 items       No.       Cluster Name       Cluster ID       Project       Availability Zone	~	emr-cygjqjkd ा⊓ wendanli-using ≠	Cluster is rur	1 Ch	eck Info of the Item to T	erminate )	2) Confirm Termination		2	019-07-23 16:0	Details
emr-95bcowling       management page. A terminated instance will be retained in the recycle bin for 7 days. Please back up the data in advance.       2019-07-18 15.2         edwin-test-maxCluster is rule       I node selected. View Details *       I node selected. View Details *         Total 3 items       No.       Cluster Name       Cluster ID       Project       Availability Zone		emr-8dgmahbj 🖻 wendanli-spark	Cluster is ru	If the idle IF	re are EIPs (including the Ps will continue to incur fo	IPs on the secondary ENI) ees. If you don't need to ret	, the instance will be retai tain it, please release it or	ined after returned, and the the corresponding resource	2	019-07-22 11:3	Details
Total 3 items No. Cluster Name Cluster ID Project Availability Zone		emr-95bccwil to edwin-test-max	Cluster is rur	mana data i 1 node se	igement page. A ter n advance. elected. View Details ▼	rminated instance will be r	etained in the recycle bin	for 7 days. Please back up the	2	019-07-18 15:2	Details
Churter ID k	Total	l 3 items		No.	Cluster Name	Cluster ID	Project	Availability Zone		K K 1	/1 page
1 v g Gun 3				1	v g	Cluster ID		Gu			

After confirming that everything is correct on the **Confirm Termination** tab, click **Start Termination** to terminate the cluster.

EMR - Instance List	All pr	ojects « Sou	uth China (Gua	nazhou) x	,						Clu
	+ (	Create Cluster	Assian to P	Termina	ate Cluster				×		
		ID/Name +	Instance St	🕑 Che	eck Info of the Item to	Terminate	2 Confirm Term	ination		Time Cre +	Action
		emr-cygjqjkd wendanli-usi	Cluster is running	This res 8 instan	sources to be termina ces,Collapse ×	n <mark>ted</mark> are as follows	E			2019-07-23 1	Details More *
		1		No.	Instance name	Instance ID	Current upper li	Instance Type			
		emr-8dgmahbj	Cluster is	1	emr-	0	5	EMR StandardS2	-	2019-07-22 1	Details
		P P	runnig	2	emr-ma		0	EMR StandardS2			WOI'e -
		emr-95bccwll edwin-test-m	Cluster is	3	emr-coi L	ž	0	EMR StandardS2	•	2019-07-18 1	Details More *
		1		16 cloud	d disks,Collapse 👻						
	Tota	al 3 items		No.	Cloud disk name	Cloud disk ID	Attributes	Configuration		K 4 1	/1 page
				1			Data disk	Efficient cloud disk 200G			
				2			System disk	Efficient cloud disk 50G			
								Efficient cloud	*		
				This res	sources to be retained	are as follows:					
				0 cloud 0 public	EIPs						
						Back Stat	rt Termination				
						BUCK OLD					

# **Operation Logs**

Last updated : 2023-12-27 11:08:30

# Feature Description

Operation logs record the operations you performed on clusters in the EMR console, such as creating and scaling clusters.

# Directions

After successfully creating a cluster, log in to the EMR console, select a region in the top-left corner of the cluster list, click the ID/Name f the target cluster to enter the cluster details page, and click Operation Logs on the left.

← antriktpopela 1. júnyor_1330656	Operation Log	s					Heir	Documentatio
Basic Information	Operation Ot Component	iject	Operation Deliver Configuration	Operation Details Cluster Level	Security Level	Operation Time 2021-07-23 15:08:5	Operator 52 100013118462	
Cluster Service	Component		Add Component	Add Component: hive-2.3.7 tez-0.1 impala-2.10.0	9.2	2021-07-21 10:21:4	47 100013118462	
Cluster Resource     Node Status	Cluster		Scale Out Cluster	-		2021-06-15 20:12:5	58 100013118462	
Resource Management	Cluster		Add Node Specification	Node Type: TASK Node Specifical SA2Model 2Core8G Efficient Clou 100G * 1	ion: d Disk	2021-06-15 20:12:4	47 100013118462	
Cluster Monitoring ~ Auto Scaling	Cluster		Create Cluster			2021-06-15 20:00:	53 100010380631	
Cluster Scripts	Total 5 items						Lines per page 10 v H 🖪 1	/1 page →
operation Logs								

# Task Center

Last updated : 2023-12-27 11:08:48

EMR has a **Task Center** where you can view the execution status of all tasks in your clusters and retry, cancel, and display error information for tasks.

## Task List

Log in to the EMR console and click **Task Center** on the left sidebar to view the status, start time, end time, and operation information of all tasks in one or all clusters in the current region.

Today Last 15 days	2022-04-09 15:30:23 to 2022-0	4-24 15:30:23 🖽 All clusters	Ŧ		
sk Name	Cluster ID/Name	Status T	Start Time \$	End Time \$	Operation
ld Component		Running 1%	2022-04-24 15:28:30		Task details
eate Cluster		Done	2022-04-24 15:21:12	2022-04-24 15:27:03	Task details
tal 2 items				Lines per page 10 -	I ✓ 1 /1 pag

Log in to the EMR console, click **Cluster List** on the left sidebar, select a cluster, and click **Task** in the top-right corner to view the status, start time, end time, and operation information of all tasks in the cluster.



e	Instance Information		Help Documentation
Basic information	Basic Configuration		
Instance Information	Instance ID	Region Info Guangzhou(Guangzhou Zone 7)	High Availability Cluster Enabled
Bootstrap Actions	Network information (	Security Group	roups Cloud disk encryption Not enabled
Cluster Service	Creation time 2022-04-24 15:21:12	Master Public IP Check for Update	s Billing mode Pay-as-you-go
Cluster Resource	COS <sup>①</sup> Not authorized Authorize	Custom Service Role (1) Not set Set	Host Login (1) Password Set
Cluster Monitoring			
Auto Scaling	Software Info	Hardware	Info
Cluster Scripts	Cluster Type Hadoop	Master Node	2
Operation Logs	Use Cases Hadoop-Default	Core Node	3
	Product Version EMR-V2.2.0.tlinux	Common No	de 3
	Components to Deploy hdfs-2.8.5,yarn-2.8.5,zookeeper-3.5.5,kn	10x-1.2.0,hive-2.3.5	
	MetaDB cdb-kcx3dajr		
	Hive Metadatabase cdb-kcx3dajr		
	Kerberos mode Not enabled		

## Task Details

Click **Task Details** to view task parameters and completion status of execution steps.

Task Parame	ters					
ClusterName	ClusterName ceshi123		2		MasterSize 1	
SoftInfo	SoftInfo zookeeper-3.5.5,hadoop-2.8.5,knox-1.2.0		false			
No.	Step Name	Status		Start Time	End Time	Operation
0	Confirm Order	Done		2020-05-07 18:12:46	2020-05-07 18:12:46	-
1	Create CVM	Done		2020-05-07 18:12:47	2020-05-07 18:13:04	-
2	Create TencentDB	Done		2020-05-07 18:12:47	2020-05-07 18:12:47	-
3	Bind tag	Done		2020-05-07 18:13:06	2020-05-07 18:13:06	-
4	Initialize security group	Done		2020-05-07 18:13:10	2020-05-07 18:13:10	-

Task details					
Task Parameters					
CoreSize 1			DeployProcess DataNo	de,NodeManager,Filebeat	
Node Type Core					
No.	Step Name	Status	Start Time	End Time	Operation
1	Create Order	Done	2022-04-24 15:36:17	2022-04-24 15:36:18	-
2	Apply For Resource	Done	2022-04-24 15:36:19	2022-04-24 15:36:37	-
3	Bind tag	Done	2022-04-24 15:36:38	2022-04-24 15:36:39	-
4	Initialize security group	Done	2022-04-24 15:36:39	2022-04-24 15:36:41	-
5	Initialize resource	Done	2022-04-24 15:36:41	2022-04-24 15:36:42	-
6	Prepare for scale-out	Done	2022-04-24 15:36:42	2022-04-24 15:36:42	-
▶ 7	Initialize service	Done	2022-04-24 15:36:43	2022-04-24 15:40:01	Run Details
8	End task	Done	2022-04-24 15:40:03	2022-04-24 15:40:03	-
		Confirm	Cancel		

## Run Details

In some operation steps in **Task Details**, you can click **Run Details** to view the completion status of the corresponding steps.

Elastic MapReduce	← Task details (Scale Out C	Core)					
Cluster List							
Container-Based Cluster	Task Parameters						
Task Center	CoreSize 1		DeployProcess DataNode,N	lodeManager,Filebeat	Node Type Core		
	No.	Step Name	Status	Start Time	End Time	Operation	
	1	Create Order	Done	2022-04-24 15:36:17	2022-04-24 15:36:18	-	
	2	Apply For Resource	Done	2022-04-24 15:36:19	2022-04-24 15:36:37	-	
	3	Bind tag	Done	2022-04-24 15:36:38	2022-04-24 15:36:39	-	
	4	Initialize security group	Done	2022-04-24 15:36:39	2022-04-24 15:36:41	-	
	5	Initialize resource	Done	2022-04-24 15:36:41	2022-04-24 15:36:42	-	
	6	Prepare for scale-out	Done	2022-04-24 15:36:42	2022-04-24 15:36:42	-	
	<del>~</del> 7	Initialize service	Done	2022-04-24 15:36:43	2022-04-24 15:40:01	See Less	
	Node's IP	State Y	Start Ti	ime End Time		ErrorDetails	
				No data yet			
	8	End task	Done	2022-04-24 15:40:03	2022-04-24 15:40:03	-	

Task details					
Task Parameters			DeployProcess DataNo	de,NodeManager,Filebeat	
Node Type Core					
No.	Step Name	Status	Start Time	End Time	Operation
1	Create Order	Done	2022-04-24 15:36:17	2022-04-24 15:36:18	-
2	Apply For Resource	Done	2022-04-24 15:36:19	2022-04-24 15:36:37	-
3	Bind tag	Done	2022-04-24 15:36:38	2022-04-24 15:36:39	-
4	Initialize security group	Done	2022-04-24 15:36:39	2022-04-24 15:36:41	-
5	Initialize resource	Done	2022-04-24 15:36:41	2022-04-24 15:36:42	-
6	Prepare for scale-out	Done	2022-04-24 15:36:42	2022-04-24 15:36:42	-
▼ 7	Initialize service	Done	2022-04-24 15:36:43	2022-04-24 15:40:01	See Less
Node's IP	State <b>Y</b>	Start Time	End Tim	ie Erro	prDetails
		No data	a yet		
		Confirm	Cancel		

# Disk Management Scaling up Cloud Disks

Last updated : 2023-12-27 11:06:40

## Overview

If the data storage space of a node becomes insufficient as your business grows, you need to expand the space. This document describes how to scale up cloud disks in the EMR console.

### Caution

You can scale up only cloud data disks but not system disks and local disks.

You cannot batch scale up cloud data disks on multiple nodes.

For your data security, we recommend you create snapshots of your cloud disks before scale-up.

To prevent data loss, a disk can be scaled up only but not scaled down.

# Scaling up Cloud Disks

### Directions

1. Log in to the EMR console and click the **ID/Name** of a target cluster in the cluster list to go to the cluster details page.

2. Click the **Resource Management** tab and then click **Scale up cloud data disk** in the **Operation** column of the target node to enter the settings page.

Resource Management						🖽 Help Document
The current cluster resources were purchased based o current page shall prevail. It is not allowed to adjust the shall be carried as the shall be carried as a statement of	n the official EMR billing ne fee policy on the CVM	rules and the renewal statu: Console.	ses of CVM resources are managed by EMR in a unified	manner. For monthly-sub	scription CVM resources, th	e billing statuses shown on the
All nodes         Master         Core         Common           Scale Out         Scale In         More         Image: Common Scale In	Task Router	Metadb Recyle	Renew			ı.
Resource ID Node type	Resource type <b>T</b>	IP	Configuration	Creation time 1	Expiration time	Operation
Master	HOST		EMR StandardS5 CPU: 8-core; memory: 16G8 System disk: Premium cloud disk 50G x 1 Data Disk: Premium cloud disk 200G x 1	2023-04-25 17:19:39	Pay-as-you-go	Change Configuration More
Core	HOST		EMR StandardS5 CPU: 4-core; memory: 8GB System disk: Premium cloud disk 50G x 1 Data Disk: Premium cloud disk 200G x 1	2023-04-25 17:19:40	Pay-as-you-go	Change Configuration More
Core	HOST		EMR StandardS5 CPU: 4-core: memory: 8G8 System disk: Premium cloud disk 50G x 1 Data Disk: Premium cloud disk 200G x 1	2023-04-25 17:19:41	Pay-as-you-go	Change Configuration More Start/Stop Servi
Total 3 items					Lines per page 20 🔻	Renew Set auto-renew Scale up cloud

3. If multiple cloud data disks are mounted to the current node, you can batch scale them up to the same capacity.

Scale up cloud data dis	k			>
Select cloud data disk	> 2 Cha	ange size		
<ul> <li>Notes: We strongly recon file system. Only disk scale-up</li> </ul>	nmend that you snapshot the c	loud disks before scale-up to avoid miso s.	peration-related damage to the	
Cloud data disks selected: 1	See Less 🔻			
Cloud disk ID	Billing mode	Expiration time	Current size	
disk-gea2qs74	Pay-as-you-go	-	200GB	
Common target size	200 GB	500 GB 1000 GB	- 250 32000 GB	+ G
Cost				
	Back	Confirm Disable		

4. Disks will be initialized automatically after scale-up, and you don't need to manually update the disk information.

# **Disk Update Check**

Last updated : 2023-12-27 11:06:24

# Feature overview

This feature checks whether the cluster disk information is consistent with that displayed in the console and supports performing updates to ensure consistency.

# Directions

1. Log in to the EMR console and click the ID/Name of the target cluster in the Cluster list to enter the cluster details page.

On the cluster details page, select Cluster Resource to enter the Resource Management page. Select More > Check disk update. If the cluster disk information is inconsistent with the information displayed in the console, click Update, and the disk usage, total disk space, and disk space usage of the cluster will all be updated.



# **Repairing Disks**

Last updated : 2023-12-27 11:02:43

# Overview

Local disk replacement events are automatically monitored in the EMR console. After disk replacement, you can initialize the new disk in the console on your own.

### Caution

After you receive a faulty disk notification from the CVM and repair or replace the physical disk as instructed in the notification, the disk repair operation is triggered in the EMR console.

When a disk is replaced, all data on the disk is lost. Make sure that the data on the disk is backed up prior to replacement.

## Directions

1. Log in to the EMR console and click the ID/Name of the corresponding cluster in the cluster list.

2. On the cluster details page, choose **Cluster Resource** > **Resource Management** and then repair the disk on the node where the disk is replaced.

3. In the disk repair process, the services on the current node will be restarted, during which both the services and node are unavailable. Therefore, we recommend you repair a disk during off-peak business hours.

# Kudu Service Recovery

### Caution

If there are multiple local disks and one or more of them use the EMR disk repair feature, the node where the disks are located will run the KuduServer service.

Due to the limitation of Kudu's fs\_data\_dirs feature, to ensure normal startup of KuduServer after one or more disks are formatted, all data directories on the KuduServer node are empty. You should confirm that these directories are not used by any businesses other than Kudu.

Scenario:

Under **Cluster Service** in the EMR console, you can view the health status of the KuduServer on the node where the disk is replaced. The health status is unavailable:

Data consistency check and recovery:



1.1 Ensure that the directory (as described below) is used exclusively for Kudu. If the directory is used for any other purpose, move the relevant data to another directory that is not configured as an fs\_data\_dirs directory before proceeding with the following steps.

Specific directory: View the /usr/local/service/kudu/conf/tserver.gflags file:



1.2 Log in to the abnormal node of the local disk to view logs: /data/emr/kudu/log/kudu-tserver.INFO :

10110 14:16:50.918087 96777 minidump.cc:244] Setting minidump size limit to 20M 10110 14:16:50.918085 96777 env\_posix.cc:2179] Not raising this process' open files per process limit of 10000000; it is already as high as it can go 10110 14:16:50.918025 96777 hybrid\_clock.cc:349] constructed file cache file cache with capacity 400000 10110 14:16:50.918053 96777 hybrid\_clock.cc:349] auto-selected time source: bulltin 10110 14:16:50.912734 96777 hybrid\_clock.cc:349] witing up to --ntp\_initial\_sync\_wait\_secs=60 seconds for the clock to synchronize 10110 14:16:50.927374 95777 hybrid\_clock.cc:383] waiting up to --ntp\_initial\_sync\_wait\_secs=60 seconds for the clock to synchronize 10110 14:16:50.927374 95777 hybrid\_clock.cc:349] Webserver started at http://0.0.0.0:8058/ using document root /usr/local/service/kudu/www and password file <none> 10110 14:16:50.927525 96777 server\_base.cc:612] This appears to be a new deployment of Kudu; creating new FS layout 10110 14:16:50.923141 96777 tablet\_come\_anin\_ccidel\_Alleredy present: RunTabletServer() failed: FS layout already exists; not overwriting existing layout: unable to create file system roots: FSManager eady exist://data/emr/kudu/tserver./datal/emr/kudu/tserver

Run the following commands as a root user to remove inconsistent data:

```
rm -rf /data/emr/kudu/tserver/*
rm -rf /data1/emr/kudu/tserver/*
```

The commands assume that /data/emr/kudu/tserver/ and /data1/emr/kudu/tserver/ are

configured in fs\_data\_dirs . Fore more information, view

/usr/local/service/kudu/conf/tserver.gflags .

1.3 Observe the service status of KuduServer in the console.

#### Note

If you encounter any issues, submit a ticket.

# Managing Service Managing Users

Last updated : 2023-12-27 11:09:05

## Overview

This feature allows you to manage (add, edit, or delete) users in an EMR cluster easily. New users can submit Hadoop cluster tasks.

### Note:

User management is available only for EMR-v2.6.0 and EMR-v3.2.1 and above.

You need to manually trigger the delivery of configurations in ranger-ugsync-site.xml and restart the EnableUnixAuth service to sync user configurations.

Deleting users and resetting passwords may cause a running task to fail.

## Directions

Log in to the EMR console and click the ID/Name of a cluster in the cluster list to enter the cluster details page.
 Click Users. On the displayed page, you can add and delete users in batches, reset passwords, and download keytabs.

emr-edqx53yb	Users				🖽 Help docur
EMR-Jhaoy/Jm					
Cluster overview	The user list shows on	ly the user information added via the Users page.			
Instance info					
Cluster services	Create user Delete				Separate keywords with $"\! ";$ press Enter to separate filter tags
Cluster resources	Username	User group	Remarks	Creation time	Operation
Cluster monitor	root	root	-/	2023-05-17 17:46:42	Reset password Download Keytab
* DashBoard BETA	hadoop	hadoop	-1	2023-05-17 17:46:42	Reset password Download Keytab
Cluster events					
* Logs	Total 2 items				Lines per page 10 • II / I page
Cluster inspection					
Alarm records					
* Java analysis					
Auto-scaling					
Scripts ~					
Users					
Operation logs					



3. Click **Create user** to create a user. **Username**, **User group**, and **Password** are required and **Remarks** is optional.

Create user			×
<ul> <li>1. Support</li> <li>2. A usern</li> </ul>	ts creating multiple users at a ame or user group name mu	a time. st be 1–30 characters long and contain letters, digits, hyphens, and und	erscores.
User Info			
Username*	User group*(j)	Remarks	Operation
Enter	Enter	Max 128 characters	Delete
+ Add			
Password *	Enter a password		
	8–30 characters; supports le	tters, digits, -, _, !, @, #, \$, and %	
Confirm password *	Enter the password again		
		Confirm	

4. Manually trigger the delivery of configuration for ranger-ugsync-site.xml and restart the service.

Log in to the EMR console, select the target cluster in the cluster list, and click \**Cluster services* to enter the cluster service list.

Select **Operation > Configurations** in the upper-right corner of the Ranger service panel.

On the **Configurations** page, select ranger-ugsync-site.xml, click **Edit configuration** and make no

modification. Then, click Save configuration below, and select Save and deliver in the pop-up window.

Return to the **Roles** page, select the EnableUnixAuth service, and click **Restart service**.

5. Reset a password.

On the user management page, find the user whose password needs to be changed and click **Reset password** on the right. Enter the new password and confirm, and click **Confirm**.

Reset password		×
Resetting the p operation.	assword may cause running tasks to fail. Please confirm this	
New password	Enter a password 8–30 characters; supports letters, digits, -, _, !, @, #, \$, and %	
Confirm new password	Enter the password again	
	Confirm Cancel	

6. Download a keytab.

On the user management page, find the user for whom you want to download the keytab and click **Download keytab**. The keytab can be used to log in to the cluster.

Users				🗉 Help document
<ol> <li>The user list shows only</li> </ol>	y the user information added via the Users page.			
Create user Delete				Separate keywords with "]"; press Enter to separate filter tags
Username	User group	Remarks	Creation time	Operation
test-11	hadoop	panjian 🎤	2023-07-13 15:47:54	Reset password Download Keytab De
root	root	-1	2023-05-17 17:46:42	Reset password Download Keytab De
hadoop	hadoop	-1	2023-05-17 17:46:42	Reset password Download Keytab De
Total 3 items				Lines per page 10 💌 🖂 🐇 1 /1 page 🕨

### Caution

Keytab downloading is available only for Kerberos clusters.

7. Delete a user.

On the user management page, find the user to be deleted and click **Delete** on the right. Then click **Delete**.

Confirm deletion	×
Deleting a user may cause running tasks to fail. Please proceed with caution.	
Are you sure you want to delete user test-11?       Confirm     Cancel	

# Adding Components

Last updated : 2023-12-27 11:09:22

# Feature Description

This document describes how to add components in the EMR console.

#### Caution

You need to manage the components through **Cluster Service**. Component changes made directly by logging in to a node (such as adding a component) will not be synced to the console and thus cannot be further managed. Components on a different EMR version or component version cannot be added. Instead, you can only add components on the current EMR version.

### Directions

1. Log in to the EMR console and click the **ID/Name** of the target cluster in the cluster list to enter the cluster details page.

2. On the cluster details page, select **Cluster Service** > **Add Component** to add a component not installed in the cluster.

Component Mana	gement East Ch	nina (Shanghai) 🔻 panjian12	3 🔻	Component m	nanagement document 🕑
Add Component	Rolling Restart	Reset Native UI Password			
Component name	Status	Version	Native WebUI Access Address(j)	Access Info	Operation
ZOOKEEPER	Running	3.4.9		QuorumPeerMain IPC:	Configuration More
HDFS	Running	3.1.2	https:// ateway/emr/hdfs	NameNode IPC: zkfc IPC: JournalNode IPC: DataNode IPC:	Configuration More
YARN	Running	3.1.2	https:// /gateway/emr/yarn	ResourceManager IPC: NodeManager IPC: JobHistoryServer IPC:	Configuration More

3. If there is no metadatabase in the cluster and Hue, Ranger, Oozie, Druid, and Superset components are to be installed, you need to purchase a TencentDB instance as the metadata store. You cannot add the Kudu component to a non-HA cluster.

4. There are two storage options for Hive metadata: you can store the metadata in the MetaDB (default option) or associate the metadata with EMR-MetaDB or a self-built MySQL database. In the latter case, metadata will be stored in the associated database and will not be deleted when the cluster is terminated.

The current cluster purchase a new Ter	has no metadatabase. If you icentDB instance to store m	ı want to add a component sı etadata.	uch as Hive, Sqoop, Hue, Ranger,	Oozie, and Presto, you need to
Cluster Version	EMR-V3.0.0.tlinux			
Optional Components	zookeeper-3.4.9	hadoop-3.1.2	knox-1.2.0	<b>sys-1.0</b>
	hbase-2.2.0	✓ hive-3.1.1	✓ hue-4.4.0	oozie-5.1.0
	presto-0.222	ranger-2.0.0	spark_hadoop3.1-2.	4.3 sqoop-1.4.7
	tez-0.9.2	flume-1.9.0	impala-2.10.0	alluxio-1.8.1
	flink-1.8.1			
Hive Metadatabase	Default Assoc	iate EMR-MetaDB Asso	ociate self-built MYSQL	
Instance ID	Select data 💌			
MetaDB ⑦	Pay-as-you-go 1 high IO T	encentDB with 4,000 MB men	nory and 100 GB disk	
Cost				

5. The selection for MetaDB purchase and Hive metadatabase is the same as when you purchase a new cluster.

6. After selecting the component, click **OK**.

# **Restarting Service**

Last updated : 2023-12-27 11:09:40

# Overview

After modifying the configuration of a component, you need to restart the service for the configuration to take effect. To minimize or avoid the impact of the service restart on your business, restart the service on a rolling basis. For instances with active and standby attributes, standby instances will be restarted first before active instances. Rolling restart takes longer than regular restart.

Services can be restarted in the console, and rolling restart is selected by default. Note that when rolling restart is disabled, restarting all nodes at the same time may make the service unavailable.

There are two policies for failure handling when node restart fails: waiting for processing in case of failure or continuing processing in case of single node failure.

## Directions

1. Log in to the EMR console and click the **ID/Name** of the target cluster in the cluster list to enter the cluster details page.

2. To restart the entire component, select **Operation** > **Restart Service** on the target component block of the **Cluster Service** page, or go to the component details page and select **More** > **Restart All Services** in the top-right corner. To restart a role instance, select **Operation** > **Role Management** on the target component block of the **Cluster Service** page, select the target service role, and click **Restart Service**.

When restarting a service on the component block or component details page, you need to select the target service role, restart method, rolling restart switch status, and failure handling policy.

If you set **Service Role** to **All**, the entire component will be restarted.

Restart Service		×
Service name *	HDFS	
Service Role *	NameNode	
Restart Range	All eligible nodes (total 2)	
Restart Method	Quick restart mode	
Rolling Restart	Each restart: 1 nodes, interval: 5 seconds	
	Up to 1 nodes can be restarted and the maximum interval is 5 minutes.	
Failure-Handling Policy	Wait for processing in case of failure	
Operation Reason		
	Enter no more than 200 characters	
	Confirm Cancel	

When restarting the service on the **Role Management** page, you only need to select the restart method and failure handling policy.

Restart Service		×
Service name *	HDFS	
Service Role *	NameNode •	
Restart Range	All eligible nodes (total 2)	
Restart Method	Quick restart mode	
Rolling Restart	Each restart: 1 nodes, interval: 5 seconds	
	Up to 1 nodes can be restarted and the maximum interval is 5 minutes.	
Failure-Handling Policy	Wait for processing in case of failure	
Operation Reason		
	Enter no more than 200 characters	
	Confirm Cancel	

### 3. The restart methods supported by service components are as follows:

Component	Service	Restart Mode	Description	Remarks
HDFS	NameNode	Quick restart mode	It can be restarted by running hadoop- daemon.sh stop   start namenode.	-
	NameNode	Safe restart mode	For an HA cluster, run `saveNameSpace` on the standby NameNode.Then, restart the component by	Only rolling restart is supported.



			running hadoop- daemon.sh stop   start namenode. For a non- HA cluster, the quick restart mode is used.	
	DataNode	Default restart mode	It can be restarted by running hadoop- daemon.sh stop   start datanode.	-
	JournalNode	Default restart mode	It can be restarted by running hadoop- daemon.sh stop   start journalnode.	-
	zkfc	Default restart mode	It can be restarted by running hadoop- daemon.sh stop   start zkfc.	-
YARN	ResourceManager	Default restart mode	It can be restarted by running sbin/yarn- daemon.sh stop   start resourcemanager.	Only rolling restart is supported.
	NodeManager	Default restart mode	It can be restarted by running sbin/yarn- daemon.sh stop   start nodemanager.	-
	JobHisotryServer	Default restart mode	It can be restarted by running sbin/yarn- daemon.sh stop  start historyserver.	-
	TimeLineServer	Default restart mode	It can be restarted by running sbin/yarn- daemon.sh stop   start timelineserver.	-
HBase	HbaseThrift	Default restart mode	It can be restarted by running hbase- daemon.sh stop   start thrift.	-
	HMaster	Default restart	It can be restarted by	-



		mode	running hbase- daemon.sh stop   start master.	
	HRegionServer	Quick restart mode	It can be restarted by running hbase- daemon.sh stop   start regionserver.	-
	HRegionServer	Safe restart mode	It can be restarted by running graceful_stop.sh restartreload.	-
	HiveMetaStore	Default restart mode	It can be restarted by running hcat_server.sh stop   strat.	-
Hive	HiveServer2	Default restart mode	It can be restarted by running hive-daemon.sh stop-h2   start-h2.	-
	HiveWebHcat	Default restart mode	It can be restarted by running webhcat_server.sh stop   start.	-
Presto	PrestoCoordinator	Default restart mode	It can be restarted by running bin/launcher stop   start.	Only rolling restart is supported.
	PrestoWorker	Default restart mode	It can be restarted by running bin/launcher stop   start.	-
ZooKeeper	QuorumPeerMain	Default restart mode	It can be restarted by running bin/zkServer.sh stop   start.	-
Spark	SparkJobHistoryServer	Default restart mode	It can be restarted by running sbin/stop- history-server.sh   sbin/start-history- server.sh.	-
Hue	Hue	Default restart mode	It can be restarted by running build/env/bin/start.sh	-



			and build/env/bin/stop.sh.	
Oozie	Oozie	Default restart mode	It can be restarted by running oozied.sh stop   start.	-
	Nimbus	Default restart mode	It can be restarted by running bin/storm- daemon.sh nimbus stop   start.	-
Ctorm	Supervisor	Default restart mode	It can be restarted by running bin/storm- daemon.sh supervisor stop   start.	-
Storm	LogViewer	Default restart mode	It can be restarted by running bin/storm- daemon.sh nimbus stop   start.	-
	UI	Default restart mode	It can be restarted by running bin/storm- daemon.sh nimbus stop   start.	-
Ranger	Ranger	Default restart mode	It can be restarted by running sbin/ranger- daemon.sh stop   start.	-
Alluxio	AlluxioMaster	Default restart mode	It can be restarted by running bin/alluxio- stop.sh master and bin/alluxio-start.sh master.	-
	AlluxioWorker	Default restart mode	It can be restarted by running bin/alluxio- stop.sh worker and bin/alluxio-start.sh worker.	-
Ganglia	Httpd	Default restart mode	It can be restarted by running /sbin/service http stop   start.	-



Gmetad	Default restart mode	It can be restarted by running /sbin/service gmetad stop   start.	-
Gmond	Default restart mode	It can be restarted by running /sbin/service gmon stop   start.	-

# Starting/Stopping Services

Last updated : 2023-12-27 11:09:56

# Overview

The Start/Stop Services feature can stop or start all services on a node. This document describes how to use Start/Stop Services in the EMR console.

## Directions

1. Log in to the EMR console and click the **Cluster ID/Name** of the target cluster in the **Cluster list** to enter the cluster details page.

2. Select Cluster Resource and click Resource Management. Then, select the target node and click More > Start/Stop Services.



3. Check the service information in the Start/Stop Services window.

4. After confirming that the information is correct, click Start All or Stop All.

# WebUI Access

Last updated : 2023-12-27 11:10:15

# Overview

The software WebUI entry feature is the accessibility to native component UIs provided by EMR. Such UIs can be quickly accessed through the public IP address of a master node (it is recommended to configure a security policy in a timely manner). If the private network of your cluster is interconnected with your organizational public network, you can disable this public IP address and access the UIs directly through the private network.

# Viewing WebUI Address

1. Log in to the EMR console and click the **ID/Name** of the target cluster in the cluster list to enter the cluster details page.

2. Click **Cluster Services** and then click **WebUI address** below a target component to access its WebUI.

The access address requires authentication. The username is "root", and the default password is the one entered when the cluster was created. To change the password, click **Reset WebUI password** on the page.

# WebUI Access over Private Network

If **Enable public network for master nodes** is not selected during the cluster creation, you cannot access the WebUIs of components via the WebUI addresses given on the **Cluster Services** page. You can access these WebUIs with a browser via the following links in a private network environment.

Component	Link
HDFS UI	http://{cluster IP in private network}:4008
YARN UI	<pre>http://{cluster IP in private network}:5004</pre>
HBase UI	<pre>http://{cluster IP in private network}:6001</pre>
Hive UI	<pre>http://{cluster IP in private network}:7003</pre>
Hue UI	<pre>http://{cluster IP in private network}:13000</pre>
Ranger UI	http://cluster IP in private network}:6080

Storm UI	<pre>http://{cluster IP in private network}:15001</pre>
Oozie UI	<pre>http://{cluster IP in private network}:12000</pre>
Ganglia UI	<pre>http://{cluster IP in private network}:1800</pre>
Presto UI	<pre>http://{cluster IP in private network}:9000</pre>
Allexio UI	<pre>http://{cluster IP in private network}:19999</pre>

To access the WebUI of a component over a public network after the cluster is created, you can set an EIP for the master node as instructed below:

1. Log in to the EMR console, click the Cluster ID of the target cluster to enter the cluster details page, then select Cluster Resources > Resource Management, and click the Resource name/ID of the target master node to go to the CVM console.

← emr-ldjnqr4a v_pjianpan_官网截图集	Resource Manageme	nt								н	lelp Documentatio
Basic Information	The current cluster resou shall prevail. It is not allow	rces were purchased wed to adjust the fee	I based on the official EMI policy on the CVM Conso	R billing rules and the rene le.	ewal statuses of CVM re	sources are managed by EMR in a unified m	anner. For monthly-sub	oscription CVM resources, i	the billing	statuses shown	on the current pag
<ul> <li>Bootstrap Actions</li> </ul>	All Nodes Master	Core Co	ommon Task	Router Metadb	Recyle Rene	w					
Cluster Service	Scale Out Scale	n More -	•	-						Search IP	Q
Cluster Resource	Node Type	Resource T	Resource Name/.	. IP	Configuration	Deployment Process	Expiration Time	Creation Time	Tag		Operation
Node Status     Resource Management Cluster Monitoring	Master	HOST	an on Stranja No Sectory	HL HOLD HE (private) HL F1.38 (HS (private) (private)	EMR StandardSA2 OF Land; memory: 8GB Premium cloud disk: 100G x 1	knox, HiveMetaStore, HiveServer2, HiveW ebHcat, Impala-Catlog, Impala- StateStore, JobHistoryServer, NameNode , Zookeeper, ResourceManager	Pay-as-you-go	2021-06-15 20:01	0		Change Configura
Auto Scaling Cluster Scripts Operation Logs	Core	HOST	ana ana gilikana ng tangangaj	NUMBER (DOM	EMR StandardSA2 CPU: 2-core; memory: 8GB Premium cloud disk: 100G x 1	DataNode,Impala-Server,NodeManager	Pay-as-you-go	2021-06-15 20:01	0	:	Change Configura Start/Stop Service:
	Core	HOST	ana considerativo Pre-tonologo	NL KOLO IO - (DOMINI)	EMR StandardSA2 CPU: 2-core; memory: 8GB Premium cloud disk: 100G x 1	DataNode,Impala-Server,NodeManager	Pay-as-you-go	2021-06-15 20:01	0	:	Change Configura Start/Stop Service:
	Task	HOST	ana conspirațețe a managentică	NUMBER OF	EMR StandardSA2 CPU: 2-core; memory: 8GB Premium cloud disk: 100G x 1	Impala-Server,NodeManager	Pay-as-you-go	2021-06-15 20:13	¢		Change Configura Start/Stop Service More 👻
	Total 4 items							Lines per pa	ge 20 <del>v</del>	н н	1 /1 page

2. Change the network bandwidth setting of the CVM instance to be bound with an EIP, so that its bandwidth is not 0 to ensure connection.

In the CVM console, select **More actions > Resource adjustment > Adjust network** against the target instance in the CVM instance list.



Bandw	vidth Package	Default P	roject	Log Ir	More 🔻			
			Purc	hase with	h same con	figurations		
			Insta	ance Stati	us	►		
Bandw	vidth Package	Default P	rc Insta	Instance Settings				
			Rein	stall the	system			
			Pass	word/key	/	Þ		
Bandy	Adjust Configurat	ion	Reso	ource Adj	ustment	Þ		
banar	Expand Data Disk		Crea	ite Image	:			
	Expand System Di	isk	IP/E	NI		•		
	Change Disk Med	ia Type	Secu	urity Grou	ıps	Þ		
Band	Adjust Network		oject	Log In	More 🔻			
	Switch VPC							
	Add to Bandwidth	n Package						

Set an appropriate target bandwidth cap to ensure the bandwidth of the CVM instance is greater than 0. 3. Click **Instance ID** of the CVM instance to enter its basic information page, and select the ENI tab.

1 result found for "Instance ID:ins-m5drck4m" Back to list									
emr-master-	Guangzhou Compute C3	4-core 8GB 5Mbps System disk: <b>Et</b> Premium Cloud (receate) Storage Network: Default-VPC	- Pay as you go Created at 2020- 02-27 17:15:51	Bandwidth Log In More * Package					

4. Click **Bind** to bind an existing EIP or create a new EIP for the CVM instance.

Bind EIPs		×
Please select the EIP that to b You can select an existing EIP	e bound with "private or create a new one Refresh List	
Search by name/ID		Q
ID/Name	EIP	
0		Î
o ein lachs0as		
	127	
0		
		*
	OK Close	

After the EIP is bound, you can see the EIP information beside the public network IP bound to the primary ENI on the ENI page.

5. Check the CVM access over public network.

6. Use the ping or ssh command to check whether the EIP is effective. Ensure the security group inbound rule is available to ICMP and port 22. Then, access WebUIs of components.

EMR v1.3.1, EMR v2.0.1, EMR v2.1.0, and EMR v3.0.0 now support Apache Knox. By default, you can access WebUIs of components over a public network with Knox as the proxy. For UI details of all components and Knox usage, see Knox Development Guide.

#### Note

After an EIP is bound, the WebUI addresses given in the EMR console will remain unchanged. To change these addresses, please contact us.

## Resetting WebUI Password

The access address requires authentication. The username is "root", and the default password is the one entered when the cluster was created. You can change the password as instructed below.

1. Log in to the EMR console and click the **ID/Name** of a target cluster in the cluster list to enter the cluster details page.

2. Click **Cluster Services** on the cluster details page and click **Reset WebUI password** in the top-left corner to reset the password.

#### Caution

The password must contain 8 to 16 characters in at least three of the following character types: uppercase letters, lowercase letters, digits, and special symbols (! @ # % ^ \*). The first character cannot be a special symbol. If OpenLDAP is installed in the cluster (EMR v2.6.0 and EMR v3.2.1 or later), you can change the password only on the **User Management** page. To reset the WebUI password, go to **User Management** and create a user.

Reset Native UI Passwo	rd ×
Enter a new password *	
	It must be 8-16 characters containing at least three of the following types: uppercases, lowercases, numbers, and special characters !@#%^*. It cannot begin with a special character.
Confirm the new password *	
_	

# **Role Management**

Last updated : 2023-12-27 11:12:20

# Overview

Role management provides Ops and operation features at the service role level. You can restart, pause, and maintain service roles at the node level. Role status can be monitored to help you stay on top of the real-time status of role processes.

### Glossary

Restart: The nodes under the selected service role will be restarted on a rolling basis.

Pause: The nodes under the selected service role will be paused. They can be resumed by using the **Start** feature. Maintenance: The process daemon will be stopped on the nodes under the selected service role, so when a process becomes abnormal for various reasons, no alarm or automatic recovery will occur. This is suitable for node debugging. You can use the **Exit Maintenance** feature to resume the daemon.

## Directions

1. Log in to the EMR console and click the **ID/Name** of the target cluster in the cluster list to enter the cluster details page.

2. On the cluster details page, select **Cluster Service** and click **Operation** > **Role Management** in the top-right corner of the target component block. The following uses HDFS as an example.

3. The **Role Management** list displays the current service role's health status, operation status, configuration group, node type, maintenance status, server IP, and last restarted time. You can select a role and **restart, maintain, start, or pause** it.

### Note

The **Health Status** column displays the running status of the current role. The **Operation Status** column displays user operations. The **Maintenance Status** column displays whether the role is under maintenance.

Cluster Services	/ HDFS								
Service Status	Role Management	Configuration N	Management Co	onfiguration Record					
			Restart Service	Enter Maintenance Exit Maintenance	Start Pause			Enter the server IP to search	Q Nodes pending restart
			Roles T	Role Status (j)	Configuration Gro $T$	Node Type 🔻	Maintenance Status (;)	Server IP	Last Restarted
			DataNode	Information of role status, which is divided	into 3 faultGroup	Core	Normal mode	192.168.252.17	
			DataNode		nuis-oefaultGroup	Core	Normal mode	192.168.252.2	
			NameNode	Running 🥥	hdfs-defaultGroup	Master	Normal mode	192.168.252.15	
			Total 3 items					Lines per page 20 v	I /1 page → →

# Health Status

Health Status	Description
Good	Port check is responded to within 5s.
Risky	Port check is responded to within 5s–10s.
Unavailable	Port check is not responded to within 10s.
Not checked	Roles in the maintenance mode or in Stopped operation status are not checked.
Unknown	The checker is abnormal or down.

# **Client Management**

Last updated : 2023-12-27 11:50:33

# Overview

Client management provides the information of the client deployed in the component, including the configuration status, configuration group, node type, and node IP.

# Directions

1. Log in to the EMR console and click the **ID/Name** of the target cluster in the cluster list to enter the cluster details page.

2. On the cluster details page, select **Cluster Service** and click **Operation** > **Client Management** in the top-right corner of the target component block to enter the **Client Management** page.

Cluster Service /	YARN *						More 🔻   📄 Help Docur
Service Status	Job Query	Role management	Client management	Configuration Management	Resource Scheduling		
Client			Configuration state	us T	Configuration Group	Node type 🗡	Node IP
YARN			Synced		yarn-core-defaultGroup	core	
YARN			Synced		yarn-master-defaultGroup	master	
YARN			Synced		yarn-core-defaultGroup	core	
Total items: 3							10 🕶 / page 🛛 🕅 🔄 1 🔢 / 1 page

# Configuration Management Configuration Management

Last updated : 2023-12-27 14:12:56

## Overview

Configuration management allows you to modify key configuration parameters of popular open-source components like HDFS, YARN, Hive, and Spark. The configuration of services can be modified as needed at the level of cluster, node, or configuration group. This document describes how to configure service parameters in the console.

#### Caution

For security reasons, when a custom configuration file is deleted in **Configuration Management** in the console, it will not be deleted from the client.

If you need to delete a custom configuration file from the client, you can use a cluster script to batch perform the operation on existing clusters.

After configuration modification, if the configuration status in **Role Management** changes to **Configuration expired**, you need to restart the service to make the configuration effective.

## Directions

### Editing a configuration item

1. Log in to the EMR console, select the target cluster in the cluster list, and click **Service** to enter the cluster service list.

2. In the cluster service list, select **Operation** > **Configuration Management** in the top-right corner of the target service panel.

Cluster Services						
	Add Component Restart Servi	ce Reset Native UI Password				
	ZOOKEEPER Running   Version3.5.5	Operation v	HDFS Running   Version2.8.5	Operation 👻	YARN Running   Version2.8.5	Operatio
	WebUI Address 🕦	Role Man	WebUI Address ()		WebUI Address ()	
	KNOX Running   Version1.2.0	View port				
	WebUI Address 🕦					

3. On the **Configuration Management** page, select the **Level** as needed. By default, **Cluster level** is selected.


To configure the parameters of all nodes of the selected service in the cluster, select Cluster level.

To configure the parameters of multiple specified nodes of the selected service, create a configuration group and use it as the level.

To configure the parameters of a specified node of the selected service, select **Node level**.

Cluster Services	ZOOKEEPER				
Service Status	Role Management	Configuration Management Con	iguration Record		
		Level ⑦ Cluster level Cluster level Sever level Sever level configuration standard to	voup such as ∞ & when you modify configuration items on the console, the console will not escape these characters. To ensure that the special characters	s are processed properly,	please follow the XM
		Configuration File	File ZOOKEEPER: log4; properties Only valid in this cluster	Restart Service	Modify Configura
		log4j.properties	Logsdir /data/emr/zookeeper/logs		
		zkEnv.sh			
		zoo.cfg			

4. You can search for a configuration item or narrow down the search scope with the filter on the left.

5. Select the target configuration file, click **Edit Configuration**, and perform operations such as adding, editing, and deleting configuration items as needed.

Enter the new parameter value for the target parameter. If needed, you can click **Recover** to recover a parameter to its original value. You can also click **Default Value** to restore the parameter to the default value recommended by the system.

You can delete certain configuration items by clicking **Delete** > **OK**.

If the file does not have the parameter you want to configure, click **Add a configuration item** and enter the new parameter name and value in the pop-up window.

6. After confirming that the modification is correct, click **Save configuration**. After the configuration is delivered successfully, click **Restart Service**.

#### Note

If the service process configuration is modified, the configuration will not take effect unless you restart the service. If the client configuration is modified, the configuration will take effect immediately after being saved.

#### Adding a configuration file

1. Log in to the EMR console, select the target cluster in the cluster list, and click **Service** to enter the cluster service list.

2. In the cluster service list, select **Operation** > **Configuration Management** in the top-right corner of the target service panel.

3. If the configuration file you want to configure does not exist, you can click **Add configuration file** on the right to configure a configuration file.



4. After you click **Save configuration**, the parameters will be delivered, and the configuration file name will be

updated in the configuration file list.

Level () Cluster lev	el ·		
Note: if the values contain special characters such as <> & when you modify configuration items on the console, the console will not escape these characters. To ensure that the special characters are processed properly, please follow the XML standard to set configurations.			
Configuration File	File Name*         New file name         Configuration Path*         Configuration path ii		
afasdfasdf	Permission Property *       Permission property       User Group *       User group name in       Username *       Username in the opt		
log4j.properties			
zkEnv.sh			
zoo.cfg			
Create Configuration File			
+ Create Configuration File	Save configuration Cancel		

5. After the custom configuration file is delivered and takes effect, it can be modified and deleted.

#### Comparing configurations at different levels

After configurations are modified at different levels, click **Switch to dimension-based comparisons** to compare them at the cluster, configuration group, and node levels. To make the configuration at a level the same as that at the higher level, overwrite the difference and save the modification.



Cluster Service / HDFS *			More 🔻   🖽 Help Docu
Service Status Role mana	agement	Client management Configuration Management	
For configuration adjustme	nt in the cons	sole, special characters contained in values such as "<", ">", and "&" are not escaped. To ensure proper processing of special characters, please follow the XML standard to set configurations.	~
Configuration Record			⇒ Switch t
Configuration filter		All core-site.xml hadoop-env.sh hdfs-site.xml httpfs-env.sh httpfs-site.xml log4j.properties	
Enter a parameter name	Q,	Sive configuration Garcel	
Range	Clear all	File Parameter Value (cluster/configuration group/node dimension) (3)	
NameNode			
DataNode		No data	
Client			50 x / page H 4 1 /1 page
Category	Clear all		See These to a Transfer
About Address Or Port			
Advanced			
Basic			
CHdfs			
Compression			
About Cos			
About Data Store			
About Disk			
group			
About HA			
Host			
About JVM OR Memory			
About Log			
Path			
About Performance			
Pid			
Proxy			
About Security			
About Time			
uri			

# **Configuration Status**

Last updated : 2023-12-27 14:30:32

### Overview

After the component configuration is modified, saved, and delivered, you can view the configuration status to check for service configuration delivery failures or configuration expiration. You can view the configuration failure and expiration details to understand the modified configuration items of each role.

### Specification items

Synced: No configuration item is modified, or the service is restarted after the modification.

Configuration expired: The service is not restarted after the configuration item is modified and delivered, which means the configuration does not take effect.

Configuration failed: The modified configuration item fails to be delivered to the node configuration.

### Status

1. A role instance has three configuration statuses: **Synced**, **Configuration failed**, and **Configuration expired**. The initial configuration status is **Synced**. If the role process configuration fails to be delivered, the configuration status of the role will become **Configuration failed**.

After the configuration is delivered successfully, the configuration status of the role will become **Configuration expired**.

In the absence of a configuration failure, the configuration status of the role will become **Synced** after the service is restarted successfully.

2. A client has two configuration statuses: **Synced** and **Configuration failed**.

The initial configuration status is **Synced**. If the client configuration or role process configuration fails to be delivered, the configuration status of the client will become **Configuration failed**.

After the configuration is delivered successfully, the configuration status of the client will become **Synced**.

3. A component has two configuration statuses: **Configuration failed** and **Configuration expired**.

The initial status is not displayed. If the role instance or client fails to be configured, the configuration status of the component will become **Configuration failed**.

If the role instance configuration expires, the configuration status of the component will become **Configuration expired**.

If the role instance or client has both failed configuration and expired configuration, the configuration status of the component will be displayed as **Configuration failed**.

### Directions

1. Log in to the EMR console and click the **ID/Name** of the target cluster in the cluster list to enter the cluster details page.

2. On the cluster details page, select **Cluster Service** and click the component block whose configuration status is **Configuration failed** or **Configuration expired**.

3. Here takes **Configuration expired** as an example. On the component details page, click the configuration status icon in the top-right corner to view the expired configuration details.

Cluster Service / HDFS *	More 🔻   🗐 Help Docu
Service Status Role management Client management Configuration Management	Configurations ex Learn more >
There are expired configurations. Click to view change details and restart the service.	Configurations fail Learn more >
For configuration adjustment in the console, special characters contained in values such as 1 < ; >>; and 1%1 are not escaped. To ensure proper processing of special characters, please follow the XML standard to set configurations.	

4. Confirm the target expired configuration and restart the service.

# **Configuration Rollback**

Last updated : 2023-12-27 14:30:49

### Overview

This document describes how to roll back operations on component configuration items, such as adding, modifying, and deleting configuration items, in the Elastic MapReduce (EMR) console.

### Directions

Log in to the EMR console and click ID/Name of the target cluster in the cluster list to go to the cluster details page.
 On the cluster details page, click Cluster Service and choose Operation > Configuration Management in the top-right corner of the target component block. The following uses HDFS as an example:

← emr-ldjnqr4a	Cluster Service		Help Documentation
Basic Information	To receive or forward data in batches, it is recommended to use this service with	<u>CKafka</u>	
Instance Information     Bootstrap Actions	Add Component Restart Service Reset Native UI Password		
Cluster Service Cluster Resource	ZOOKEEPER ⊘ Running   Version 3.6.1	HDFS Operation • Operation •	YARN ⊘ Running   Version2.8.5
Cluster Monitoring	WebUI URL ①	WebUI URL () may also a solution and a solution process and the	WebUI URL () http://www.analysia.com
Auto Scaling Cluster Scripts	HIVE Operation • Operation •	IMPALA Ø Running   Version2.10.0	KNOX Operatio Operatio
Operation Logs	WebUI URL ()	WebUI URL 🖉 🛶 🖬 🛯 🖉 🔤 🔤 🖉	WebUI URL 🕕 🛛
	TEZ Operation   Operation		
	WebUI URL 🛞		

3. Select **Configuration Record** on the **Configuration Management** page to display the configuration record of the current component by default. You can also switch to view the configuration record of all components. Click **Details** to view the parameter values before and after the configuration change. Click **Rollback** > **Confirm rollback** to roll back the configuration change of a parameter. After the rollback is successful, the component will be restarted, and the previous configuration will take effect in a short time.

#### Note

Added, modified, or deleted configuration items can be rolled back, while added or deleted configuration files cannot.

Elastic MapReduce	← emr-	and the second second	) / HDFS			
🗉 Cluster	Role Management	Configuration Management	Configuration Record	]		
⑦ Monitor	Component Name	Configuration File	Level	Modified At	Reason for Change	Operation
18 Hardware	HDFS	core-site.xml	Cluster	2020-02-27 18:50:04	No relevant remarks	Details Rollback
Components	Total 1 item				Lines per page 10 💌 🕅	1 /1 page ▶

# **Configuration Group Management**

Last updated : 2023-12-27 14:31:16

### Overview

A configuration group is used for group-based configuration management of nodes that have different specifications or purposes but are deployed with the same component. This document describes how to use a configuration group to manage service configuration.

A configuration item modified at the cluster level will overwrite that at the configuration group and node levels if you have not separately modified it at the configuration group and node levels.

A configuration item modified at the configuration group level will overwrite that at the node level within the configuration group if you have not separately modified it at the node level. After the modification, the configuration at the cluster level will not overwrite the configuration at the configuration group level.

A configuration item modified at the node level will be updated only for the node. After the modification, the configuration delivered from the cluster and configuration group levels will not overwrite the configuration at the node level.

#### Note

A node can belong to only one configuration group, and nodes in different groups are different.

### Directions

Log in to the EMR console and click ID/Name of the target cluster in the cluster list to go to the cluster details page.
 On the cluster details page, click Cluster services and select Operation > Configurations in the top-right corner of the target component block. The following uses HDFS as an example:

3. Click **Configurations**, set **Level** to **Configuration group**, select the target configuration group, and click **Manage configuration groups**.

#### Note

After a cluster is created, several default configuration groups based on node type will be generated for each component, which can only be modified but not deleted and will inherit the cluster configuration upon initialization. You can add configuration groups as needed, with a maximum of 15 configuration groups allowed for each component.



Configuration filter	All <b>core-site.xml</b> hadoop-env.sh hdfs-site.xml httpfs-env.sh httpfs-site.xml log4j.properties	
Enter a parameter name Q	Restart Service Edit configuration	
Dimension ①	Parameter Value	Description
Configuration Group v	delegation.token.identifier.serializ	
hdfs-core-defaultGroup 🔻	ation.version	
Manage Configuration Group	delegation token identifier transm ission version	

4. On the **Manage configuration groups** page, you can view, add, modify, and delete configuration groups for the current component.

#### Note

You can select an inherited configuration group for a new configuration group. If no inherited configuration group is selected, the cluster-level configuration is inherited by default.

After a custom configuration group is deleted, nodes in it will be transferred to the default configuration group for the same node type.

Role Mar	nagement	Configuration Management	Configuration Record
Level 🕐	Configuration	Group 🔻	
hdfs-defaultGroup   Manage Configuration Group			

5. If the configuration of a configuration group is different from that of the cluster, click **Save** at the cluster level to overwrite the configuration group configuration. Similarly, if the configuration of a node is different from that of the configuration group to which the node belongs, click **Save** at the configuration group level to overwrite the node configuration.

# Service List

Last updated : 2024-10-30 10:40:57

The service list displays the services installed on the cluster, along with their health status, configuration status, version information, and more. It also provides convenient service Ops and management tools, including general service operations as well as service-specific command-based Ops operations.

### Service Health Status

The service health status indicates whether the current service is operating normally. It is aggregated from the health statuses of various roles.

Cluster Service					Help Documentation
Add Component Reset WebUI Pa	assword				
ZOOKEEPER Version 3.5.5	Operation ▼	HDFS Version 2.8.5   WebUI URL (	Operation 🔻	Version 2.8.5   WebUI URL ()	Operation
⊘ HIVE Version 2.3.5   WebUI URL ③	Operation ▼	C KNOX Version 1.2.0	Operation *		

There are four main types of service health statuses: Good, At Risk, Unavailable, and Unknown or Undetected. Each status type is displayed with a different color.

Service Health Status	Status Description	Status Aggregation Rule
Green: good.	The service is running normally.	All role instances have a health status of Good.
Orange: At risk.	The service is available, but some role instances are unavailable or at risk and need attention.	Some instances of a specific role within this component are either unavailable or at risk. For example, HDFS has one NameNode role instance



		and two DataNode role instances. Among them, one DataNode instance is unavailable, while the other DataNode instance and the NameNode instance are healthy, resulting in an At Risk health status for HDFS.
Red: unavailable.	The service is unavailable. All instances of a certain role are in an unhealthy status and require immediate handling.	All instances of a particular role in this component are in an unhealthy status. For example, HDFS has one NameNode role instance and two DataNode role instances. Among them, both DataNode instances are unavailable, while the NameNode instance is healthy, resulting in Unavailable health status for HDFS.
Gray: unknown or undetected.	The service health status is unknown or has not been detected. Components without processes are marked as Undetected for health status. For components with processes, if they enter maintenance mode or their operation status has stopped, they are also marked as Undetected. If the health status information of a role instance cannot be correctly obtained, it is marked as Unknown. If there are no issues found with the business, no further attention is needed.	<ol> <li>All role instances of this component are neither at risk nor unavailable, and at least one role instance has an unknown health status. For example, HDFS has one NameNode role instance and two DataNode role instances, with one DataNode role instance having an unknown health status, while the other DataNode role instance and the NameNode role instance have a good health status. Therefore, the health status of HDFS is unknown.</li> <li>And all role instances of this service have an undetected health status. When all role instances of the service enter maintenance mode or have stopped operation, their health status will not be detected.</li> <li>If the component has no processes, its health status will not be detected, such as Iceberg, Hudi, and Flink.</li> </ol>

### Service Operations

Common service operations include service restart, start, pause, enter/exit maintenance mode, and view ports. Command-based service operations include HDFS NameNode primary/secondary switch, HDFS data balancing, YARN ResourceManager primary/secondary switch, and YARN queue refresh. The operation descriptions are as follows:

Service Operation	Description
HDFS NameNode	Abbreviated as NN primary/secondary switch, it switches the current Active
primary/secondary	NameNode to StandBy status and the previously StandBy NameNode to Active
switch	status.



HDFS data balancing	Usually executed when new DataNodes are added. This operation distributes data evenly, avoiding hotspot issues and ensuring a more balanced read/write load across the cluster.
HDFS management status switch	Only supports switching the DataNode to Maintenance Status (IN_MAINTENANCE). This feature is typically used when a DataNode needs to be temporarily taken offline without migrating data. Currently, this feature is supported in Hadoop 3.x and later versions. For detailed directions, see Best Practices for HDFS DataNode Maintenance Status Switching.
YARN ResourceManager primary/secondary switch	Abbreviated as RM primary/secondary switch, it switches the current Active ResourceManager to StandBy status and the previously StandBy ResourceManager to Active status. RM primary/secondary switch is only allowed when yarn.resourcemanager.ha.automatic-failover.enabled is disabled. If the RM primary/secondary switch option is not displayed in the YARN card operations dropdown list, locate yarn.resourcemanager.ha.automatic-failover.enabled in YARN Configuration Management - yarn-site.xml and disable it.
YARN queue refresh	<ul> <li>When new content is added or existing content is updated in capacity-scheduler.xml or fair-scheduler.xml, this operation allows the changes to take effect in ResourceManager.</li> <li>Note: Do not delete active queues defined in capacity-scheduler.xml or fair-scheduler.xml.</li> </ul>
Using Ranger for modifying metabase	When changing the underlying database of Ranger, you need to modify the conf/install.properties file and execute the setup.sh script locally. This operation provides a one-click metabase configuration feature, preventing service issues due to incomplete configuration changes when modifying the Ranger metabase address. This operation currently only supports MySQL databases, and the Test Connection feature is only for testing the connection of the administrator user. This operation synchronizes the database information to the local ranger-admin-site.xml configuration file, but it will not update the content of ranger-admin-site.xml in the configuration management. If the user modifies and publishes ranger-admin-site.xml through the configuration management page due to additional requirements, it may result in the database information being overwritten, causing exceptions.

To perform a service operation, follow these steps:

1. Log in to the EMR console and click the corresponding Cluster ID/Name in the cluster list to enter the cluster details page.

2. On the cluster details page, select **Cluster Services**, and then select the component card you want to operate on.

3. For example, to perform an HDFS NN primary/secondary switch, select **HDFS component card operation > NN Failover** in the Cluster Services to proceed with the operation.

### Service Pause Methods List

The supported pause methods for each service component are as follows:

Component	Service	Pause Method	Description	Remarks
	NameNode	Quick pause	Directly stops service	-
LIDES	DataNode	Quick pause	Directly stops service	-
HDF3	JournalNode	Quick pause	Directly stops service	-
	zkfc	Quick pause	Directly stops service	-
	ResourceManager	Quick pause	Directly stops service	-
	NodeManager	Quick pause	Directly stops service	-
TANN	JobHistoryServer	Quick pause	Directly stops service	-
	TimeLineServer	Quick pause	Directly stops service	-
	HbaseThrift	Quick pause	Directly stops service	-
	HMaster	Quick pause	Directly stops service	-
HBASE	RegionServer	Quick pause	Directly stops service	-
	RegionServer	Safe pause	Before RegionServer is stopped, the regions on it will be migrated.	Supports setting thread concurrency.
	HiveMetaStore	Quick pause	Directly stops service	-
HIVE	HiveServer2	Quick pause	Directly stops service	-
	HiveWebHcat	Quick pause	Directly stops service	-
PRESTO	PrestoCoordinator	Quick pause	Directly stops service	-
THEOTO	PrestoWorker	Quick pause	Directly stops service	-
ZOOKEEPER	QuorumPeerMain	Quick pause	Directly stops service	-
SPARK	SparkJobHistoryServer	Quick pause	Directly stops service	-
HUE	Hue	Quick pause	Directly stops service	-

OOZIE	Oozie	Quick pause	Directly stops service	-
	Nimbus	Quick pause	Directly stops service	-
CTODM	Supervisor	Quick pause	Directly stops service	-
STONM	Logviewer	Quick pause	Directly stops service	-
	Ui	Quick pause	Directly stops service	-
RANGER	Ranger	Quick pause	Directly stops service	-
	AlluxioMaster	Quick pause	Directly stops service	-
ALLUXIO	AlluxioWorker	Quick pause	Directly stops service	-
	Httpd	Quick pause	Directly stops service	-
GANGLIA	Gmetad	Quick pause	Directly stops service	-
	Gmond	Quick pause	Directly stops service	-

# YARN Resource Scheduling Overview

Last updated : 2024-10-21 16:06:19

### Overview

YARN resource scheduling supports interactive YARN resource queue scheduling management, which is more convenient than file-based configuration management. Currently, it supports two types of scheduling configurations: Fair Scheduler and Capacity Scheduler.

### Notes

1. Fair Scheduler is used by default. Therefore, you need to configure parameters in the fair-scheduler.xml configuration file for the YARN component. If you switch to Capacity Scheduler, configure parameters in the capacity-scheduler.xml configuration file. No matter which scheduler you use, parameter configurations must be consistent with those on the **Resource Scheduling** page.

2. After switching the scheduler or setting the policy on the **Resource Scheduling** page, you need to click **Apply** to deliver the configuration, which may restart the ResourceManager.

3. If you have enabled label-based scheduling in **Configuration Management**, the sync operation will be performed when Capacity Scheduler is enabled.

4. Subpools can be nested inside a resource pool, and the nested subpools are subject to the rules set for the parent pool.

5. If you specify to sync the modified configuration when disabling the resource scheduler, the configuration file and parameters of the scheduler in **Configuration Management** will be overwritten.

# **Configuring Fair Scheduler**

Last updated : 2024-10-21 16:07:11

### Overview

Fair Scheduler allocates resources fairly to each job on YARN based on weight.

### Definitions

Configuration Set: Specifies the resource assignment among active resource pools during a specified period. Execution Plan: Specifies when the configuration set will be active. Placement Rule: Automatically allocates jobs submitted by different users to specified resource pools. User Limits: Specifies the maximum number of applications that can be submitted by a user at a time.

### Directions

#### Creating a resource pool

1. Log in to the EMR console and click **Details** of the target Hadoop cluster in the cluster list to enter the cluster details page.

2. On the cluster details page, select **Cluster services** and click **Operation** > **Resource scheduling** in the top-right corner of the YARN component block to enter the **Resource scheduling** page.

Cluster Service / YARN *	More *	Help Docur						
Service Status Job Query Role management Client management Configuration Management Resource Scheduling								
1. After Resource Scheduler is enabled, duster configurations will be synced to the resource scheduler.     After Resource Scheduler is enabled, you will be unable to configure the fair-scheduler.aml files via the Configuration Management page.     After Resource Scheduler is disabled, you will be able to configure the fair-scheduler.aml files via the Configuration Management page.								
Resource scheduler 🕕								

3. Toggle on **Resource scheduler** and configure the scheduler.

cl + c - :										
Cluster Service	/ YARN *								Mo	re 🔻   🖃 Help Docu
Service Status	Job Query	Role management Client r	management Configuration N	lanagement Resource Sch	eduling					
<ul> <li>1. Fair Scheduler is used by default. Please configure relevant parameters in the configuration file fair scheduler.xml in "Configuration management" to keep them consistent with those on this resource scheduling page, or configure relevant parameters in the capacity-scheduler.xml if Capacity Scheduler is used.</li> <li>2. After setting policies here. dick (Apply) to update configuration file fair scheduler.xml in "Configuration management". Executing Qapity will deliver the comporting policy configurations.</li> <li>3. If you have enabled tag scheduling in "Configuration management", the settings will be synced when you enable the Capacity Scheduler.</li> </ul>										
Resource schedu	Resource scheduler 💽									
Type 🔹 🔵 Fair S	Scheduler 🚫 Capac	ity Scheduler								
Resource Po	ols Execution	Plans Placement Rules	User Limits							
Create Resou	rce Pool									
			Configuration Set default •	Configuration Set default •						
Resource Pool	Name	Weight	Min Resources		Max Resources			Scheduling Policy	Preemption Mode	Operation
			CPU	Memory	CPU	Memory	max kunning Apps			
root										Edit Create Subpool
default		-	-							Edit Clone

4. Create a resource pool for Fair Scheduler.

Select **Fair Scheduler**. On the displayed page, click **Create resource pool**. You can edit, clone, and delete an existing resource pool as well as create a subpool for it.

Create Resource Pool	
Basic Settings	
Resource Pool Name() *	Enter a resource pool name
Parent Pool	
Resource Limits	
Configuration Set	default
Weight (	- 1 +
Min Resources()	vCore Memory MB
Max Resources 🛈	vCore Memory MB
Max Concurrent Running Apps 🛈	- +
Max Share for App Master()	- 0.50 +
Scheduling Policy	
	$\bigcirc$ DDE: Dominant Decourse Epimone. This fairly schedul



#### Fields and parameters:

Field	Parameter	Description
Resource Pool Name	name	Name of the resource pool or queue. It can contain only letters, digits, hyphens, and underscores and cannot start with a hyphen or underscore.
Parent Pool	The value of `type` is `parent`.	Means that although the resource pool has no subpools, it is not a leaf node. A parent pool cannot have subpools in Hadoop v2.8 and later.
Configuration Set	None	YARN does not have this parameter, which means a collection of scheduled tasks.
Weight	weight	Percentage of resources in the parent pool. A greater weight means more resources allocated.
Min Resources	minResources	The minimum amount of resources guaranteed. When the minimum amount of resources guaranteed for a queue is not met, it has a higher priority than other queues at the same level to obtain resources.



Max Resources	maxResources	The maximum amount of resources that can be used. The amount of resources available for each queue cannot exceed this value.
Max Concurrent Running Apps	maxRunningApps	The maximum number of concurrent running applications allowed. This limitation can prevent intermediate outputs generated by excessive concurrent running map tasks from filling up the disks.
Max Share for App Master	maxAMShare	The maximum percentage of resources that can be used to run Application Master. This attribute only applies to leaf queues.
Scheduling Policy	schedulingPolicy	You can set a scheduling policy for any queue. Valid values include `Fifo`, `Fair`, and `Drf`. If the value is `Fifo` or `Fair`, only memory is taken into account in resource allocation. If `Drf`, both memory and the number of cores are taken into account.
Preemption Mode	allowPreemptionFrom	This field applies only to Hadoop v3.x and later. In v2.x, you can only configure `yarn.scheduler.fair.preemption` to set the preemption mode.
Fair Share Preemption Threshold	fairSharePreemptionThreshold	The fair share preemption threshold for the queue. If the queue waits `fairSharePreemptionTimeout` without receiving `fairSharePreemptionThreshold*fairShare` resources, it is allowed to preempt resources from other queues. If this field is not set, the queue will inherit the value from its parent queue.
Fair Share Preemption Timeout	fairSharePreemptionTimeout	Number of seconds the queue is under its fair share threshold before it will try to preempt resources from other queues. If this field is not set, the queue will inherit the value from its parent queue.
Min Share Preemption Timeout	minSharePreemptionTimeout	Number of seconds the queue is under its minimum share before it will try to preempt resources from other queues. If this field is not set, the queue will inherit the value from its parent queue.
Submission	aclSubmitApps	List of users that can submit apps to the queue
Management	aclAdministerApps	List of users that can manage the queue

### Configuring an execution plan

1. In the **Policy settings** section, click **Execution plans** > **Create execution plan** to create an execution plan.

**Configuration set status** indicates whether the execution plan is enabled, which defaults to "Enabled". When you don't want to enable an execution plan but still want to retain the configuration set, you can set the plan status to "Disabled".

2. In the Create execution plan window, select/enter the Configuration set, Name, and Plan validity.

Create Execution	n Plan		
Configuration Set	O New O Existing		
Name *	Enter a name	Copy From	default 💌
Plan Validity 🚯 *	Daily •	00:00:00	③ - 00:00:00 ⑤

#### Sample configuration set

1. Log in to the EMR console and click **Details** of the target Hadoop cluster in the cluster list to enter the cluster details page.

2. On the cluster details page, select **Cluster services** and click **Operation** > **Resource scheduling** in the top-right corner of the YARN component block to go to the **Resource scheduling** page.

- 3. Toggle on **Resource scheduler** and select **Fair Scheduler**.
- 4. Click **Create resource pool** and configure parameters as needed.

Resource scheduler C Reset									
Type * 🔵 Fair Scheduler 🔷 Capa	icity Scheduler								
Resource Pools Executio	Resource Pools Execution Plans Placement Rules User Limits								
Create Resource Pool		Configuration Set							
Resource Pool Name	Weight	Min Resources	Min Resources		Max Resources		Scheduling Policy	Preemption Mode	Operation
		CPU	Memory	CPU	Memory	wax kunning Apps			
root		-		-	-				Edit Create Subpoo
flink	4						drf		Edit Clone Create Subpool Dele
hive	4					100	fair		Edit Clone Create Subpool Dele
default	2			50	150	10	fair		Edit Clone



5. On the **Resource Scheduling** page, click **Execution Plans** > **Create Execution Plan** and adjust the **Plan** 

#### Validity as needed.

#### Note

If scheduled scale-out is configured for the EMR cluster, we recommend you set the **Plan Validity** to a time period after the scheduled scale-out.

Create Execution	Plan		>
Configuration Set	• New • Existing		
Name *	08to12	Copy From default 💌	
Plan Validity(i) *	Daily <b>•</b>	00:00:00 🕓 - 12:00:00 🔇	
		Confirm Cancel	

6. On the **Resource Scheduling** page, click **Resource Pools** and select an item from the **Configuration Set** dropdown list.

#### Note

The resource limits of a resource pool can be configured by configuration set.

Resource scheduler 💽 Reset										Reset Schedul
Type * O Fair Scheduler Capacity Scheduler										
	Resource Pools Execution Plans Placement Rules User Limits									
	Create Resource Pool									
			Configuration Set default •							
	Resource Pool Name	Weight	Min Resources 08to12 default		Max Resources		Max Running Apps	icheduling Policy	Preemption Mode	Operation
			CPU	Memory	CPU	Memory				
	root									Edit Create Subpool
	flink	4	-	-	-		-	drf	-	Edit Clone Create Subpool Dele
	hive	4	-	-	-	-	100	fair	-	Edit Clone Create Subpool Dele
	default	2	-		50	150	10	fair	-	Edit Clone

7. Select the created resource pool from **Configuration Set** and adjust the **Resource Limits** as needed.

Resource Pools Execution Plans Placement Rules User Limits									
		Configuration Set 08to12							
Resource Pool Name	Weight	Min Resources CPU Memory		Max Resources		Max Running Apps	Scheduling Policy	Preemption Mode	Operation
root	-		-		-		-		Edit Create Subpoo
flink	2			-			drf		Edit Clone Create Subpool Dele
hive	2		-	-		100	fair		Edit Clone Create Subpool Dele
default	2			50	150	10	fair		Edit Clone

资源池 计划模式 彭	<b>资源地</b> 计划模式 欲重规则 用户限制									
新建资源油	新建设资源									
		配置集 20to24 👻								
资源池名称	权重	最小资源数		最大资源数		导查办工资经济本Application和目	调度策略	抢占模式	操作	
		CPU	内存	CPU	内存	WHOM 2 YES 2010 (b) (concernent)				
root	-	-	-	-	-	-	-	-	编辑 新建子池	
flink	8	-		-	-	-	dıf		编辑 克隆 新建子池	
hive	1					100	fair		编辑 克隆 新建子油	
default	1	-		50	150	10	fair		编辑 克隆	

资源池 计划模式 邡	置规则 用户限制									
新建资源地	新建党政治									
		EE旗集 default 🔻								
资源池名称 权重	权重	最小资源数		最大资源数		最高处于运行扩充Application数目	调度策略	抢占模式	操作	
		CPU	内存	CPU	内存					
root	-	-			-				编辑 新建子池	
flink	4	-	-		-	-	drf		编辑 克隆 新建子池	
hive	4	-	-		-	100	fair	-	编辑 克隆 新建子池	
default	2			50	150	10	fair		编辑 克隆	

8. After the resource pool adjustment, click **Apply** to make the settings effective.

#### Configuring a placement rule

In the **Policy settings** section, click **Placement rules** > **Create placement rule** to create a placement rule.

Cluster Service	/ YARN *						,	More 🔻   📰	Help Docu
Service Status	Job Query	Role management	Client management	Configuration Management	Resource Scheduling				
<ol> <li>1. Fair Scholl, 2. After set</li> <li>3. If you had</li> </ol>	<ol> <li>1. Fair Scheduler is used by default. Please configure relevant parameters in the configuration management. Securing (Apply) will deliver the consistent with those on this resource scheduling page, or configure relevant parameters in the capacity-scheduler sumd.</li> <li>2. After setting policies here, click (Apply) to update configuration file fair-schedulers in "Configuration management". Executing (Apply) will deliver the corresponding policy configurations.</li> <li>3. If you have enabled tag scheduling in "Configuration management", the settings will be spreed when you enable the Capacity Scheduler.</li> </ol>								
Resource schedu	ler 🚺						Apply	Reset	Scheduli
Type 🔹 🔵 Fair :	Scheduler 🔷 Caj	vacity Scheduler							
Resource Poo	ols Executio	n Plans Placement	Rules User Limits						
Create Place	nent Rule								
Placemer	rt Rules					Execution Order			Operat
						No data yet			

#### Set Placement type and Pool name.

新建放置规则							
放置类型 *	请选择放置类型	Ŧ					
池名称*	请输入池名称						
	在池不存在时创建池。						
	确定取消						

**Configuration rule type description:root.[pool name]**: The rule always needs to be met and applies when other rules do not match. Therefore, it needs to be placed after others.

**root.[pool name].[username]**: The rule checks whether the pool name exists in the resource pool, and if so, a resource pool with the same name as the username will be created (when **Create one if the pool does not exist** is selected).

**root.[primary group]**: The rule uses the resource pool matching the user's primary group. On Linux, the user's default primary group has the same name as the username. The names of the user's primary group and resource pool are used for match.

**root.[primary group].[username]**: The rule first uses the resource pool matching the user's primary group and then the subpool matching the username. If the selected pool does not exist but a pool is created, a subpool with the same name as the username will be created under the pool.

**root.[secondarygroup]**: The rule matches the user's secondary groups and uses the resource pool matching one of the secondary groups.

**root.[secondarygroup].[username]**: The rule first matches the user's secondary groups and then the resource pool matching the username.

root.[username]: The rule matches the resource pool with the same name as the username (not recommended).

Specified at runtime: The rule mainly uses the resource pool specified at runtime.

Placement rules are matched in the order of 1, 2, 3... If a rule is matched, rules after it will not be matched.

#### **Configuring user limits**

In the **Policy settings** section, click **User limits** > **Create user limit** to create a user limit.

Cluster Service / YARN *		More 🔻   📃 Help Docu				
Service Status Job Query Role management Client management Configuration Management	Resource Scheduling					
I. Fair Scheduler is used by default. Please configure relevant parameters in the configuration file fair-scheduler and in "Configuration 2. After setting policies here, click (Apply) to update configuration files and parameters in "Configuration management". Executing 3. If you have enabled tag scheduling in "Configuration management", the settings will be spriced when you enable the Capacity Sci 2. Science 2.	on management' to keep them consistent with those on this resource scheduling page, or configure relevant parameters in the capacity-scheduler.xml if Capacity Scheduler is used. (Apply) will deliver the corresponding policy configurations. Beduler.					
Resource scheduler 💽						
Type - <b>O</b> Fair Scheduler Capacity Scheduler						
Resource Pools Execution Plans Placement Rules User Limits						
Create User Limit Default 🖋						
Username	Max Concurrent Running Apps	Operat				
	No data yet					

Set Username and Max Concurrent Running Apps.

新建用户限制	
用户名称*	请输入用户名称
同时运行应用程序上限③*	请输入同时运行应用程序上限
	确定取消

# **Configuring Capacity Scheduler**

Last updated : 2023-12-27 14:32:16

### Overview

Capacity Scheduler organizes resources in a hierarchical manner, allowing multiple users to share cluster resources based on multi-level resource restrictions.

### Directions

#### Creating a resource pool

1. Log in to the EMR console and click **Details** of the target Hadoop cluster in the cluster list to go to the cluster details page.

2. On the cluster details page, select **Cluster Service** and choose **Operation** > **Resource Scheduling** in the topright corner of the YARN component block to go to the **Resource Scheduling** page.

Cluster Service	/ YARN *									More 🔻   📃 Help Docur
Service Status	Job Query	Role management	Client management	Configuration Management	Resource Scheduling					
<ol> <li>1. After Re</li> <li>2. After Re</li> <li>3. After Re</li> </ol>	1. After Resource Scheduler is enabled, cluster configurations will be synced to the resource scheduler.     After Resource Scheduler is enabled, you will be unable to configure the fair-scheduler.ml and capacity-scheduler.ml files via the Configuration Management page.     After Resource Scheduler is disabled, you will be able to configure the fair-scheduler.ml files via the Configuration Management page.									
Resource schedu	ler 🔵									

3. Toggle on **Resource scheduler** and configure the scheduler.

4. Create a resource pool for Capacity Scheduler.

Select **Capacity Scheduler** and click **Create Resource Pool** to create a resource pool. You can edit or clone an existing resource pool as well as create a subpool for it. You can also click **Default settings** to set the number of times of delayed scheduling.

Cluster Service /	YARN *										Mor	re 🔻   🖽 Help Doct
Service Status	Job Query	Role mana	agement	Client manag	gement Co	onfiguration Management	Resource Scheduling	_				
<ol> <li>1. Fair Schec</li> <li>2. After setti</li> <li>3. If you hav</li> </ol>	I. Fair Scheduler is used by default. Please configure relevant parameters in the configuration management. To keep them consistent with those on this resource scheduling page, or configure relevant parameters in the capacity-scheduler is used. 2. After setting policies here, click (Apply) to update configuration files and parameters in "Configuration management". Executing (Apply) will deliver the corresponding policy configurations. 3. If you have enabled tag scheduling in "Configuration management", the setting solid be synced when you enable the Capacity Scheduler.											
Resource schedule	source scheduler 💽											
Type 🔹 🔿 Fair Sc	Type * 🔿 Fair Scheduler 💿 Capacity Scheduler											
Resource Poo	s Resource	e Pool Mappi	ings									
Create Resour	ce Pool De	fault settings										
Tag scheduling												
Resource Pool I	Name	Capacity	Max C	Min User Ca	User Resour	Max vCores per Container			Max Memory per Cor	tainer	Resource Pool	Operation
root												Edit Create Subp
default		100%	100%		1							Edit Clone

Create Resource Pool				×			
Basic Settings							
Resource Pool Name (	Enter a resource pool name						
Resource Limits							
Tag settings	Tag	Capacity	Max Capacity	Op			
	<default_partition th="" •<=""><th>Enter %</th><th>Enter %</th><th></th></default_partition>	Enter %	Enter %				
Min User Capacity 🕄	Enter %						
User Resource Factor 🛈	Enter %						
Max Memory per Container 🛈	Enter MB						
Max vCores per Container 🛈	Enter						
Resource Pool Status 🛈	<b>O</b> RUNNING OSTOPPED						
App Limits							
Max Apps	<b>-</b> 10000 <b>+</b>						
Max Resources for AM 🛈	25 %						

Resource pool priority 🛈	Enter
Access control	
Submission 🛈	<ul> <li>Allow all users to submit jobs to this resource pool</li> <li>Allow the following users to submit jobs to this resource pool</li> </ul>
Management	<ul> <li>Allow all users to manage this pool</li> <li>Allow the following users to manage this pool</li> </ul>
	<b>Confirm</b> Cancel

#### Fields and parameters:

Field Name	Parameter	Description
Resource Pool Name	yarn.scheduler.capacity. <queue- path&gt;.queues</queue- 	Name of the resource pool or queue
Label Settings	N/A	Set the specified label that can be accessed by the queue.
Capacity	yarn.scheduler.capacity. <queue- path&gt;.capacity</queue- 	Available resource amount. The total capacity of the subpools of a parent pool is 100. Available resource amount = resource amount of the parent pool * percentage set here. The queue can consume more resources than the queue's capacity if there are idle resources in other queues.
Max Capacity	yarn.scheduler.capacity. <queue-path>.maximum- capacity</queue-path>	Maximum queue capacity in percentage. Because of resource sharing, the amount of resources used by a queue may exceed its capacity, and this field specifies the maximum amount of resources that can be used by the queue.
Default Label Expression	yarn.scheduler.capacity. <queue-path>.default-node- label-expressionpath&gt;</queue-path>	If a resource request does not have a node label specified, the application will be submitted to the corresponding partition specified by this configuration item. By default, the value is empty, i.e., applications will be allocated to containers on nodes with no label.
Min User Capacity	yarn.scheduler.capacity. <queue-path>.minimum-user- limit-percent</queue-path>	Minimum resources in percentage guaranteed for each user. Each queue enforces a limit on the percentage of resources allocated to a user at any given time. When multiple users' applications are running in a queue concurrently, the amount



		of resources used by each user varies between a minimum and maximum value. The minimum value depends on the number of running applications, and the maximum value is determined by `minimum-user-limit-percent`.
User Resource Factor	yarn.scheduler.capacity. <queue-path>.user-limit- factor</queue-path>	Maximum amount of resources in percentage that can be used by each user. For example, if the value is `30`, the amount of resources for each user cannot exceed 30% of the queue capacity at any given time.
Max Memory per Container	yarn.scheduler.capacity. <queue-path>.maximum- allocation-mb</queue-path>	Maximum memory that can be allocated to each container. The value will overwrite and cannot be greater than that of the system's `yarn.scheduler.maximum-allocation-mb`.
Max vCores per Container	yarn.scheduler.capacity. <queue-path>.maximum- allocation-vcorespath&gt;</queue-path>	Maximum number of CPU cores that can be allocated to each container. The value will overwrite and cannot be greater than that of the system's `yarn.scheduler.maximum- allocation-vcores`.
Resource Pool Status	yarn.scheduler.capacity. <queue-path>.statepath&gt;</queue-path>	Status of the queue. The value can be `Running` or `Stopped`. If a queue is in the `Stopped` status, new applications cannot be submitted to it or any of its subqueues.
Max Apps	yarn.scheduler.capacity. <queue-path>.maximum- applications</queue-path>	Maximum number of concurrent active (both running and pending) applications allowed in the system
Max Resources for AM	yarn.scheduler.capacity. <queue-path>.maximum-am- resource-percentpath&gt;</queue-path>	Maximum percentage of resources in the cluster which can be used to run application masters. It controls the number of concurrent active applications.
Resource Pool Priority	yarn.scheduler.capacity.root. <leaf-queue-path>.default- application-priorityqueue-path&gt;</leaf-queue-path>	Configure the priority of the resource queue, which is `0` by default. The larger the value, the higher the priority.
Submission	yarn.scheduler.capacity.root. <queue- path&gt;.acl_submit_applications</queue- 	List of users that can submit apps to the queue
Management	yarn.scheduler.capacity.root. <queue- path&gt;.acl_administer_queue</queue- 	List of users that can manage the queue
Delayed	yarn.scheduler.capacity.node-	Set the allowed number of times of delayed scheduling to



Scheduling	locality-delay	ensure the local execution of tasks. If the value is `-1`,
		delayed scheduling will be disabled.

#### Configuring resource pool mappings

1. In the **Policy Settings** section, choose **Resource Pool Mappings** > **Create Resource Pool Mapping** to create a resource pool mapping.

	Resource Pools Resource Pool Mappings						
	Create Resource Fool Mapping Overwrite Specified Queues()						
	Mapping Type	Username/Group Name		Resource Pool Name	Operation		
			No data yet				
L							

Mapping Type *	User •	
Username *	Enter a username	
	Specific users: %user	
Resource Pool Name *	Enter a resource pool name	
	%user %primary_group	
Mapping Result		

#### 2. Configure **Overwrite Specified Queues**.

This feature is disabled by default. For example, you have defined a mapped queue in **Resource Pool Mappings** and specified a queue other than the mapped queue when submitting a job; if the specified queue is default and **Overwrite Specified Queues** is enabled, the mapped queue will be used; otherwise, the specified queue will be used.

#### Sample label-based scheduling

1. Log in to the EMR console and click **Details** of the target Hadoop cluster in the cluster list to go to the cluster details page.

2. On the cluster details page, select Cluster Service and choose Operation > Resource Scheduling in the top-

right corner of the YARN component block to go to the **Resource Scheduling** page.

- 3. Toggle on **Resource scheduler** and select **Capacity Scheduler**.
- 4. Toggle on **Label-based scheduling** and click **Label management** to go to the label management\*\* page.

Resource Pool Mappings									
Create Resource Pool Defa	Create Resource Fool Default settings								
Tag scheduling O	PARTITION	,			Tag management				
Resource Pool Name	Capacity	Max C	Min User Ca	User Resour	Max vCores per Container	Max Memory per Container	Resource Pool	Operation	
root						-		Edit Create Subp	
default	100%	100%		1	-			Edit Clone	

5. Click Create Label, enter the label name, and set the label type and node to be bound to as needed.

Create tag	×
Tag name	It must be 1–128 characters long and contain letters, digits, hyphens, and u
Tag type	O Exclusive O Non-exclusive
Node to be bound	Please select
	Confirm Cancel

6. After the label is set, click **Apply**. Then, you can view and edit the resource queue of the label in the resource pool.

Resource Pools Create Resource Pool	Resource Defa	Pool Mapp ult settings	ings	/				
Tag scheduling	< DEFAULT	PARTITION:		-	Tag management			
Resource Pool Name	A1	_PARITION:	>	 	Hax vCores per Container	Max Memory per Container	Resource Pool	Operation
root				 		*		Edit Create Subp
default		100%	100%	 1		**		Edit Clone

7. On the **Resource Scheduling** page, click **Create Resource Pool** and set **Label**, **Capacity**, **Max Capacity**, and other parameters as needed.

#### Note

The capacity and maximum capacity of a resource pool can be configured by label.

Resource Pools Resource	tesource Pool Mappings							
Create Resource Pool Defa	Create Resource Ptod Default settings							
Tag scheduling	T_PARTITION:				ag management			
Resource Pool Name	Capacity	Max C	Min User Ca	User Resour	Max vCores per Container	Max Memory per Container	Resource Pool	Operation
root					-			Edit Create Subp
test	10%	10%			-	-	RUNNING	Edit Clone Create Subpool De
default	90%	100%		1		**		Edit Clone

8. After the resource pool is set, click **Apply** to submit the task to the backend.

#### Caution

Proceed with caution when restarting the ResourceManager. If you are prompted that the ResourceManager will be restarted after you click **Apply**, check whether the operation is successful in **Scheduling History** and whether the ResourceManager is healthy in **Role Management**.

# Label Management

Last updated : 2023-12-27 14:34:47

### Overview

Label management allows you to create, edit, or delete a label or bind it to a node. Adding different labels to cluster nodes helps allocate resources at a finer granularity based on Capacity Scheduler and specify the running locations of applications.

### Prerequisites

You have enabled YARN resource scheduling, switched to Capacity Scheduler, and enabled label-based scheduling.

### Directions

1. Log in to the EMR console and click **Details** of the target Hadoop cluster in the cluster list to enter the cluster details page.

2. On the cluster details page, select **Cluster Service** and click **Operation** > **Resource Scheduling** in the top-right corner of the YARN component block to enter the **Resource Scheduling** page.

- 3. Toggle on **Resource scheduler** and select **Capacity Scheduler**.
- 4. Toggle on Label-based scheduling and click Label Management to enter the Label Management page.

Resource Pools Resource	Resource Pools Resource Pool Mappings								
Greate Resource Pool De	Create Resource Pool Default settings								
Tag scheduling C	T_PARTITION				▼ Tag management				
Resource Pool Name	Capacity	Max C	Min User Ca	User Resour	Max vCores per Container	Max Memory per Container	Resource Pool	Operation	
root								Edit Create Subp	
default	100%	100%		1		-		Edit Clone	

5. On the **Label Management** page, view all labels of the cluster, create, edit, delete, or sync a label, or redirect to the web UI to view the node to which the label is bound.

Tag management	Apply Reset
Create tag Sync tag	View nodes
Tag	Operation
<default_partition></default_partition>	Edit
A1	Edit Delete
Total items: 2	10 ▼ / page 🛛 ◀ 1 / 1 page ▶ ▶

# Viewing Scheduling History

Last updated : 2023-12-27 14:35:03

### Overview

The scheduling history displays information such as operation records and task status of the resource queue configuration.

### Directions

1. Log in to the EMR console and click **Details** of the target Hadoop cluster in the cluster list to enter the cluster details page.

2. On the cluster details page, select **Cluster Service** and click **Operation** > **Resource Scheduling** in the top-right corner of the YARN component block to enter the **Resource Scheduling** page.

3. On the **Resource Scheduling** page, click **Scheduling History** to view the start time, end time, status, and operation information of a task.

-	-					
Today	Yesterday	Last 7 days	Last 30 days	2022-12-22 00:00:00	~ 2022-12-22 15:18:46	
Туре 🝸	C	peration type	Start Time	End Time	Status 🝸	Ope atio
			No d	data yet		
lotal items:	0			<b>10 💌</b> / page	<b>◄ 1</b> /1p	age 🕨 🕨

4. You can filter scheduling records by execution time and click **Details** in the **Operation** column to view detailed information.

# HBase RIT Fixing

Last updated : 2023-12-27 14:35:21

### Feature overview

 The HBase RIT fixing feature is suitable for fixing regions in the FAILED\_OPEN
 FAILED\_CLOSED
 or

 CLOSING
 status when an HBase 2.2.0 or later cluster encounters prolonged Region-In-Transition (RIT) issues.

### Directions

1. Log in to the EMR console and click a cluster ID/name in the Cluster list to enter the cluster details page.

2. On the cluster details page, click **Cluster Service** and select **Operation** > **RIT fixing** in the top-right corner of the HBase component panel to view and fix regions in the RIT status.

3. In the list, click **Fix** in the **Operation** column, or select the target regions and click **Batch fix**, and confirm the information in the **RIT fixing** pop-up window.

- 4. After confirming that everything is incorrect, click **Confirm**.
- 5. You can view the progress and result of RIT fixing in the **Task Center**.

# **Component Port Information**

Last updated : 2023-12-27 14:35:38

This document describes common component ports.

### Common HDFS ports

Component	Port Number	Description
HDFS	4007	NameNode RPC port, which is used for: 1. The communication between the HDFS client and NameNode. 2. The connection between the DataNode and NameNode.
HDFS	4008	<ul><li>HDFS HTTP port (NameNode), which is used for:</li><li>1. Point-to-point NameNode checkpoint operations.</li><li>2. Connecting the remote web client to the NameNode UI.</li></ul>
HDFS	4009	<ul><li>HDFS HTTPS port (NameNode), which is used for:</li><li>1. Point-to-point NameNode checkpoint operations.</li><li>2. Connecting the remote web client to the NameNode UI.</li></ul>
HDFS	4004	DataNode IPC server port, which is used for connecting the client to the DataNode to perform RPC operations.
HDFS	4001	DataNode data transfer port, which is used for: 1. The data transfer from or to the DataNode to or from the HDFS client. 2. Point-to-point DataNode data transfer.
HDFS	4002	DataNode HTTP port, which is used for connecting the remote web client to the DataNode UI in safe mode.
HDFS	4003	DataNode HTTPS port, which is used for connecting the remote web client to the DataNode UI in safe mode.
HDFS	4005	JournalNode RPC port, which is used for client communication to access multiple types of information.
HDFS	4006	JournalNode HTTP port, which is used for connecting the remote web client to the JournalNode in safe mode.
HDFS	4032	The listener port of the HttpFS HTTP server, which is used for connecting the remote RESTful API to HttpFS.
# Common YARN ports

Component	Port Number	Description
YARN	5024	The web HTTP port of the JobHistory server, which is used for viewing the webpage of the JobHistory server.
YARN	5022	JobHistory server port, which is used for: 1. Recovering task data on the MapReduce client. 2. Getting the task report on the JobHistory client.
YARN	5003	The web HTTP port of the ResourceManager service.
YARN	5002	The web HTTPS port of the ResourceManager service, which is used for connecting to the ResourceManager web application in safe mode.
YARN	5004	The web HTTP port of the NodeManager.
YARN	5005	The web HTTPS port of the NodeManager, which is used for connecting to the NodeManager web application in safe mode.
YARN	5001	ResourceManager scheduler port.
YARN	5000	ResourceManager RPC port, which is used for submitting tasks on the client.

# Common Hive ports

Component	Port Number	Description
Hive	7010	The port for WebHCat to provide RESTful services, which is used for the communication between the WebHCat client and WebHCat server.
Hive	7001	The port for HiveServer to provide Thrift services, which is used for the communication between the HiveServer client (Beeline) and HiveServer.
Hive	7004	The port for MetaStore to provide Thrift services, which is used for the communication between the MetaStore client and MetaStore, i.e., the communication between HiveServer and MetaStore.
Hive	7003	The web UI port of Hive, which is used for the HTTPS/HTTP communication between web requests and the Hive UI server.

#### Common Spark ports

Component	Port Number	Description
Spark	10000	HTTP port.

#### **Common Presto ports**

Component	Port Number	Description
Presto	9000	The HTTP port for the Presto coordinator or worker to provide external services.

## Common PrestoSQL ports

Component	Port Number	Description
PrestoSQL	9000	The HTTP port for the PrestoSQL coordinator or worker to provide external services.

# Common Trino ports

Component	Port Number	Description
Trino	9000	The HTTP port for the Trino coordinator or worker to provide external services.

# Common Impala ports

Component	Port Number	Description
Impala	27000 and 27009	The ports for Impala application communication. The former is for Impala 2.x, and the latter is for Impala 3.x.
Impala	27001	The port for impala-shell communication.

Impala	27004	The web port of the Impala server.
Impala	27007	The web port of the Impala Catalog.
Impala	27005	The web port of the Impala StateStore.

# Common Kudu ports

Component	Port Number	Description
Kudu	7051	The PRC port of the Kudu master.
Kudu	7050	The RPC port of the Kudu server.
Kudu	8051	The HTTP port of the Kudu master.
Kudu	8050	The HTTP port of the Kudu server.

# Common ClickHouse ports

Component	Port Number	Description
ClickHouse	9000	The TCP port for accessing the business client.
ClickHouse	8123	The HTTP port for accessing the business client.
ClickHouse	9009	The HTTPS port for accessing the business client.

# Common Kylin ports

Component	Port Number	Description
Kylin	16500	Kylin HTTP port.

# Common Doris ports

Component	Port Number	Description
-----------	----------------	-------------



Doris	8000	The Thrift server port on the broker.
Doris	9060	The Thrift server port on the backend, which is used for receiving frontend requests.
Doris	8060	The BRPC port on the backend, which is used for the communication between backend instances.
Doris	9050	The heartbeat service port (Thrift) on the backend, which is used for receiving frontend heartbeats.
Doris	8040	The Thrift server port on the broker, which is used for receiving requests.
Doris	9010	The port on the frontend for the communication between metadata management modules (BDBJE).
Doris	8030	The HTTP server port on the frontend.
Doris	9020	The Thrift server port on the frontend.
Doris	9030	The port on the frontend for receiving MySQL client access requests.

# Common StarRocks ports

Component	Port Number	Description
StarRocks	8000	The Thrift server port on the broker.
StarRocks	9060	The Thrift server port on the backend, which is used for receiving frontend requests.
StarRocks	8060	The BRPC port on the backend for the communication between backend instances.
StarRocks	9050	The heartbeat service port (Thrift) on the backend, which is used for receiving frontend heartbeats.
StarRocks	8040	The Thrift server port on the broker, which is used for receiving requests.
StarRocks	9010	The port on the frontend for the communication between metadata management modules (BDBJE).
StarRocks	8030	The HTTP server port on the frontend.



StarRocks	9020	The Thrift server port on the frontend.
StarRocks	9030	The port on the frontend for receiving MySQL client access requests.

# Common Druid ports

Component	Port Number	Description
Druid	8082	The listener port of the broker server ( broker-runtime.properties ), which is used for receiving client queries.
Druid	8081	The listener port of the coordinator server ( coordinator - runtime.properties ) for the communication with other components.
Druid	8083	The listener port of the historical server ( historical-runtime.properties ) for the communication with other components.
Druid	8091	The listener port of the middleManager server ( middleManager- runtime.properties ) for the communication with other components.
Druid	8090	The listener port of the Overlord server ( overlord-runtime.properties ) for the communication with other components.
Druid	18888	The listener port of the router server ( router-runtime.properties ), which is used for routing client queries to the broker.

# Common HBase ports

Component	Port Number	Description
HBase	6000	HMaster RPC port, which is used for connecting the HBase client to HMaster.
HBase	6001	HMaster HTTPS port, which is used for connecting the remote web client to the HMaster UI.
HBase	6002	The RPC port of the RegionServer (RS), which is used for connecting the HBase client to the RegionServer.
HBase	6003	RegionServer HTTPS port, which is used for connecting the remote web client to the RegionServer UI.



HBase	6005	The Info Server listener port of the HBase Thrift server.
HBase	6004	The HBase Thrift server listener port.

# Common Flink ports

Component	Port Number	Description
Flink	16001	Flink web UI port, which is used for the HTTP/HTTPS communication between client web requests and the Flink server.

#### Common Storm ports

Component	Port Number	Description
Storm	15002	Logviewer service port.
Storm	15000	Nimbus service port.
Storm	15001	Storm web UI port.

#### Common Hue ports

Component	Port Number	Description
Hue	13000	The port for Hue to provide HTTPS web services, which can be modified.

#### Common Oozie ports

Component	Port Number	Description
Oozie	12000	HTTP port, which is used for client access.

# Common Superset ports



Component	Port Number	Description
Superset	18088	Superset service port, which is used for client connection.

## Common Zeppelin ports

Component	Port Number	Description
Zeppelin	18000	Zeppelin service port, which is used for client connection.

# Common Kafka ports

Component	Port Number	Description
Kafka	9092	The port for the broker to provide data reception and acquisition services.

# Common Kafka Manager ports

Component	Port Number	Description
Kafka Manager	9000	The listener port of the Kafka Manager server, which is used for client connection.

# Common Ranger ports

Component	Port Number	Description
Ranger	6080	Ranger admin web UI port.
Ranger	5151	The port of Ranger UserSync for authentication services.

#### Common COSRanger ports

Component	Port Number	Description
COSRanger	9999	The listener port of the COSRangerServer server, which is used for client connection.

## Common KRB5 ports

Component	Port Number	Description
KRB5	749	Kadmin service port.
KRB5	754	Kprop service port.
KRB5	88	Kerberos server port for component authentication in the Kerberos service, which may be used in configuring mutual trust of clusters.

# Common Knox ports

Component	Port Number	Description
Knox	30002	HTTP port, which is usually used for browser connection.
Knox	33389	The LDAP port of Knox, which is used for Knox authentication.

# Common ZooKeeper ports

Component	Port Number	Description
ZooKeeper	2181	ZooKeeper client port, which is used for connecting the ZooKeeper client to the ZooKeeper server.

# Common OpenLDAP ports

Component	Port Number	Description



OpenLDAP	389	The port for client access.

# Common Tez ports

Component	Port Number	Description
Tez	2000	The web UI port of Tez.

#### Common Livy ports

Component	Port Number	Description
Livy	8998	The listener port of the Livy server, which is used for client connection.

# Common Kyuubi ports

Component	Port Number	Description
Kyuubi	10009	The listener port of the Kyuubi server, which is used for client connection.

# Common Alluxio ports

Component	Port Number	Description
Alluxio	20001	The RPC port of the Alluxio Job Master, which is used by the Alluxio Master to assign tasks to the Alluxio Job Master.
Alluxio	30001	The RPC port of the Alluxio Job Worker, which is used by the Alluxio Job Master to assign tasks to the Alluxio Job Worker.
Alluxio	19998	The RPC port of the Alluxio Master, which is used for connecting the client to the Alluxio Master.
Alluxio	29998	The RPC port of the Alluxio Worker, which is used by the Alluxio Master to assign read-write tasks to the Alluxio Worker.

# Common GooseFS ports

Component	Port Number	Description
GooseFS	9206	The RPC port of the Alluxio Job Master, which is used by the Alluxio Master to assign tasks to the Alluxio Job Master.
GooseFS	9210	The RPC port of the Alluxio Job Worker, which is used by the Alluxio Job Master to assign tasks to the Alluxio Job Worker.
GooseFS	9201	The RPC port of the Alluxio Master, which is used for connecting the client to the Alluxio Master.
GooseFS	9211	The listener port of the GooseFSProxy server, which is used for client proxy.
GooseFS	9204	The RPC port of the Alluxio Worker, which is used by the Alluxio Master to assign read-write tasks to the Alluxio Worker.

# Common Ganglia ports

Component	Port Number	Description
Ganglia	1602	Gmetad server port.
Ganglia	1603	Gmetad server port, which is used for receiving HTTPd data queries.
Ganglia	1601	The port for gmond process communication and gmetad access.

# **RSS Uniffle Operation Guide**

Last updated : 2025-07-03 15:55:13

## Feature Introduction

Shuffle, as a core operator of big data computing, provides high coverage rate but witnesses high resource consumption and prominent stability problems, as the EMR customers and the user scale of Spark continuously grow, and containerized applications become popular.

Based on such background, Tencent Cloud Elastic MapReduce (EMR) has added a new RSS cluster type and supports the Uniffle component. It allows Spark clusters of both EMR on CVM and EMR on TKE versions to bind to RSS clusters with one click, providing remote Shuffle services for Spark jobs. This aims to improve the performance of large data Shuffle tasks or enhance task stability in node recycling scenarios.

## Creating an RSS Cluster

1. Log in to the EMR console, click **Creating a Cluster** on the EMR on CVM cluster list page, and select **RSS** as the cluster type when creating the cluster. You need to configure the software, region and hardware, basic settings, and confirm the configuration information in 4 steps.

2. In the software configuration step, carefully select the region, as Shuffle services do not support cross-region operations.

3. In the region and hardware configuration step, the minimum number of nodes for a high-availability cluster is 4, including 2 master nodes and at least 2 core nodes. Non-HA clusters have only one replica stored and is suitable for testing but not recommended for production environments.

#### Note :

1. Evaluate the memory configuration, disk IO bandwidth, etc., for the new RSS cluster based on the number of parallel Spark SQL tasks and the Shuffle data volume of the tasks.

2. An RSS cluster supports shared use for all Spark clusters in the same region.

4. After completing the above 4 configuration steps, click Purchase to make the payment. Once the payment is successful, the EMR cluster will enter the creation process. You can find the newly created cluster in the EMR console in about 10 minutes.

# **Operation Steps**

#### **Using Uniffle in Spark Clusters**



#### Using Uniffle in EMR on CVM Spark Clusters

1. In the EMR on CVM cluster list, select the **cluster that needs to enable the Shuffle service**, and click **Manage** under **Service**.

2. Select Cluster Service, and in the Spark Service panel, click Operation and then Set Shuffle Service.

3. Enable the service and select the required RSS cluster ID that is to be bound.

By default, spark.shuffle.manager is disabled. If it is enabled, all submitted Spark tasks will use the RSS-Shuffle service. It is recommended to enable it at the task level. The steps are as follows:

1. Modify the spark-defaults.conf configuration file for the Spark component on the Spark configuration management page.

```
spark.shuffle.manager=org.apache.spark.shuffle.RssShuffleManager
spark.driver.extraClassPath=/usr/local/service/spark/rss-client/* (If this paramete
spark.executor.extraClassPath=/usr/local/service/spark/rss-client/* (If this parame
```

2. Spark dynamic allocation.

2.1 If you use dynamic resource allocation (spark.dynamicAllocation.enabled=true) with Spark and need to use Uniffle, you need to apply a patch on Spark according to the specific version, and submit a ticket.

2.2 Modify the spark-defaults.conf configuration file: spark.shuffle.service.enabled=false.

#### Note:

If the spark version is 3.5 or above, you need to additionally configure the parameter spark.shuffle.sort.io.plugin.class=org.apache.spark.shuffle.RssShuffleDataIo.

#### **Disabling Uniffle in EMR on CVM Spark Clusters**

1. In the Spark Service panel, click Operation, Set Shuffle Service, and click the button to disable the service.

2. Remove the Spark parameter spark.shuffle.manager=org.apache.spark.shuffle.RssShuffleManager on the configuration management page and apply the change.

#### Using Uniffle in EMR on TKE Spark Clusters

The operations are the same as those for using Uniffle in EMR on CVM spark clusters, but the paths for associating and unbinding the RSS cluster are different. The operation paths are as follows:

1. In the EMR on TKE cluster list, select the cluster that needs to enable the Shuffle service, and click the cluster ID to enter the Cluster Information page.

2. Select Cluster Service. In cluster service Spark Service, click Operation in the upper right corner, and click Set Shuffle Service.

3. Enable the service and select the required RSS cluster ID that is to be bound.

# Monitoring and Alarms Cluster Overview

Last updated : 2023-12-27 14:37:51

## Feature overview

Cluster overview displays a summary of cluster operation status, where you can view the cluster operation status, core service metrics, core node metrics, and top 10 nodes.

#### Directions

1. Log in to the EMR console and click the **ID/Name** of the target cluster in the cluster list to enter the cluster details page.

2. On the cluster details page, select **Cluster Monitor** > **Cluster Overview** to view the overall status of the current cluster. The **Cluster Overview** page provides four monitoring views in the cluster dimension, namely, overall status of cluster, key cluster metrics, cluster deployment status, and top 10 nodes.

Overall cluster conditions: This view displays the current cluster status, node quantity, metadatabase online status, component health status, and events in the cluster within one hour.



Key cluster metrics: This view displays the total utilization of HDFS, YARN, CPU, memory, and disk as well as total network traffic of the current cluster. You can click **View Trend** in the top-right corner of each metric and select a target time period to view the corresponding usage trend.







Cluster deployment status: This view displays abnormal, missing, and invalid processes deployed on nodes and the number of nodes in the current cluster, so that you can correct the configuration accordingly.

eployment Status				
Node Type	Standard process (deployed)	Standard process (missing)	Invalid Process	Nodes
MASTER	HMaster,HistoryServer,Knox,NameN ode,QuorumPeerMain,ResourceMan ager,SparkJobHistoryServer,ThriftSe rver			1
CORE	DataNode,NodeManager,RegionServ er			2

Top 10 nodes: This view displays the trend of node loads in the current cluster under core metrics. You can select multiple nodes to compare their load trends under the same metric.





# Node Status

Last updated : 2023-12-27 14:38:20

# Overview

The **Node status** page displays the monitoring overview and list of all nodes in the current cluster. It enables you to view heat maps of all nodes. You can manage the status information and metrics of nodes in the EMR console.

# Directions

1. Log in to the EMR console and click the **ID/Name** of the target cluster in the cluster list to enter the cluster details page.

2. Select **Cluster resources > Node status** to view the monitoring information of all nodes in the cluster.

3. On the Node status page, view aggregated monitoring metrics and list of all nodes in the cluster.



**Overview**: Displays aggregated monitoring metrics of all nodes during a specified period of time and the statistical rule of each metric. You can click **Set metrics** to customize the displayed metrics.

Set metrics				
Overview				
✓ CPU Usage Rate	Preview			
CPU Load(15min)	Preview			
<ul> <li>CPU Load(1min)</li> </ul>	Collapse			





**Heat map**: Displays the loads of nodes. You can view the maps for a specified period of time and under specified load conditions. Load heat maps consist of two parts: the aggregated map of all nodes in the current cluster and individual heat map of each node.





other items. You can click the IP of a node to view its basic configurations, deployment status, load status, and monitoring information.

Basic configurations: Display basic information of the node, such as type, resource name, resource ID, billing mode, and spec.

Basic configuration			
IP:			
Node type: Master Resource name: emi	-vm-hrmm178p Resource ID: ins-90oeejs2 Billing m	node: Pay as you go	
Node specification: EMR StandardSA2 CF	U: 4-core MEM: 16GB Cloud SSD 200GB * 1		
Disk	Mount point	Usage (j)	Disk read-write speed
vda1	/	68% 35.87GB/52.71GB	Read: 0 MB/s Write: 0.09 MB/s
vdb	/data	0% 0.9GB/214.64GB	Read: 0 MB/s Write: 0 MB/s

Deployment status: Display services deployed in the node, process type, and process status.

leployment status					
Standard process deployed: 10 Standard process deployed: 0 Standard process	deployed: 0				
Process name	Process type	Process status			
HMaster	Standard process	Running			
HiveMetaStore	Standard process	Running			
HiveServer2	Standard process	Running			
HiveWebHcat	Standard process	Running			
JobHistoryServer	Standard process	Running			
Knox	Standard process	Running			

Load status: Displays top N processes in the node. You can view the data on a specified date.

.oad status			
Top CPU Processe	s Top Memory Processes	Top IO Processes	Process List
top = 15:43:25 m Tasks: 150 total %Cpu(s): 1.6 u: KiB Man : 16128	p 56 days, 21:57, 0 users, , 1 running, 76 sleeping, , 1.6 sy, 0.0 ni, 96.8 id, 00 totol 2500512 from 22	load average: 0.02, 0.14 0 stopped, 0 zombie 0.0 wa, 0.0 hi, 0.0 ; 21460 wed, 4346128 buf	), 0.15 si, 0.0 st f/aasba
KiB Mem : 161281	UU total, 2500512 free, 92	281460 used, 4346128 but:	t/cache
KiB Swap:	O total, O free,	O used, 4787540 ava	il Mem
PID USER	PR NI VIRT RES SH	ERS & CPU & MEM TIME+	java
11179 hadoop	20 0.6531132.997.1m 9672	18 S 6.2 6.3 96:35.59	
1 root	20 0 129064 6668 429	6 S 0.0 0.0 20:18.39	systemd
2 root	20 0 0 0	0 S 0.0 0.0 0:04.59	kthreadd
3 root	0-20 0 0	0 I 0.0 0.0 0.00 0.00.00	rcu_gp
4 root	0-20 0 0	0 I 0.0 0.0 0.0 0:00.00	rcu_par_gp
6 root	0-20 0 0	0 I 0.0 0.0 0:00.00	kworker/0+
8 root	0-20 0 0	0 I 0.0 0.0 0:00.00	mm_percpu+
9 root	20 0 0 0	0 S 0.0 0.0 2:51.30	ksoftirqd+

Node monitoring: Displays load maps of metrics (6 by default and max 12) of the node. You can click **Set metrics** to customize the displayed metrics.



# Service Status

Last updated : 2023-12-27 14:38:36

## Overview

The **Service status** page provides detailed monitoring for main services installed in the cluster, including HDFS, YARN, Hive, ZooKeeper, Spark, HBase, and Presto. This document describes how to view cluster service status in the console.

#### Directions

1. Log in to the EMR console and click the ID/Name of the target cluster in the cluster list.

2. On the cluster details page, select **Cluster services** and click **Operation** > **Service status** in the top-right corner of the target component block. The following uses HDFS as an example.

Cluster services					🖽 Help docume
Add component Reset WebUI password					
OPENLDAP Version 2.4.44 WebUI URL ①	Operation 🔻	COOKEEPER Version 3.63 WebUI URL ①	Operation 🔻	HDFS Version 2.8.5   WebUI URL ①	Operati Restart service
⊘ YARN Version 2.8.5   WebUI URL <sup>①</sup>	Operation 🔻	KNOX Version 1.6.1 WebUI URL ③	Operation 🔻	⊘ KUDU Version 1.15.0   WebUI URL ③	Start service Pause service Enter maintena
HBASE Version 2.4.5   WebUI URL ①	Operation 🔻	Wersion 2.3.9   WebUI URL ①	Operation 🔻		Exit maintenan View port Service status File storage an
					Roles Configurations Federations Rebalance NN Failover

The service status page provides monitoring views of three service dimensions, namely, service summary, health status, and service overview. As different service components have different services, the dimensions of display vary.
 Service overview displays 6 metrics of service components by default and their statistical rules during a specified time period. You can click Set metrics to customize the displayed metrics.





et metrics		
uster Overview		
Snapshots Preview		
Block Capacity Preview		
Blocks Preview		
Cluster Storage Capacity Preview		
Cluster Data Nodes Preview		
Total Cluster Files Preview		
Cluster Load Preview		
Master/slave Conditions Preview		
Storage Usage Rate Preview		
sk		

5. Service summary displays the current overall status of the service.

6. **Health status** displays running overview of each component in the current service. Clicking a role name or running overview will redirect you to **Roles** or **Role status** page.

Cluster services / H	DFS *				More 🔻   🖽 Help document
Service status	File storage analysis Roles	Client management Configurat	tions Federations		
Service summary	10.0.0.135(Active) 10.0.0.100(Standby)	DataNode Status	3 Alive 0 Dead 0 Decommissioning	Missing Report	0 Blocks 0 BlocksWithReplicationFactorOne
Health status					Create instant inspectio
NameNode(2)	⊘ 2 Good		zkfc(2)	🕑 2 Good	
JournalNode(3)	⊘ 3 Good		DataNode(3)	🕑 3 Good	

On a role status page, you can click **Set metrics** to customize the displayed metrics.

ter services / HDFS / NameN	Node(10.0.0.135) *				🖽 Help doc
<b>asic info</b> with Status: ⊘Good ate: Active		Node IP: 10.0.0.135 NameService: HDFS1006715		JVM Memory: Heap Memory Used 292M Heap Memory	Total 4096M
Real time Last 24 hours	Last 7 days 2023-07-13 13:26:50 to 2023-0	17-13 14:26:50 回 Time granularity: 30 seconds マ Se	t metrics Role lo	g -	
Blocks (count)	[]	SinceLastCheckpoint (count)	[]	LastCheckpoint (ms) (j)	
1500		300		1689232500000	
1000		200		1689227500000	
500				168922500000	
I3:30 13:40 13:50     I3:60     I3:50     I3:50	14:00 14:10 14:20 plicated – PendingDeletionB locks – PendinoReolicatio *	U 13:30 13:40 13:50 14:00 14:10 — SinceLastCheckpol	14:20	- LastCheckpoint	14:10 14:20

7. Service overview enables you to view cluster-level metrics. You can click Set metrics to customize the displayed metrics.



#### Caution

**Service monitoring** displays the HDFS service information by default. You can manually adjust to view other service components.

Service monitoring dimensions vary by service component due to different service types. For example, HBase supports table-level monitoring.





# **Cluster Event**

Last updated : 2025-06-11 16:39:36

# Overview

Cluster events include event lists and event policies.

Event list: It records key change events and abnormal events occurring in the cluster.

Event policy: Event monitoring trigger policies can be customized based on the actual business conditions. Events with monitoring enabled can be set as cluster inspection items.

# Viewing Event List

1. Log in to the EMR console and click the **ID/Name** of the target cluster in the cluster list to enter the cluster details page.

2. On the cluster details page, select **Cluster monitoring** > **Cluster events** > **Event list** to view all operation events in the current cluster.

The severity divides into the following:

Fatal: Exception events of a node or service that require manual intervention and will cause service interruption if left unattended. Such events may last for a period of time.

Severe: Alert events that currently have not caused service or node interruption but will cause fatal events if left unattended.

Moderate: Regular events occurring in the cluster that generally do not require special processing.

3. Click the value in the **Triggers today** column to view the event's trigger records, metrics, logs, and snapshots.

# Setting Event Policies

1. Log in to the EMR console and click the ID/Name of the target cluster in the cluster list.

2. On the cluster details page, select **Cluster monitoring** > **Cluster events** > **Event policy** and you can customize the event monitoring trigger policies.

3. The event configuration list contains the event name, event trigger policy, severity (fatal, severe, and moderate), and option to enable/disable monitoring, which can be modified and saved.

4. Event trigger policies cover two types of events: fixed system policy events (which cannot be modified) and custom events (which can be configured based on the business standards).

5. You can select whether to enable event monitoring in an event policy. Only events with monitoring enabled can be selected as cluster inspection items. Monitoring is enabled by default for some events and is enabled by default and cannot be disabled for some other events. The following are the specific rules:

Category	Event Name	Description	Recommendations and Measure	Default Value
Node	The CPU utilization exceeds the threshold continuously	The server CPU utilization has been greater than or equal to m for t ( $300 \le t \le$ 2,592,000) seconds continuously	Expand the node capacity or upgrade the node	m=85, t=1,800
	The average CPU utilization exceeds the threshold	The average server CPU utilization has been greater than or equal to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Expand the node capacity or upgrade the node	m=85, t=1,800
	The average CPU iowait utilization exceeds the threshold	The average CPU iowait utilization of the server in the last t (300 $\leq$ t $\leq$ 2,592,000) seconds has been greater than or equal to m	Manually troubleshoot the issue	m=60, t=1,800
	The 1-second CPU load exceeds the threshold continuously	The 1-minute CPU load has been greater than or equal to m for t $(300 \le t \le 2,592,000)$ seconds continuously	Expand the node capacity or upgrade the node	m=8, t=1,800

The 5-second CPU load exceeds the threshold continuously	The 5-minute CPU load has been greater than or equal to m for t $(300 \le t \le 2,592,000)$ seconds continuously	Expand the node capacity or upgrade the node	m=8, t=1,800
The memory utilization exceeds the threshold continuously	The memory utilization has been greater than or equal to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Expand the node capacity or upgrade the node	m=85, t=1,800
The swap space exceeds the threshold continuously	The server swap memory has been greater than or equal to m for t ( $300 \le t \le$ 2,592,000) seconds continuously	Expand the node capacity or upgrade the node	m=0.1, t=300
The total number of system processes exceeds the threshold continuously	The total number of system processes has been greater than or equal to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Manually troubleshoot the issue	m=10,000, t=1,800
The average total number of fork subprocesses exceeds the threshold	The average total number of fork subprocesses in the last t ( $300 \le t \le$ 2,592,000) seconds has been greater than or equal to m	Manually troubleshoot the issue	m=5,000, t=1,800
No process OOM	An OOM error occurred in the process	Adjust the process heap memory size	-
A disk I/O error occurred (this event is not supported currently)	A disk I/O error occurred	Replace the disk	-
The average disk utilization exceeds the threshold continuously	The average disk space utilization has been greater than or equal to m for t (300 ≤	Expand the node capacity or upgrade the node	m=85, t=1,800

	$t \le 2,592,000)$ seconds continuously		
The average disk I/O utilization exceeds the threshold continuously	The average disk I/O utilization has been greater than or equal to m for t ( $300 \le t \le$ 2,592,000) seconds continuously	Expand the node capacity or upgrade the node	m=85, t=1,800
The node file handle utilization exceeds the threshold continuously	The node file handle utilization has been greater than or equal to m for t ( $300 \le t \le$ 2,592,000) seconds continuously	Manually troubleshoot the issue	m=85, t=1,800
The number of TCP connections to the node exceeds the threshold continuously	The number of TCP connections to the node has been greater than or equal to m for t $(300 \le t \le 2,592,000)$ seconds continuously	Check whether there are connection leaks	m=10,000, t=1,800
The configured node memory utilization exceeds the threshold	The memory utilization configured for all roles on the node exceeds the node's physical memory threshold	Adjust the allocated node process heap memory	90%
The node process is unavailable	The node service process is unavailable	View the service logs to find out why the service failed to be pulled	-
The node heartbeat is missing	The node heartbeat failed to be reported regularly	Manually troubleshoot the issue	-
The hostname is incorrect	The node's hostname is incorrect	Manually troubleshoot the issue	-
Failed to ping the metadatabase	The TencentDB instance heartbeat failed to be reported regularly	-	-

	The utilization of a single disk exceeds the threshold continuously	The single disk space utilization has been greater than or equal to m for t ( $300 \le t \le$ 2,592,000) seconds continuously	Expand the node capacity or upgrade the node	m=0.85, t=1,800
	The I/O utilization of a single disk exceeds the threshold continuously	The single disk I/O device utilization has been greater than or equal to m for t ( $300 \le$ t $\le 2,592,000$ ) seconds continuously	Expand the node capacity or upgrade the node	m=0.85, t=1,800
	The single disk inodes utilization exceeds the threshold continuously	The single disk inodes utilization has been greater than or equal to m for t ( $300 \le t \le$ 2,592,000) seconds continuously	Expand the node capacity or upgrade the node	m=0.85, t=1,800
	The difference between the UTC time and NTP time of the server exceeds the threshold	The difference between the UTC time and NTP time of the server exceeds the threshold (in ms)	<ol> <li>Make sure that the NTP daemon is running</li> <li>Make sure that the network</li> <li>communication with the NTP server is normal</li> </ol>	Difference=3000(
	Automatic node replenishment	If automatic node replenishment is enabled, when any exceptions in task and router nodes are detected, the system automatically purchases nodes of the same model to replace the affected nodes.	<ol> <li>If the replenishment is successful, no more attention is required.</li> <li>If the replenishment fails, manually terminate the affected nodes in the console and purchase new nodes to replace them.</li> </ol>	-
	Node failure	Faulty nodes exist in a cluster	Handle the failure in the console or submit a ticket to contact us.	-

HDFS	The total number of HDFS files exceeds the threshold continuously	The total number of files in the cluster has been greater than or equal to m for t ( $300 \le$ t $\le 2,592,000$ ) seconds continuously	Increase the NameNode memory	m=50,000,000, t=1,800
	The average total number of HDFS files exceeds the threshold	The average total number of files in the cluster in the last t $(300 \le t \le 2,592,000)$ seconds has been greater than or equal to m	Increase the NameNode memory	m=50,000,000, t=1,800
	The total number of HDFS blocks exceeds the threshold continuously	The total number of blocks in the cluster has been greater than or equal to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Increase the NameNode memory or the block size	m=50,000,000, t=1,800
	The average total number of HDFS blocks exceeds the threshold	The average total number of HDFS blocks in the last t (300 $\leq$ t $\leq$ 2,592,000) seconds has been greater than or equal to m	Increase the NameNode memory or the block size	m=50,000,000, t=1,800
	The number of HDFS data nodes marked as dead exceeds the threshold continuously	The number of data nodes marked as dead has been greater than or equal to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Manually troubleshoot the issue	m=1, t=1,800
	The HDFS storage space utilization exceeds the threshold continuously	The HDFS storage space utilization has been greater than or equal to m for t ( $300 \le$ t $\le 2,592,000$ ) seconds continuously	Clear files in HDFS or expand the cluster capacity	m=85, t=1,800
	The average HDFS storage space	The average HDFS storage space	Clear files in HDFS or expand the cluster	m=85, t=1,800



utilization exceeds the threshold	zation exceeds the sholdutilization has been greater than or equal to m for t $(300 \le t \le 2,592,000)$ seconds continuously		
Active/Standby NameNodes were switched	Active/Standby NameNodes were switched	Locate the cause of NameNode switch	-
The NameNode RPC request processing latency exceeds the threshold continuously	The RPC request processing latency has been greater than or equal to m milliseconds for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Manually troubleshoot the issue	m=300, t=300
The number of current NameNode connections exceeds the threshold continuously	The number of current NameNode connections has been greater than or equal to m for t ( $300 \le t \le$ 2,592,000) seconds continuously	Manually troubleshoot the issue	m=2,000, t=1,800
A full GC event occurred on a NameNode	A full GC event occurred on a NameNode	Fine-tune the parameter settings	-
The NameNode JVM memory utilization exceeds the threshold continuously	The NameNode JVM memory utilization has been greater than or equal to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Adjust the NameNode heap memory size	m=85, t=1,800
The DataNode RPC request processing latency exceeds the threshold continuously	The RPC request processing latency has been greater than or equal to m milliseconds for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Manually troubleshoot the issue	m=300, t=300
The number of current DataNode connections	The number of current DataNode connections	Manually troubleshoot the issue	m=2,000, t=1,800



eds the threshold nuously	has been greater than or equal to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously		
GC event red on a Node	A full GC event occurred on a NameNode	Fine-tune the parameter settings	-
DataNode JVM bry utilization eds the threshold huously	The NameNode JVM memory utilization has been greater than or equal to m for t ( $300 \le$ t $\le 2,592,000$ ) seconds continuously	Adjust the DataNode heap memory size	m=85, t=1,800
NameNodes of S are in Standby se status	Both NameNode roles are in Standby status at the same time	Manually troubleshoot the issue	-
umber of HDFS ng blocks eds the threshold	The number of missing blocks in the cluster has been greater than or equal to m for t (300 $\leq$ t $\leq$ 604,800) seconds continuously	We recommend you troubleshoot HDFS data block corruption and run the hadoop fsck / command to check the HDFS file distribution	m=1, t=1,800
IDFS NameNode ed the safe mode	The NameNode entered the safe mode (for 300 seconds continuously)	We recommend you troubleshoot HDFS data block corruption and run the hadoop fsck / command to check the HDFS file distribution	-
umber of ntly missing Managers in the er exceeds the nold continuously	The number of currently missing NodeManagers in the cluster has been greater than or equal to m for t ( $300 \le t \le$ 2,592,000) seconds continuously	Check the NodeManager process status and check whether the network connection is smooth	m=1, t=1,800
	eds the threshold nuously GC event red on a Node DataNode JVM ory utilization eds the threshold nuously NameNodes of S are in Standby se status Umber of HDFS ng blocks eds the threshold DFS NameNode ed the safe mode d the safe mode	adds the threshold huouslyhas been greater than or equal to m for t (300 $\leq t \leq 2,592,000$ ) seconds continuouslyGC event red on a NodeA full GC event occurred on a NameNodeDataNode JVM ory utilization hods the threshold nuouslyThe NameNode JVM memory utilization has been greater than or equal to m for t (300 $\leq$ t $\leq 2,592,000$ ) seconds continuouslyNameNodes of B are in Standby se statusBoth NameNode roles are in Standby status at the same timeumber of HDFS ng blocks do the thresholdThe number of missing blocks in the cluster has been greater than or equal to m for t (300 $\leq t \leq 604,800$ ) seconds continuouslyIDFS NameNode ed the safe mode ad the safe modeThe NameNode entered the safe mode (for 300 seconds continuously)umber of huby missing Managers in the r exceeds the hold continuouslyThe number of currently missing NodeManagers in the cluster has been greater than or equal to m for t (300 $\leq t \leq$ 2,592,000) seconds continuously	add the threshold nuouslyhas been greater than or equal to m for t (300 $\le t \le 2,592,000$ ) seconds continuouslyFine-tune the parameter settingsGC event red on a lodeA full GC event occurred on a NameNodeFine-tune the parameter settingsAbataNode JVM bataNode JVM memory utilization rds the threshold nuouslyThe NameNode JVM memory utilization has been greater than or equal to m for t (300 $\le$ $t \le 2,592,000$ ) seconds continuouslyAdjust the DataNode heap memory sizeNameNodes of to are in Standby set statusBoth NameNode roles are in Standby status at the same timeManually troubleshoot the issueThe number of missing blocks in the cluster has been greater than or equal to m for t (300 $\le t \le 604,800$ )We recommend you troubleshoot HDFS data block corruption and run the hadoop fsck / command to check the HDFS file distributionHDFS NameNode ed the safe mode dot he safe mode of the safe mode of the safe mode ontinuously)We recommend you troubleshoot HDFS data block corruption and run the hadoop fsck / command to check the HDFS file distributionHDFS NameNode ed the safe mode if the safe mode of the safe mode if the safe mode if neared the safe mode (for 300 seconds continuously)We recommend you troubleshoot HDFS data block corruption and run the hadoop fsck / command to check the HDFS file distributionumber of ntity missing Managers in the receeds the nold continuouslyThe number of currently missing NodeManagers in the custer has been greater than or equal to m for t (300 $\le t \le$ 2,592,000) sec
The number of pending containers exceeds the threshold continuously	The number of pending containers has been greater than or equal to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Reasonably specify resources that can be used by YARN jobs	m=90, t=1,800
--	---	--	---------------------
The cluster memory utilization exceeds the threshold continuously	The memory utilization has been greater than or equal to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Scale out the cluster	m=85, t=1,800
The average cluster memory utilization exceeds the threshold	The average memory utilization in the last t $(300 \le t \le 2,592,000)$ seconds has been greater than or equal to m	Scale out the cluster	m=85, t=1,800
The cluster CPU utilization exceeds the threshold continuously	The CPU utilization has been greater than or equal to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Scale out the cluster	m=85, t=1,800
The average cluster CPU utilization exceeds the threshold	The average CPU utilization in the last t $(300 \le t \le 2,592,000)$ seconds has been greater than or equal to m	Scale out the cluster	m=85, t=1,800
The number of available CPU cores in each queue is below the threshold continuously.	The number of available CPU cores in each queue has been less than or equal to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Assign more resources to the queue	m=1, t=1,800
The available memory in each queue is below the threshold continuously	The available memory in each queue has been less than or equal to m for t (300 ≤	Assign more resources to the queue	m=1,024, t=1,800

		t ≤ 2,592,000) seconds continuously		
	Active/Standby ResourceManagers were switched	Active/Standby ResourceManagers were switched	Check the ResourceManager process status and view the standby ResourceManager logs to locate the cause of active/standby switch	_
	A full GC event occurred in a ResourceManager	A full GC event occurred in a ResourceManager	Fine-tune the parameter settings	-
	The ResourceManager JVM memory utilization exceeds the threshold continuously	The ResourceManager JVM memory utilization has been greater than or equal to m for t ( $300 \le t \le$ 2,592,000) seconds continuously	Adjust the ResourceManager heap memory size	m=85, t=1,800
	A full GC event occurred in a NodeManager	A full GC event occurred in a NodeManager	Fine-tune the parameter settings	-
	The available memory in NodeManager is below the threshold continuously	The available memory in a single NodeManager has been less than or equal to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Adjust the NodeManager heap memory size	m=1, t=1,800
	The NodeManager JVM memory utilization exceeds the threshold continuously	The NodeManager JVM memory utilization has been greater than or equal to m for t ( $300 \le t \le$ 2,592,000) seconds continuously	Adjust the NodeManager heap memory size	m=85, t=1,800
HBase	The number of regions in RIT status in the	The number of regions in RIT status in the	If the HBase version is below 2.0, run	m=1, t=60

cluster exceeds the threshold continuously	cluster has been greater than or equal to m for t ( $300 \le t \le$ 2,592,000) seconds continuously	hbase hbck - fixAssigment	
The number of dead RegionServers exceeds the threshold continuously	The number of dead RegionServers has been greater than or equal to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Manually troubleshoot the issue	m=1, t=300
The average number of regions in each RegionServer in the cluster exceeds the threshold continuously	The average number of regions in each RegionServer in the cluster has been greater than or equal to m for t ( $300 \le t \le$ 2,592,000) seconds continuously	Expand the node capacity or upgrade the node	m=300, t=1,800
A full GC event occurred on HMaster	A full GC event occurred on HMaster	Fine-tune the parameter settings	m=5, t=300
The HMaster JVM memory utilization exceeds the threshold continuously	The HMaster JVM memory utilization has been greater than or equal to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Adjust the HMaster heap memory size	m=85, t=1800
The number of current HMaster connections exceeds the threshold continuously	The number of current HMaster connections has been greater than or equal to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Manually troubleshoot the issue	m=1000, t=1800
A full GC event occurred in RegionServer	A full GC event occurred in RegionServer	Fine-tune the parameter settings	m=5, t=300
The RegionServer JVM memory utilization exceeds the threshold continuously	The RegionServer JVM memory utilization has been greater than or equal	Adjust the RegionServer heap memory size	m=85, t=1800



		to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously		
	The number of current RPC connections to RegionServer exceeds the threshold continuously	The number of current RPC connections to RegionServer has been greater than or equal to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Manually troubleshoot the issue	m=1000, t=1800
	The number of RegionServer StoreFiles exceeds the threshold continuously	The number of RegionServer StoreFiles has been greater than or equal to m for t ( $300 \le t \le$ 2,592,000) seconds continuously	Run the major compaction	m=50,000, t=1,800
	A full GC event occurred in HBase Thrift	A full GC event occurred in HBase Thrift	Fine-tune the parameter settings	m=5, t=300
	The HBase Thrift JVM memory utilization exceeds the threshold continuously	The HBase Thrift JVM memory utilization has been greater than or equal to m for t (300 $\leq$ t $\leq$ 2,592,000) seconds continuously	Adjust the HBase Thrift heap memory size	m=85, t=1,800
	Both HMasters of HBase is in Standby service status	Both HMaster roles are in Standby status at the same time	Manually troubleshoot the issue	-
Hive	A full GC event occurred in HiveServer2	A full GC event occurred in HiveServer2	Fine-tune the parameter settings	m=5, t=300
	The HiveServer2 JVM memory utilization exceeds the threshold continuously	The HiveServer2 JVM memory utilization has been greater than or equal to m for t ( $300 \le$ t $\le 2,592,000$ ) seconds continuously	Adjust the HiveServer2 heap memory size	m=85, t=1,800
	A full GC event	A full GC event	Fine-tune the	m=5, t=300



	occurred in HiveMetaStore	occurred in HiveMetaStore	parameter settings	
	A full GC event occurred in HiveWebHcat	A full GC event occurred in HiveWebHcat	Fine-tune the parameter settings	m=5, t=300
ZooKeeper	The number of ZooKeeper connections exceeds the threshold continuously	The number of ZooKeeper connections has been greater than or equal to m for t ( $300 \le t \le$ 2,592,000) seconds continuously	Manually troubleshoot the issue	m=65,535, t=1,800
	The number of ZNodes exceeds the threshold continuously	The number of ZNodes has been greater than or equal to m for t ( $300 \le t \le$ 2,592,000) seconds continuously	Manually troubleshoot the issue	m=2,000, t=1,800
Impala	The ImpalaCatalog JVM memory utilization exceeds the threshold continuously	The ImpalaCatalog JVM memory utilization has been greater than or equal to m for t ( $300 \le t \le$ 604,800) seconds continuously	Adjust the ImpalaCatalog heap memory size	m=0.85, t=1,800
	The Impala daemon JVM memory utilization exceeds the threshold continuously	The Impala daemon JVM memory utilization has been greater than or equal to m for t ( $300 \le t \le$ 604,800) seconds continuously	Adjust the Impala daemon heap memory size	m=0.85, t=1,800
	The number of Impala Beeswax API client connections exceeds the threshold	The number of Impala Beeswax API client connections has been greater than or equal to m	Adjust the value of fs_sevice_threads in the impalad.flgs configuration in the console	m=64, t=120
	The number of Impala HiveServer2 client	The number of Impala HiveServer2 client	Adjust the value of fs_sevice_threads	m=64, t=120



	connections exceeds the threshold	connections has been greater than or equal to m	in the impalad.flgs configuration in the console	
	The query execution duration exceeds the threshold	The query execution duration exceeds m seconds	Manually troubleshoot the issue	-
	The total number of failed queries exceeds the threshold	The total number of failed queries has been greater than or equal to m for t seconds ( $300 \le t \le$ 604,800)	Manually troubleshoot the issue	m=1, t=300
	The total number of committed queries exceeds the threshold	The total number of committed queries has been greater than or equal to m for t seconds ( $300 \le t \le$ 604,800)	Manually troubleshoot the issue	m=1, t=300
	The query execution failure rate exceeds the threshold	The query execution failure rate has been greater than or equal to m for t seconds (300 $\leq$ t $\leq$ 604,800)	Manually troubleshoot the issue	m=1, t=300
PrestoSQL	The current number of failed PrestoSQL nodes exceeds the threshold continuously	The current number of failed PrestoSQL nodes has been greater than or equal to m for t ( $300 \le t \le 604,800$ ) seconds continuously	Manually troubleshoot the issue	m=1, t=1,800
	The number of queuing resources in the current PrestoSQL resource group exceeds the threshold continuously	The number of queuing tasks in the PrestoSQL resource group has been greater than or equal to m for t ( $300 \le t \le$ 604,800) seconds continuously	Fine-tune the parameter settings	m=5,000, t=1,800
	The number of failed	The number of failed	Manually troubleshoot	m=1, t=1,800

	PrestoSQL queries exceeds the threshold	PrestoSQL queries is greater than or equal to m	the issue	
	A full GC event occurred in a PrestoSQLCoordinator	A full GC event occurred in a PrestoSQLCoordinator	Fine-tune the parameter settings	-
	The PrestoSQLCoordinator JVM memory utilization exceeds the threshold continuously	The PrestoSQLCoordinator JVM memory utilization has been greater than or equal to m for t ( $300 \le t \le$ 604,800) seconds continuously	Adjust the PrestoSQLCoordinator heap memory size	m=0.85, t=1,800
	A full GC event occurred on a PrestoSQL worker	A full GC event occurred on a PrestoSQL worker	Fine-tune the parameter settings	-
	The PrestoSQLWorker JVM memory utilization exceeds the threshold continuously	The PrestoSQLWorker JVM memory utilization has been greater than or equal to m for t ( $300 \le t \le$ 604,800) seconds continuously	Adjust the PrestoSQLWorker heap memory size	m=0.85, t=1,800
Presto	The current number of failed Presto nodes exceeds the threshold continuously	The current number of failed Presto nodes has been greater than or equal to m for t (300 $\leq$ t $\leq$ 604,800) seconds continuously	Manually troubleshoot the issue	m=1, t=1,800
	The number of queuing resources in the current Presto resource group exceeds the threshold continuously	The number of queuing tasks in the Presto resource group has been greater than or equal to m for t (300 $\leq$ t $\leq$ 604,800) seconds continuously	Fine-tune the parameter settings	m=5,000, t=1,800
	The number of failed Presto queries exceeds the threshold	The number of failed Presto queries is	Manually troubleshoot the issue	m=1, t=1,800

		greater than or equal to m		
	A full GC event occurred on a PrestoSQL coordinator	A full GC event occurred on a PrestoSQL coordinator	Fine-tune the parameter settings	-
	The Presto coordinator JVM memory utilization exceeds the threshold continuously	The Presto coordinator JVM memory utilization has been greater than or equal to m for t ( $300 \le t \le$ 604,800) seconds continuously	Adjust the Presto coordinator heap memory size	m=0.85, t=1,800
	A full GC event occurred on a Presto worker	A full GC event occurred on a Presto worker	Fine-tune the parameter settings	-
	The Presto worker JVM memory utilization exceeds the threshold continuously	The Presto worker JVM memory utilization has been greater than or equal to m for t ( $300 \le t \le$ 604,800) seconds continuously	Adjust the Presto worker heap memory size	m=0.85, t=1,800
Alluxio	The current total number of Alluxio workers is below the threshold continuously	The current total number of Alluxio workers has been smaller than or equal to m for t ( $300 \le t \le$ 604,800) seconds continuously	Manually troubleshoot the issue	m=1, t=1,800
	The utilization of the capacity on all tiers of the current Alluxio worker exceeds the threshold	The utilization of the capacity on all tiers of the current Alluxio worker has been greater than or equal to the threshold for t $(300 \le t \le 604,800)$ seconds continuously	Fine-tune the parameter settings	m=0.85, t=1,800
	A full GC event occurred on an Alluxio	A full GC event occurred on an Alluxio	Manually troubleshoot the issue	-

	master	master		
	The Alluxio master JVM memory utilization exceeds the threshold continuously	The Alluxio master JVM memory utilization has been greater than or equal to m for t ( $300 \le t \le$ 604,800) seconds continuously	Adjust the Alluxio worker heap memory size	m=0.85, t=1,800
	A full GC event occurred on an Alluxio worker	A full GC event occurred on an Alluxio worker	Manually troubleshoot the issue	-
	The Alluxio worker JVM memory utilization exceeds the threshold continuously	The Alluxio worker JVM memory utilization has been greater than or equal to m for t ( $300 \le t \le$ 604,800) seconds continuously	Adjust the Alluxio master heap memory size	m=0.85, t=1,800
kudu	The cluster replica skew exceeds the threshold	The cluster replica skew has been greater than or equal to the threshold for t ( $300 \le t$ $\le 3,600$ ) seconds continuously	Run the rebalance command to balance the replicas	m=100, t=300
	The hybrid clock error exceeds the threshold	The hybrid clock error has been greater than or equal to the threshold for t ( $300 \le t$ $\le 3,600$ ) seconds continuously	Make sure that the NTP daemon is running and the network communication with the NTP server is normal	m=5,000,000, t=300
	The number of running tablets exceeds the threshold	The number of running tablets has been greater than or equal to m for t ( $300 \le t \le$ 3,600) seconds continuously	Too many tablets on a node can affect the performance. We recommend you clear unnecessary tables and partitions or expand the capacity as needed.	m=1,000, t=300
	The number of failed	The number of failed	Check whether any	m=1, t=300

tablets exceeds the threshold	tablets has been greater than or equal to m for t ( $300 \le t \le$ 3,600) seconds continuously	disk is unavailable or data file is corrupted	
The number of failed data directories exceeds the threshold	The number of failed data directories has been greater than or equal to m for t (300 $\leq$ t $\leq$ 3,600) seconds continuously	Check whether the path configured in the fs_data_dirs parameter is available	m=1, t=300
The number of full data directories exceeds the threshold	The number of full data directories has been greater than or equal to m for t ( $120 \le t \le$ 3,600) seconds continuously	Clear unnecessary data files or expand the capacity as needed	m=1, t=120
The number of write requests rejected due to queue overloading exceeds the threshold	The number of write requests rejected due to queue overloading has been greater than or equal to m for t (300 $\leq$ t $\leq$ 3,600) seconds continuously	Check whether the number of write hotspots or worker threads is small	m=10, t=300
The number of expired scanners exceeds the threshold	The number of expired scanners has been greater than or equal to m for t ( $300 \le t \le$ 3,600) seconds continuously	Be sure to call the method for closing a scanner after reading data	m=100, t=300
The number of error logs exceeds the threshold	The number of error logs has been greater than or equal to m for t $(300 \le t \le 3,600)$ seconds continuously	Manually troubleshoot the issue	m=10, t=300
The number of RPC requests that timed out while waiting in the queue exceeds the threshold	The number of RPC requests that timed out while waiting in the queue has been greater than or equal	Check whether the system load is too high	m=100, t=300



		to m for t (300 $\leq$ t $\leq$ 3,600) seconds continuously		
Kerberos	The Kerberos response time exceeds the threshold	The Kerberos response time has been greater than or equal to m (ms) for t $(300 \le t \le 604,800)$ seconds continuously	Manually troubleshoot the issue	m=100, t=1,800
Cluster	The auto scaling policy has failed	<ol> <li>The scale-out rule failed due to insufficient subnet EIPs associated with the cluster.</li> <li>The scale-out rule failed due to insufficient expansion resource inventory of the preset specifications.</li> <li>The remaining resource quota of the current account is insufficient, causing the scale-out rule execution to fail.</li> <li>The scale-out rule failed due to insufficient account balance.</li> <li>The cluster process conflicts.</li> <li>An internal error occurred.</li> </ol>	<ol> <li>Switch to another subnet in the same VPC.</li> <li>Switch to specifications of resources that are sufficient or submit a ticket to contact developers.</li> <li>It is recommended to adjust resource quotas or replace with models that have sufficient quotas to ensure sufficient remaining quotas for the account.</li> <li>Top up the account to ensure that the account balance is sufficient.</li> <li>Try again later.</li> <li>Submit a ticket to contact developers.</li> </ol>	
	The execution of the auto scaling policy has timed out	<ol> <li>Scaling cannot be performed temporarily as the cluster is in the cooldown period.</li> <li>Scaling is not triggered because the retry period upon expiration is too short.</li> <li>The cluster in the</li> </ol>	<ol> <li>Adjust the cooldown period for the scaling rule.</li> <li>Extend the retry period upon expiration.</li> <li>Try again later or submit a ticket to contact developers.</li> </ol>	-



	current status cannot be scaled out.	
The auto scaling policy is not triggered	<ol> <li>The scale-out rule cannot be triggered because no expansion resource specification is set.</li> <li>The scale-out rule cannot be triggered because the maximum number of nodes for elastic resources is reached.</li> <li>The scale-in rule cannot be triggered because the minimum number of nodes for elastic resources is reached.</li> <li>The time range for scaling has expired.</li> <li>The scale-in rule cannot be triggered because there are no elastic resources in the cluster.</li> </ol>	<ol> <li>Set at least one elastic resource specification for the rule.</li> <li>Modify the maximum number of nodes to continue scaling out if the upper limit is reached.</li> <li>Modify the minimum number of nodes to continue scaling in if the lower limit is reached.</li> <li>Modify the effective time range of the rule if you want to continue using the rule for auto scaling.</li> <li>Execute the scale-in rule after adding elastic resources.</li> </ol>
Auto scaling partially succeeded	<ol> <li>Only partial resources were supplemented because the resource inventory was less than the required quantity for scale-out.</li> <li>Only partial resources were supplemented because the required quantity for scale-out exceeded the actual quantity of resources delivered.</li> <li>The scale-out rule was partially</li> </ol>	<ol> <li>Use the available resources for manual scaling to supplement the resources for auto scaling.</li> <li>Use the available resources for manual scaling to supplement the resources for auto scaling.</li> <li>Modify the maximum number of nodes to continue scaling out if the upper limit is reached.</li> <li>Modify the minimum number of nodes to</li> </ol>

	the maximum number of nodes for elastic resources was reached. 4. The scale-in rule was partially successful because the minimum number of nodes for elastic resources was reached. 5. The resource supplement failed due to insufficient subnet EIPs associated with the cluster. 6. The resource supplement failed due to insufficient expansion resource inventory of the preset specifications. 7. The resource supplement failed due to insufficient account balance.	the lower limit is reached. 5. Switch to another subnet in the same VPC. 6. Switch to specifications of resources that are sufficient or submit a ticket to contact developers. 7. Top up the account to ensure that the account balance is sufficient.	
The node process is unavailable	The node process is unavailable	Manually troubleshoot the issue	-
The process is killed by OOMKiller	The process OOM is killed by OOMKiller	Adjust the process heap memory size	-
A JVM or OLD exception occurred	A JVM or OLD exception occurred	Manually troubleshoot the issue	<ol> <li>The OLD utilization reaches 80% for 5 consecutive minutes Or</li> <li>The JVM memory utilization reaches 90%</li> </ol>
Timeout of service role health status occurred	The service role health status has timed out	The service role health status has timed out	t=300



	for t seconds (180 ≤ t ≤ 604,800)	for minutes. To resolve this issue, check the logs of the corresponding service role and perform necessary actions accordingly.	
A service role health status exception occurred	The service role health status has been abnormal for t seconds $(180 \le t \le 604,800)$	The service role health status has been unavailable for minutes. To resolve this issue, check the logs of the corresponding service role and perform necessary actions accordingly.	t=300
Auto scaling failed	This alert indicates that the auto scaling process has failed (either completely or partially)	Manually troubleshoot the issue	-

## Log

Last updated : 2024-01-10 09:56:47

### Overview

The log search feature collects component operation logs. Then, you can search for core service logs of the current cluster and node system logs by keyword to quickly view key service logs without logging in to nodes.

### Directions

1. Log in to the EMR console and click the ID/Name of the target cluster in the cluster list.

2. On the cluster details page, choose **Cluster Monitoring** > **Log Search** to filter and view log content by current cluster, log file, node IP, and time range.

Alternatively, choose **Cluster Service** > target component block > **Role Management**, click a **Node IP** in **Role List** to enter the page where the node monitoring metrics are displayed, and then click **Role Log** to enter the **Log Search** page.

Cluster Service	/ YARN / ResourceMar	ager(1					🖂 Help Documer
Basic info Health Status: ( State: Standby	∂Good		Node IP:			JVM Memory: Heap Memory Used 38	IM Heap Memory Total 2000M
Real time	Last 24 hours Last 7	lays 2023-05-23 18:46:19	to 2023-05-23 19:46:19 🗎	Time granularity: 30 seconds 💌	Set Metrics Ro	le Log	
MemHea	pUsedM area percentage	(%) (i) []	GC Count (j		0	GC Time (s) (j)	5
15			2.0			0.020	

Click a **Node IP** to enter the **Node Status** page, or click **Log Source** to enter the page where the node monitoring metrics are displayed.

Log Search							🖾 Help Documer
Time	Real time Last 24 hours	s Last 7 days	2023-05-16 19:47:50 to 2023-05-23 19:4	7:50 🖽			
Keywords 🚯	Enter a log keyword						
Log Source	ResourceManager 😣	▼ Node IP	▼ Min	mum log level 🚯	WARN 🔻		
Query							
Node IP	Log Source	Time ↓	Log Level	Log content			
				No data ye	t		

#### Keyword description:

Supports full-text search based on keywords.

Supports search by using special symbols, including hyphen (-), period (.), asterisk (\*), greater-than sign (>), less-than sign (<), equal sign (=), exclamation mark (!), brackets (() and {}), and forward slash (/).

Supports phrase-based search. Example: address=/ip:port .

3. During troubleshooting, you usually need to pay attention to contextual logs of keywords. On the **Log Search** page, click **View Context** to enter the **Log Context** page.

## Service Types Supported for Log Search

#### Note:

Currently, only logs in the past 15 days can be searched for.

If log collection has not been enabled on your cluster, you can contact your Tencent Cloud rep for assistance.

Component	Role	Log	Description
	NameNode	/data/emr/hdfs/logs/hadoop-hadoop- namenode.log	NameNode runtime log
	ZKFC	/data/emr/hdfs/logs/hadoop-hadoop-zkfc.log	ZKFC runtime log
HDFS	DataNode	/data/emr/hdfs/logs/hadoop-hadoop-datanode.log	DataNode runtime log
	JournalNode	/data/emr/hdfs/logs/hadoop-hadoop- journalnode.log	JournalNode runtime log
	DFSRouter	/data/emr/hdfs/logs/hadoop-hadoop-dfsrouter.log	DFRouter runtime log



YARN	ResourceManager	/data/emr/yarn/logs/yarn-hadoop- resourcemanager.log	ResourceManager runtime log
	NodeManager	/data/emr/yarn/logs/yarn-hadoop- nodemanager.log	NodeManager runtime log
	JobHistoryServer	/data/emr/yarn/logs/mapred-hadoop- historyserver.log	JobHistoryServer runtime log
	HMaster	/data/emr/hbase/logs/hbase-hadoop- master.log	HMaster runtime log
HBase	ThriftServer	/data/emr/hbase/logs/hbase-hadoop-thrift.log	ThriftServer runtime log
	RegionServer	/data/emr/hbase/logs/hbase-hadoop- regionserver.log	RegionServer runtime log
ClickHouse	ClickHouse-server	/data/clickhouse/clickhouse- server/logs/clickhouse-server.log	ClickHouse-server runtime log
	Broker	/data/emr/druid/var/log/druid/broker.log	Broker runtime log
	Coordinator	/data/emr/druid/var/log/druid/coordinator.log	Coordinator runtime log
	Router	/data/emr/druid/var/log/druid/router.log	Router runtime log
Druid	Overload	/data/emr/druid/var/log/druid/overload.log	Overload runtime log
	Historical	/data/emr/druid/var/log/druid/historical.log	Historical runtime log
	MiddleManager	/data/emr/druid/var/log/druid/middleManager.log	MiddleManager runtime log
Zookeeper	Zookeeper	/data/emr/zookeeper/logs/zookeeper-root- server.log	ZooKeeper runtime log
Hive	HiveServer2	/data/emr/hive/logs/hadoop-hive	HiveServer2 runtime log
Kudu	KuduMaster	/data/emr/kudu/logs/kudu-master.WARNING	KuduMaster runtime log
	KuduServer	/data/emr/kudu/logs/kudu-tserver.WARNING	KuduServer

Elastic MapReduce



			runtime log
Alluxio	AlluxioMaster	/data/emr/alluxio/logs/master.log	AlluxioMaster runtime log
Alluxio	AlluxioWorker	/data/emr/alluxio/logs/worker.log	AlluxioWorker runtime log
Ranger	EmbeddedServer	/data/emr/ranger/logs/ranger-admin.log	EmbeddedServer runtime log
CosRanger	CosRangerServer	/usr/local/service/cosranger/log/info.log	COSRanger runtime log
	Catalogd	/data/emr/impala/logs/catalogd.INFO	Cataloged runtime log
Impala	Statestored	/data/emr/impala/logs/statestored.INFO	Statestored runtime log
	Impalad	/data/emr/impala/logs/impalad.INFO	Impalad runtime log
Spark	HistoryServer	/data/emr/spark/logs/spark-hadoop.log	HistoryServer runtime log
Kylin	Kylin	/data/emr/kylin/logs/kylin.log	Kylin runtime log
Zeppelin	ZeppelinServer	/data/emr/zepplin/logs/zeppelin-hadoop.log	ZeppelinServer runtime log
Knox	Gateway	/data/emr/knox/logs/gateway.log	Gateway runtime log
	BrokerBootstrap	/data/emr/doris/broker/log/apache_hdfs_broker.log	BrokerBootstrap runtime log
Doris	PaloFe	/data/emr/doris/fe/log/fe.log	PaloFe runtime log
	PaloBe	/data/emr/doris/be/log/be.INFO	PaloBe runtime log
Kafka	Kafka	/user/local/service/kafka/logs/server.log	Kafka runtime log

## Minimum Log Level Supported by Services

Service	Default Minimum Log Level
Impala and Kudu	INFO
Other services	WARN

## Query Rules for the Minimum Log Level

Minimum Log Level	Queryable Log Levels
INFO	INFO, WARN, ERROR, FATAL
WARN	WARN, ERROR, FATAL
ERROR	ERROR, FATAL
FATAL	FATAL

# Application Analysis StarRocks Query Management

Last updated : 2024-10-30 10:42:26

## Feature Introduction

Provides multi-dimensional metrics and detailed displays for StarRocks queries. The query list allows quick viewing of the query statement, query start time, query status, query duration, user, number of rows retrieved, total CPU time, total memory consumed, and other detailed metrics. You can also view related SQL, query plan, profile, and execution topology.

#### Directions

1. Log in to the EMR Console, and click the corresponding StarRocks **Cluster ID/Name** in the cluster list to enter the cluster details page.

2. On the cluster details page, click **Cluster Services**, then select **Operations > Query Management** in the topright corner of the StarRocks component to view the related query list. Some column headers support filtering or sorting, and multi-dimensional composite filtering is available. You can also customize the query duration to filter for slow queries.

3. Click Operation **Operation Column > Details** in the far-right column of the query list to view the query statement, query plan, profile, execution DAG, and operator order.

#### Note:

StarRocks Query Management currently supports only StarRocks V1.4.0 and later. For StarRocks V1.4.0 clusters purchased after September 12, 2023, the query management feature will be enabled by default, and an emr\_admin user will be created for system query information collection. For StarRocks V1.4.0 clusters created before September 12, 2023, Submit a Ticket to contact support for assessment and activation of StarRocks Query Management.

Before a query is executed, enter set enable\_profile=true in the command to print the profile of the current session's SQL query. Additionally, the collection service will automatically gather the profile of the submitted SQL in that session.

The visualization feature for execution DAG and operator order in StarRocks Query Management is currently supported only in StarRocks V3.1.8 and later versions. This feature requires Profile collection to be enabled.

To ensure stable operation of cluster services, queries with an execution time of less than 200 ms will be downgraded and ignored. Parameters related to the query collection downgrade policy can be modified by

Submitting a Ticket.

# Hive Data Table Analysis

Last updated : 2024-10-30 10:43:13

### Feature Introduction

Provides data distribution and trend information based on the storage volume of databases and data tables. Provides the distribution of data tables based on their last access time, serving as a reference for hot and cold data distribution.

The proportion of small files at the data table level and the data volume at the partition level within tables can help identify issues such as small files and partition data skew.

#### Directions

1. Log in to the EMR console, and click the corresponding **Cluster ID/Name** in the cluster list to enter the cluster details page.

2. On the cluster details page, click **Cluster Services**, then select Hive component and select **Operations > Data** 

**Table Analysis** in the top-right corner. This provides the relevant data table and data information collected based on the Hive MetaStore up to the last collection point.

3. Statistical Graph

3.1 Statistics view

View metrics such as the Hive database, total number of tables, total storage volume, and daily increase and daily comparison information for related metrics.

3.2 See the last access time of data tables to view the distribution, which can be used as a reference for hot and cold data distribution.

#### Note:

Interval	Description
Within 3 months	Most recent access time less than 3 months
3 months to 1 year	Most recent access time greater than or equal to 3 months and less than or equal to 1 year
1 year to 5 years	Most recent access time greater than 1 year and less than or equal to 5 years
Over 5 years	Most recent access time greater than 5 years
Others	Collection not enabled at the COS/CHDFS side Most recent access time for stored data not collected

The trend view provides the historical growth trend of the number of databases, tables, and table storage size. It includes a dimension for tables with small files, allowing observation of the distribution and growth of small file tables.
 You can view the database to which a table belongs, storage size, number of files, proportion of small files, and the table's partition details. The cloud file storage size and the proportion of small files provide an intuitive view of Hive small file issues.

6. Click Operation to **view partitions**, where you can see information such as partition name, partition size, and total number of files. Partition size and total number of files help identify partition skew and file quantity details.

#### **Risk note**

The analysis data for data table analysis will be collected daily starting at 14:00 UTC+8.

1. It mainly collects HMS database and table, partition metadata, and NameNode directory information. This operation will slightly increase the request load for HMS and NameNode. If the request volume continuously exceeds the load capacity, you can ticket feedback to disable this feature.

2. HMS and NameNode data collection only involves metadata and does not affect specific business data.

# YARN Job Query

Last updated : 2023-12-27 14:43:45

### Overview

You can view submitted jobs, memory and vCore consumption, and other information at the user granularity, and quickly check the queue, status, duration, and other metrics of YARN jobs in the EMR console. Comparison of historical tasks by job, job insight, task execution information, and other features are also available.

### Directions

1. Log in to the EMR console and click the ID/Name of the target cluster in the cluster list.

2. On the cluster details page, click **Cluster services** and select **Operation** > **Job query** in the top-right corner of the YARN component to view job statistics and resource consumption trends, query job and task information, get insights into application execution results, and compare application monitoring data.

i. Submitted jobs, view and distribution of memory and vCores consumed at the user granularity, and recent trends of relevant metrics.

rvice Status	Job Query	Role Ma	anagement Configuratio	n Management R	esource Schedulir	ıg		
ob List								
Today	Yesterday La	st 7 days	Last 30 days 2022-11-14	00:00:00 ~ 2022-11-14	17:44:01	Application ID or I	name (fuzzy search sup	p Q
ApplicationId		User <b>T</b>	ApplicationName		Queue <b>T</b>	ApplicationType <b>T</b>	ElapsedTi \$	Task info
		hadoop	1.11		root.default	SPARK	4min10s	Applicat
		hadoop			root.default	MAPREDUCE	34s270ms	Details N
	100	hadoop			root.default	MAPREDUCE	4min51s	Details N
		hadoop			root.default	MAPREDUCE	4min51s	Details

ii. Job filtering by user, application name, queue name, job type, duration, or relevant throughput metrics.

Job list									
Today Yesterday	Last 7 days	Last 30 days	2023-07-12 00:00:00	~ 2023-07-12 20:27:15			Application ID or name (fuzzy search	n supp 🛛 🔍	¢
ApplicationId	InsightRe	esult i 🄻	User <b>T</b>	ApplicationName	Queue <b>T</b>	ApplicationType <b>T</b>	ElapsedTime <b>\$</b>	Operation	
							No data yet		
4									
Total items: 0							10 🔻 / page 🛛 🕅 🔺 1	/ 1 page	•

iii. Resource consumption statistics by user, queue, or other metrics, helping check costs (also supported by API).

tatistics list								
Today Yester	lay Last 7 days	Last 30 days	2023-07-12 00:00:00 ~ 2023-07	-12 20:27:15 💼 Select	Ψ			
Queue <b>T</b>	User <b>T</b>	Job t	ype <b>T</b> vCor	es ‡ Memory	\$	Data read from HDFS ‡	Data written to HDFS 💲	Submitted jobs
				No data yet				
Total items: 0							10 🔻 / page 🔣 📢	1 / 1 page

#### Caution

To use the new features of task information, application insight, and application comparison for applications of Spark type, you first need to check whether the Spark History version is supported by running curl

"http://localhost:10000/api/v1/applications" | json\_pp . If the returned result is in abnormal

JSON format, the Spark History version is not supported. In this case, you can submit a ticket to apply for enabling such features.

In job query, ResourceManager data is collected once every 30s. The collection operation has negligible impact on cluster running.

3. In the job list, click **More** > **Application insight** to view detailed application insight items and related insight rules, results, and suggestions.

1	Application insight			
		_		
	Insight item	Level T	Rule	Result/Recommendation
	Map Data Skew	Severe	The data processed by the task is more than twice the	$\odot$
	Map Memory waste	Medium	The idle Memory resources account for more than 20	$\odot$
	Map Slow Task	Medium	The task processing time is more than twice the avera	$\odot$
	MapperTaskGC	Medium	The GC time of executor exceeds 20%	$\odot$
	Reduce Data Skew	Severe	The data processed by the task is more than twice the	$\odot$
	Reduce Memory waste	Medium	The idle Memory resources account for more than 20	$\odot$
	Reduce Slow Task	Medium	The task processing time is more than twice the avera	$\odot$

#### Caution

To ensure stable cluster running, no collection will be performed when any of the following rules is met:

1.1 Apps for which the runtime is less than 10 minutes.

1.2 Apps for which the number of subtasks during collection is greater than 30,000.

1.3 Apps for which the collection delay time is longer than 24 hours.

You can submit a ticket to modify relevant parameters of insight collection policy.

#### **Risk description**

YARN's application insight feature collects and analyzes the application data of Spark History, Job History, and Timeline Server. If you find the number of requests of this service constantly exceeds the load threshold, you can submit a ticket to disable the feature.

4. In the job list, click **More** > **Application comparison** to compare the business metrics between the current application and another application of the same type.





#### Caution

Only applications of MAPREDUCE, Spark, or Tez type in the final status of SUCCEEDED can be compared. Applications of the same type with the same name are filtered on the page by default. Only applications of the same type can be selected and filtered for comparison, and real-time data can be filtered and queried on the backend. 5. In the job list, click **More** > **Task information** to view the job task list, hosts comparison, and task execution logs.

Task list							
Compare hosts						Enter a tas	sk ID Q
id	State <b>T</b>	Туре 🔻	Host	StartTime	FinishTime	E	Operation
ta:	SUCCEEDED	MAP		2022-11-14 17:38:44	2022-11-14 17:39	<b>):01</b> 1	stderr stdout syslog
task.	SUCCEEDED	REDUCE		2022-11-14 17:39:04	2022-11-14 17:39	):11 7	stderr stdout syslog
Total items: 2					10 ▼ / page 🛛	•	1 / 1 page 🕨

The feature support is as follows:

Job Type	Task Information	Hosts Comparison	Task Log
MAPREDUCE	Supported	Supported	Supported
Spark	Supported	Unsupported	Supported
Tez	Supported	Supported	Unsupported
Others	Unsupported	Unsupported	Unsupported

# HDFS File Storage Analysis

Last updated : 2023-12-27 14:44:48

### Overview

This feature allows you to perform an overall analysis of HDFS storage, including the total number of files, total storage usage, file distribution, and recent trends in file storage as of the day before the current day (T-1). It also provides top directory lists for large and small files.

You can view the daily changes and recent trends in the total number of files and total storage usage based on HDFS storage within a cluster.

The file count and storage usage pie charts can help you understand the proportion of small files and the storage used by them.

The feature also allows you to query and download directory information for the top 1,000 large/small files as of the last collection time.

### Directions

1. Log in to the EMR console and click the ID/Name of the corresponding cluster in the cluster list.

2. On the cluster details page, click **Cluster Service** and choose **Operation** > **File Storage Analysis** in the top-right corner of the HDFS component block to view information about relevant files stored on **HDFS** and their directories as of the last collection time.

3. Obtain statistical views.

3.1 You can view the daily increase and day-over-day change in the total number of files stored on HDFS and total storage usage.

3.2 Pie charts are provided to show you the numbers of empty files (size = 0 MB), small files (size  $\leq$  2 MB), other files (2 MB < size < 128 MB), and large files (size  $\geq$  128 MB), and the storage used by each of the four categories of files. 4. You can view recent trends in the number and storage usage of different categories of files.

5. You can query and download information about the top 1,000 small/large files as of the day before the current day

(T-1), including the file name, file path, file size, user group, owner, and last access time.

#### **Risk description**

The data required for file storage analysis is collected at 14:00 (Beijing time) every day.

1. The file storage analysis involves collecting and analyzing backup FsImage files, which can increase the memory usage of the local machine (a maximum increase of 4 GB). If the proportion of machines in the cluster that are using memory consistently remains at a high level, you can submit a ticket to disable the feature.

2.. In HA clusters, the analysis feature is executed on the Standby Master node, while in non-HA clusters, the feature is executed on the Master node.

# Impala Query Management

Last updated : 2023-12-27 14:45:04

### Feature overview

You can view the percentile distribution of various metrics for Impala queries in multiple dimensions. The Impala list allows you to quickly view detailed metrics such as query statements, query status, users, databases, scanned rows, peak memory usage, total read/sent bytes, and HDFS scanned rows.

### Directions

1. Log in to the EMR console and click the **ID/Name** of the target cluster in the cluster list to enter the cluster details page.

2. On the cluster details page, click **Cluster Service** and select **Operation** > **Query Management** in the top-right corner of the Impala component to query jobs.

Take the duration as an example. In the currently selected time range, if the TP90 percentile duration is 6.86k (ms), 90% of queries last for within 6.86s.



3. Impala query list information is provided. Certain column header fields can be filtered and sorted, and compound filtering by multiple dimensions is supported.

Query list								
Today	Yesterday	Last 7 days	Last 30 days	2022-11-09 00:00:00 ~ 2022-11-09 20	:29:24 💼	Execution statement of	or query ID	Q ¢
statement			Coordinator <b>T</b>	MaxNodePeakMem 🗘 🍸	queryType	PerNodePea 🕈 🍸	ScanHDF 🕈 🍸	Operation
select * from	n hivetest.local_N	ative_TEXT		-	QUERY		-	Details
select count	:(*),name from hiv	etest.local			QUERY			Details
select count	:(*),name from hiv	etest.local			QUERY			Details
select a.nam	ne,b.age from hiv	etest.local			QUERY			Details
select * from	hivetest.hbase_	test			QUERY			Details
SELECT get	_json_object(imp	ala_hdfs_js		199.84 KB	QUERY	352.17	2	Details
select * from	n hivetest.cos_de	fault_defaul			QUERY			Details
insert into hi	ivetest.cos_defau	lt_default			QUERY			Details
SELECT * fro	om impalatest.im	pala_hdfs_j		199.84 KB	QUERY	335.88	2	Overview Details
CREATE ext	ernal table if not	exists impa	2		DDL			Details
Total items:	41					<b>10 ▼</b> / page ⊯	◀ 1 /5 p	ages 🕨 🕨

4. Click **Overview** in the **Operation** column to view the time distribution information throughout the entire lifecycle of an Impala query, key metric information, and certain node information of the runtime.





Click **Details** in the **Operation** column to view query statements, query plans, execution overview, profiles, and memory information.



Overview Products ▼ Elas	tic MapReduce Cloud Virtual Machine	Details					
emr-51kb33a9 EMR-o9hz7opx	Cluster Service / IMPALA -	Query Statement Query Plan Execution Overview Profile Memory					
Cluster Overview							
Instance info	Today Yesterday Las	Query (ic Summary: Session ID:					
Cluster Service	statement	Session Type: BEESWAX					
Cluster Resource	insert into stu values("gary",23),("cas	Start Time: 2022-08-01 19:58:55.930712000 End Time: 2022-08-01 19:58:55.127934000 Query Type: QUERY Query State: FINISHED Impala Query State: FINISHED Query Status: OK Impala Version: impalad version 3.4.0-RELEASE RELEASE (build 66699337e89fb8f4b2d758173eab8225f User: hadoop Connected User: hadoop Delegated User: Network Address: IV Default Db: default Sql Statement: select *from stu2 Coordinator: 10.0.0.132:27002					
Cluster monitor	insert into stu values("gary",23),("cas						
* DashBoard BETA	insert into stu values("gary",23),("cas						
<ul> <li>Cluster Events</li> <li>Log Search</li> </ul>	insert into stu values("gary",23),("cas						
Cluster Inspection	insert into stu values("gary",23),("cas						
<ul> <li>Alarm Records</li> </ul>	insert into stu values("gary",23),("cas	Query Options (set by configuration): TIMEZONE Query Options (set by configuration and planner): MT_DOP=0,TIMEZONE					
Auto Scaling	// //						

Note:

For Impala queries that take more than 3s to execute, you can view profiles in **Overview** and **Details**.

# HBase Table Analysis

Last updated : 2023-12-27 14:45:41

### Overview

Table analysis provides information such as the numbers of read/write requests and storage information for an entire HBase table, each individual region within the table, and each region server. The feature also supports analyzing the read/write queries per second (QPS) and their trends for the corresponding table or region server in practical scenarios.

#### Directions

1. Log in to the EMR console and click the ID/Name of the corresponding cluster in the cluster list.

2. On the cluster details page, click **Cluster Service** and choose **Operation** > **Table analysis** in the top-right corner of the HBase component block to query HBase table loads.

## Supported Operations in the Table List

In the HBase table list, you can view information about the read QPS, write QPS, MetaStore storage, and StoreFile size for each table. You can also identify the top tables in the cluster by using the sorting button in each column.
Clu	uster Service /	HBASE 🔻							More 🔻 📔 Help Documen
S	ervice Status	Table analysis	Role management	Client management	Configuration Ma	anagement RIT fix	ing		
	Tables								
	Table		Read Requests/s 🗘	Write Requ	ests/s ‡	Memstore Size 🕈	Storefile	Size ‡	Operation
	hbase_meta		0.02	0		768	10737		Regions RegionServers
	hbase_namespace		0	0		256	4926		Regions RegionServers
	Total 2 items							Lines per page 10 💌	I I /1 page ►
	Statistics by Reg 1 hour 6 Region	ion hours 12 hours	2023-05-23 19:02:37 ·	- 2023-05-23 20:02:37 📋	negionServer <b>T</b>	Read	RequestRate ⊬ ↓	WriteRequ	Search by table nam Q
	hbase:meta,,1		hbase_meta		10.0.0.24:6003	0.01		0	
	59ecc1cbc0d857b	0397c75744c8cb66c	hbase_names	pace	10.0.0.95:6003	0		0	
	Total items: 2							<b>10 🔻</b> / page	I < 1 / 1 page ►

### Viewing table details

Click a table name to obtain a detailed view of the table. The details page displays the numbers of read/write requests, MemStore storage, and storeFile size for the entire table or each individual node. You can switch between nodes by using the node filter in the top-right corner.

Real time	Last 24 hours	Last 7 days	2023-05-23 19	:03:26 to 2023-0	5-23 20:03:26	İ	Node Filte	er V				
Number	Of Requests 🛈			23	Store S	ize (E	Byte) 🛈					
0.040					25000							
0.030					20000							
0.020					15000							
01020					10000							
0.010					5000							
0.000	9:10 19:20 19:3	0 19:40 19:5	50 20:00		0	19:1	0 19:20	19:30	19:40	19:50	20:00	
- Write	- Rea	d			— storeF	ileSiz	e -	- memsto	eSize			

#### Viewing the overall region information

Click **Regions** to view the overall region information, including the number of read/write requests for each region in the table, which helps you identify hot spot regions.

Clu	uster Service / HBASE / hbase_	meta / Region List				🖽 Help Docu
	Refresh					
	Region Name \$	RegionServer \$	Start Key 🗘	End Key ‡	Read Request Count/s \$	Write Request Count/s 🗘
	hbase:meta,,1	10.0.0.24;6003	-	-	0.007	0
	Total 1 item				Lines per page 10 💌	14 4 1 /1 page
L						

#### Viewing region details

Click a region name to obtain the detailed view of the region. The details page displays the numbers of read/write requests for the table at different time granularities. Select an option from **Time granularity** in the top-right corner to switch between different time granularities.

Keal time	Last 24 hours	Last 7 days	2023-05-2	19:04:38 to 2023-05-23 20:04:38	Time granularity: 30 seconds 🔻
Number	Of Requests 访		53		
0.015					
0.010					
0.005					
0.000	10 19:20 19:30 19:	40 19:50 20:00			
- WriteRe	quest – ReadRe	quest			

### Viewing the overall region server information

Click **RegionServers** to view the request latencies of each region server that hosts one or more regions in the table.

luster Service / HBASE	/ hbase_namespace / Reg	gionserver list				🖽 Help Docume
Refresh						
RegionServer \$	GetTimeTp99 ‡	ScanTimeTp99 \$	PutTimeTp99 \$	IncrementTimeTp99 <sup>‡</sup>	AppendTimeTp99 \$	DeleteTimeTp99 \$
10.0.0.95	0	0	0	0	0	0
Total 1 item					Lines per page 10 🔻	H

### **Region Analysis**

×

### Tencent Cloud

With region analysis, you can search for the table that a specific region belongs to or filter results by the region server that hosts the region. By examining the average request QPS and average read/write QPS, you can identify hot spots in the cluster where a large number of requests are being processed.

Real time	Last 24 hours	Last 7 days	2023-05-23	19:05:47 to 2023-05-23 20:05:47	İİ	Time granularity: 30 seconds 🔻
REQUEST	T LATENCY 🛈		53			
1.0						
0.5						
0.0						
-0.5						
-1.0 19:10	19:20 19:30 19:40	0 19:50 20:00				
- append	Time — deleteTim	ie_9 — incremer	ntTim			
- but time	99t scan time	99 dettime	991			

By clicking the view button in the "Average read QPS" or "Average write QPS" column header, you can view the trend of read/write QPS for the current region and observe sudden changes in request traffic. You can specify the time range for the information displayed.

## Kudu Table Analysis

Last updated : 2023-12-27 14:45:59

### Overview

Kudu table monitoring and Tablet analysis can help you identify data hot spots and skew in tables and Tablet deployment.

1. Kudu table analysis provides load information, including table-level, Tablet, and TabletServer read/write QPS and storage.

2. Through Tablet analysis, you can analyze the historical trends of read/write QPS information for tables or TabletServers, based on real-world scenarios.

### Directions

1. Log in to the EMR console and click the ID/Name of the corresponding cluster in the cluster list.

2. On the cluster details page, click **Cluster Service** and choose **Operation** > **Table Analysis** in the top-right corner of the Kudu component block to view the analysis of Kudu table loads.

### Supported Operations in the Table List

In the Kudu table list, you can view information about table-level request QPS, write QPS, and OnDiskDataSize storage. You can also identify the top tables in the cluster by using the sorting button in each column header.

ervice Status	Table a	analysis	Role management	Client management	Configuration Manage	ement				
lables										
									Search by table nam	
TableName		TableID		TabletCount \$	OnDiskDataSize 💲	InsertedRate(count/s) \$	UpsertedRate(count/s	<b>\$</b> Upd	datedRate(count/s) Operation	1
								No data yet	t	
•										
▲ Total items: 0	0							10 🔻 / page	H + 1 /1 pag	ge
Total items:	0							<b>10 v</b> / page	н н 1 /1 раз	ge
∢ Total items: ∶	0							10 🔻 / page	H 🔫 1 / 1 pag	ge
Total items:	0 (Tablet							<b>10 ▼</b> / page	id d 1 /1 pag	ge
Total items:	0 / Tablet							10 🔻 / page	14 4 1 /1 pag	ge
Total items:  Statistics by 1 hour	0 / Tablet 6 hours	12 hours	2023-05-23 18:55:39	~ 2023-05-23 19:55:39	1			10 ▼ /page	Id     1     / 1 page	ge
Total items:   Statistics by  1 hour  TabletID	0 / Tablet 6 hours	12 hours	2023-05-23 18:55:39 TableName	~ 2023-05-23 19:55:39 📑 TabletServer	] OnDiskDataSize ⊬ ‡	InsertedRate(count/s)	⊻↓ UpsertedRate	10 ▼ / page (count/s) 上 \$	Search by table nam	ge ‡
Total items:      Statistics by     1 hour     TabletID	0 / Tablet 6 hours	12 hours	2023-05-23 18:55:39 TableName	~ 2023-05-23 19:55:39 📋 TabletServer	0nDiskDataSize 🗠 🕈	InsertedRate(count/s)	ビ↓ UpsertedRate No	10 ▼ / page (count/s) ビ ‡ data yet	K     1     / 1 page       Search by table nam     UpdatedRate(count/s)     K	
Total items:  Statistics by  1 hour TabletID	0 / Tablet 6 hours	12 hours	2023-05-23 18:55:39 TableName	~ 2023-05-23 19:55:39 💼 TabletServer	0nDiskDataSize 🗠 🕈	InsertedRate(count/s)	ビ↓ UpsertedRate No	10 ▼ / page (count/s) 🗠 ‡ data yet	K     1     / 1 page       Search by table nam     UpdatedRate(count/s)     2	ge ∠ ‡

#### Viewing table details

Click a table name to obtain a detailed view of the table. The details page displays the numbers of read/write requests and storage size (including OnDiskDataSize) for the entire table or each individual node. You can switch between nodes by using the node filter in the top-right corner.

#### Viewing the overall Tablet information

Click **Tablets Operation** to view the overall information about the read/write request volumes of each Tablet in the table, and identify hot spots within the Tablets.

#### **Viewing Tablet details**

Click the corresponding Tablet name to view the Tablet details and metric trends. The details page provides data on request and scan metrics for the selected Tablet at various time granularities. You can change the time granularity in the top-right corner.

#### Viewing the overall TabletServer information

Click **TabletServers Operation** to view the overall information about each TabletServer, such as the request latency and data storage.

### **Tablet Analysis**

With Tablet analysis, you can search for the table that a specific region belongs to or filter results by the TabletServer that hosts the region. By examining the average request QPS and average read/write QPS, you can identify hot spots in the cluster where a large number of requests are being processed.

By clicking the view button in the "Average read QPS" or "Average write QPS" column header, you can view the trend of read/write QPS for the current Tablet and observe sudden changes in request traffic. You can specify the time range for the information displayed.

## **Cluster Inspection**

Last updated : 2023-12-27 14:46:15

### Overview

You can perform health check on the nodes and services in each cluster instantly or periodically (daily or weekly) with the selected inspection items. Only one periodic inspection task can be configured for each cluster to help you know the cluster health conditions periodically and fix the exceptions and risks in time.

When you set a one-time or periodic inspection task, general inspection items will be selected by default. If you want to check certain service features in special cases, you can select the corresponding inspection items as needed. However, service feature inspection will compromise cluster performance, which is thus not recommended during business peak hours.

After each inspection task is completed, a PDF report will be generated, which can be downloaded and deleted. Up to 50 inspection reports can be retained under each root account. If an excessive report is generated, the oldest one will be deleted.

#### Note

The inspection system supports service inspection items. Currently, only HDFS, YARN, HBase, Hive, and ZooKeeper are supported.

The configuration of a periodic inspection task cannot be changed when it is being executed.

### Directions

1. Log in to the EMR console and click the **ID/Name** of the target cluster in the cluster list to enter the cluster details page.

 On the cluster details page, select Cluster Monitoring > Cluster Inspection and you can perform health check on the nodes and services in the current cluster. Only one periodic inspection task can be configured for each cluster. You can also click One-Time Inspection to perform an inspection. To configure a periodic inspection task, click Periodic Inspection Settings.



emr-c6rw6dcf ceshi123	Cluster Inspection								
Basic Info v		One-Time Inspection	Today Ye	sterday Last 7 days	Last 30 days	2020-11-03 to 2020-11-10 📋		Periodic Inspection Settings	¢
Cluster Services		Report Time		Report name		Inspection Type	Report Status	Operation	
Cluster Resources									
Node Status						No data yet			
Resource Management									
Cluster Monitoring									
Cluster Overview									
Cluster Events									
Log Search									
Service Alarms									
Cluster Inspection									
Operation Logs									

One-Time Inspection: The health status of nodes and services in the cluster from a specified time point to the current time point will be checked, and an inspection report will be generated.

<ul> <li>(i) 1. One-time and genera</li> <li>2. Periodic services in for each clu</li> <li>3. Inspection</li> <li>Policy ☑</li> </ul>	e inspection: check the health status of a cluster's nodes and service from a specified moment to the current time ate an inspection report. inspection: when periodic inspection is enabled, the system automatically checks the health status of nodes and each preset inspection cycle and generate an inspection report. Only one periodic inspection can be configured uster. on items support all enabled event monitoring policies by default. You can adjust the inspection items in <u>Event</u>					
Inspection Interval	2020-11-10 OO:00 O to date					
Inspection Item	Select all Invert All					
	- NODE					
	CpuUsageMoreThanThreshold					
	CpulOWaitAvgMoreThanThreshold					
	✓ MemoryFreeException					
	✓ ProcessCntMoreThanThreshold					
	✓ ProcessOutOfMemory					
	✓ DiskIOError					
	✓ DiskSpaceMoreThanThreshold					
	✓ DiskUtilMoreThanThreshold					
	ProcessLost					
	✓ HeartBeatLost					
	✓ HostnameChanged					
	HDFS					
	VARN					

Periodic Inspection: After a periodic inspection policy is enabled, the system will automatically check the health status of the nodes and services in the cluster in every inspection cycle and generate an inspection report. One periodic inspection policy can be configured for each cluster.

		~
()	<ol> <li>One-time inspection: check the health status of a cluster's nodes and service from a specified moment to the current time and generate an inspection report.</li> <li>Periodic inspection: when periodic inspection is enabled, the system automatically checks the health status of nodes and services in each preset inspection cycle and generate an inspection report. Only one periodic inspection can be configured for each cluster.</li> <li>Inspection items support all enabled event monitoring policies by default. You can adjust the inspection items in Event Policy. ☑.</li> </ol>	
eriodic	Inspection	

Inspection Item: All enabled event monitoring policies are supported by default. You can adjust the inspection items as instructed in Cluster Event. All events with monitoring enabled are selected by default as the initial inspection items. When you customize the inspection items, those selected last time will be selected by default.

Periodic Inspectio	on Settings	
<ul> <li>(i) 1. One-time and generative</li> <li>2. Periodic in services in e for each cluss</li> <li>3. Inspection</li> <li>Policy. ☑ .</li> </ul>	inspection: check the e an inspection report nspection: when perio each preset inspection ster. n items support all ena	health status of a cluster's nodes and service from a specified moment to the current time dic inspection is enabled, the system automatically checks the health status of nodes and cycle and generate an inspection report. Only one periodic inspection can be configured abled event monitoring policies by default. You can adjust the inspection items in <u>Event</u>
Periodic Inspection		
Inspection Cycle	Weekly 💌	Monday - 00:00 🕓
Inspection Item	Select all  Select all  HDFS	Monday         Tuesday         Wednesday         Invert All         Thursday         Friday         Saturday         Sunday
	YARN	
		Save configuration Cancel

## Java Analysis

Last updated : 2024-10-30 10:44:06

### Feature Introduction

Supports online GC analysis for Java processes by collecting, recording, and analyzing GC logs in real time, helping users identify whether process exceptions are caused by GC.

GC view: Filter by service, role, node, and time to view the related GC distribution and point information.

GC list: Filter GC log data as needed to view multidimensional GC information.

### Directions

1. Log in to the EMR console and click the corresponding **Cluster ID/Name** in the cluster list to enter the cluster details page.

2. On the cluster details page, click **Cluster Monitoring**, and then select JAVA Analysis > GC Online Analysis to view the GC view and point information.

3. Additionally, the GC query list information is provided, with some column headers supporting filtering or sorting operations.

The services and roles that support GC log collection are as follows:

Role
Zookeeper
zkfc
NameNode
SecondaryNameNode
JournalNode
DataNode
ResourceManager
NodeManager
JobHistoryServer
TimeLineServer

Hbase	Hmaster			
	RegionServer			
	HbaseThrift			
	HiveMetaStore			
Hive	HiveServer2			
	HiveWebHcat			
Knox	Idap			
	gateway			
Oozie	Oozie			
	Nimbus			
Storm	Supervisor			
	worker			
Spark	SparkJobHistoryServer			
	AlluxioMaster			
Allunia	AlluxioJobMaster			
Alluxio	AlluxioWorker			
	Alluxio-JobWorker			
Dongor	EmbeddedServer			
nangei	EnableUnixAuth			
Livy	LivyServer			
Kylin	Kylin			
Dreate	Presto-Coordinator			
Flesto	Presto-Worker			
Prostocal	PrestoSQL-Coordinator			
Frestosyl	PrestoSQL-Worker			

Cosranger	CosRangerServer		
Starrocks	StarRocksFe		
	StarRocksBroker		
	DorisFe		
Dons	DorisBroker		
	overlord		
	coordinator		
Druid	router		
	broker		
	historical		
	middleManager		
Zeppelin	Zeppelin		
Kyuubi	KyuubiServer		
Kafka	Kafka		

# Monitoring Metrics Node Monitoring Metrics

Last updated : 2023-12-27 14:46:33

### Node - CPU

Title	Metric	Unit	Description		
	idle	%	Percentage of CPU idle time		
	irq	%	Percentage of interrupts		
	nice	%	Percentage of CPU utilization under nice priority		
	steal	%	Percentage of wait time by virtual CPUs for physical CPUs		
CPU utilization	softirq	%	Percentage of CPU soft interrupts		
	guest	%	Percentage of time spent running virtual processors		
	system	%	CPU utilization in kernel mode		
	user	%	CPU utilization in user mode		
	iowait	%	Percentage of CPU idleness due to process I/O waits		
	1m	%	1-minute load		
Load	5m	%	5-minute load		
	15m	%	15-minute load		
Cores	cpu_count	-	Number of CPU cores		

### Node - memory

Title	Metric	Unit	Description
Memory utilization MemTotal		GB	Total memory size
	MemFree	GB	Total free memory size
	MemAvailable	GB	Total available memory size
	Buffers	GB	Total memory size used by buffers



	Cached	GB	Total memory size used by file cache
	SwapCached	GB	Total swap memory size by anonymous page writes
	SwapFree	GB	Total available swap size
	AnonPages	GB	Total unmapped memory size
	SwapTotal	GB	Total swap size
	Dirty	GB	Total memory size to write to disk
	Writeback	GB	Total memory size being written back to disk
	HardwareCorrupted	GB	Total unavailable memory size due to memory hardware failure
	Shmem	GB	Total shared memory size
	MemUsed	GB	Total used memory size
Percentage of used memory	available_percent	%	Percentage of available memory size out of total memory
	used_percent	%	Percentage of used memory size out of total memory

### Node - disk

Title	Metric	Unit	Description
	Read	MB/s	Data read per second
Device read/ write rate	Write	MB/s	Data written per second
Device IOPS	all	count/s	Number of I/O operations in progress on current device
	Read	ms	Average wait time per device I/O read operation
I/O operation time	Write	ms	Average wait time per device I/O write operation
	IO	ms	Average processing time per I/O request
Device read/write QPS	Read	count/s	Read QPS
	Write	count/s	Write QPS
	Merge-Read	count/s	Merged read QPS



	Merge-Write	count/s	Merged write QPS	
I/O device utilization	all	%	Disk busyness	
	Free	GB	Free disk storage space	
Disk space	Available	GB	Available disk storage space (for unprivileged users)	
	Total	GB	Total disk storage space	
Disk space utilization	Used	%	Disk space utilization	
INODES	Free	-	Number of remaining disk inodes	
	Total	-	Total number of disk inodes	
Inode utilization	Used	%	Disk inode utilization	

### Node - file handle

Title	Metric	Unit	Description
Eile beselle	allocated	-	Number of allocated file handles
File Halidle	maximum	-	Maximum number of file handles
System interrupt	intr_total	count/s	Number of system interrupts
System context switch	context_switches_total	count/s	Number of system context switches
	forks_total	-	Number of new system processes
	procs_running	-	Number of running system processes
System process	procs_blocked	-	Number of blocked system processes
	procs_total	-	Total number of system processes
	thrds_total	-	Total number of system threads
Agent version	AgentVersionI	version	Agent version

#### Node - network

Title	Metric	Unit	Description
TCP LISTEN	ListenDrops	count/s	Number of incoming connections (SYN packets)



exception			dropped for any reason
	ListenOverflows	count/s	Number of occurrences where the upper limit of the Accept queue is exceeded after the last step of three-way handshake is completed
	SyncookiesFailed	count/s	Number of packets received with invalid SYN Cookie information
TCPSyncookies	SyncookiesRecv	count/s	Number of packets received with valid SYN Cookie information
	SyncookiesSent	count/s	Number of SYN/ACK packets sent through SYN Cookie
	TCPAbortOnTimeout	count/s	Number of connections closed because the attempts of retransmissions of various timers (RTO/PTO/keepalive) exceed the upper limit
	TCPAbortOnData	count/s	Number of sockets closed due to unknown data received
TCP connection abort exception	TCPAbortOnClose	count/s	Number of sockets closed when the user-mode program has data in the buffer
	TCPAbortOnMemory	count/s	Number of connections closed due to memory issues
	TCPAbortOnLinger	count/s	Number of connections suspended in lingering status after being closed
	TCPAbortFailed	count/s	Number of failed attempts to close connection
	ActiveOpens	count/s	Number of actively established TCP connections
TCP connection establishment	CurrEstab	count/s	Number of TCP connections currently established
	PassiveOpens	count/s	Number of passively established TCP connections
	AttemptFails	count/s	Number of connection establishment failures
	EstabResets	count/s	Number of reset connections
TCP packet	InSegs	count/s	Number of received packets, including erroneous



			ones
	OutSegs	count/s	Number of sent packets
	RetransSegs	count/s	Number of received TCP packets
	InErrs	count/s	Number of retransmitted packets
	OutRsts	count/s	Number of sent RST packets
TCP	RetransSegsRate	%	Retransmission rate at TCP layer
retransmission	ResetRate	%	RESET sending frequency
rale	InErrRate	%	Percentage of erroneous packets
	TW	count/s	Number of sockets ending TIME_WAIT status after normal timeout
	TWKilled	count/s	Number of sockets ending TIME_WAIT status through tcp_tw_recycle mechanism
TOP TIME-WAT	TCPTimeWaitOverflow	count/s	Number of TIME_WAIT sockets unable to be allocated due to limit exceeding
	TWRecycled	count/s	Number of sockets ending TIME_WAIT status through tcp_tw_reuse mechanism
TCP RTO	TCPTimeouts	count/s	Number of first RTO timer timeouts
	TCPSpuriousRTOs	count/s	Number of spurious timeouts detected through F- RTO mechanism
	TCPLossProbes	count/s	Number of Tail Loss Probe (TLP) packets sent due to Probe Timeout (PTO)
	TCPLossProbeRecovery	count/s	Number of lost packets just repaired by TLP probes
	TCPRenoRecoveryFail	count/s	Number of connections that enter the Recovery phase and then undergo RTO (SACK option not supported by the opposite)
	TCPSackRecoveryFail	count/s	Number of connections that enter the Recovery phase and then undergo RTO (SACK option supported by the opposite)



	TCPRenoFailures	count/s	Number of failures that enter the TCP_CA_Disorder phase and then undergo RTO (SACK option not supported by the opposite)
	TCPSackFailures	count/s	Number of failures that enter the TCP_CA_Disorder phase and then undergo RTO (SACK option supported by the opposite)
	TCPLossFailures	count/s	Number of connections that enter the TCP_CA_Loss phase and then undergo RTO timeout
	RtoAlgorithm	1/s	Number of delayed algorithms for forwarding unanswered objects
TCP RTO constant	RtoMax	1	Maximum number of retransmissions due to TCP latency
	RtoMin	1	Minimum number of retransmissions due to TCP latency
	TCPLostRetransmit	count/s	Number of SKB retransmissions due to loss
	TCPFastRetrans	count/s	Number of fast SKB retransmissions
TCP	TCPForwardRetrans	count/s	Number of regular SKB retransmissions
retransmission	TCPSlowStartRetrans	count/s	Number of SKB retransmissions with successful slow start
	TCPRetransFail	count/s	Number of failed retransmission attempts
	OutDatagrams	count/s	Number of sent UDP datagrams
ODP datagram	InDatagrams	count/s	Number of received UDP datagrams
ENI data	eth0-receive_bytes	MB/s	Volume of data received by ENI
rate	eth0-transmit_bytes	MB/s	Volume of data sent by ENI
ENI packet receiving/sending	eth0-receive_drop	count/s	Volume of data received and then dropped by ENI
rale	eth0-receive_errs	count/s	Volume of data failed to be received by ENI
	eth0-transmit_drop	count/s	Volume of data sent and then dropped by ENI
	eth0-transmit_errs	count/s	Volume of data failed to be sent by ENI



	eth0-transmit_packetsl	count/s	Number of packets sent by ENI
	TCP_inuse	-	Number of TCP sockets in use (listening)
	TCP_orphan	-	Number of TCP connections waiting to be closed
TCP socket	TCP_tw	-	Number of TCP sockets to be terminated
	TCP_alloc	-	Number of TCP sockets allocated (established, sk_buff obtained)
	sockets_used	-	Total number of used sockets
	ESTABLISHED	-	Number of TCP connections in Established status
	SYN-SENT	-	Number of TCP connections in SYN-SENT status
	SYN-RECV	-	Number of TCP connections in SYN-RECV status
	FIN-WAIT1	-	Number of TCP connections in FIN-WAIT1 status
TCP connection	FIN-WAIT2	-	Number of TCP connections in FIN-WAIT2 status
status	TIME-WAIT	-	Number of TCP connections in TIME-WAIT status
	CLOSE	-	Number of TCP connections in CLOSE status
	CLOSE-WAIT	-	Number of TCP connections in CLOSE-WAIT status
	LAST-ACK	-	Number of TCP connections in LAST-ACK status
	LISTEN	-	Number of TCP connections in LISTEN status
	CLOSEING	-	Number of TCP connections in CLOSEING status

### Node - event

Title	Metric	Unit	Description
CPU utilization	used	%	1 - (percentage of CPU idle time)

🔗 Tencent Cloud

15-minute CPU load	15m	-	15-minute load
1-minute CPU load	1m	-	1-minute load
5-minute CPU load	5m	-	5-minute load
Disk IOPS	all	-	Number of I/O operations in progress on current device
Disk I/O operation time	IO	-	Average processing time per I/O request
Disk space utilization	Used	-	Disk space utilization
Disk I/O device utilization	all	-	Disk busyness
Memory utilization	used_percent	-	Percentage of used memory size out of total memory
Outbound network traffic rate	*-transmit_bytes	-	Volume of data sent by ENI
Inbound network traffic rate	*-receive_bytes	-	Volume of data received by ENI
TCP connections	CurrEstab	-	Number of TCP connections currently established

# **HDFS Monitoring Metrics**

Last updated : 2023-12-27 14:48:18

### HDFS - Overview

Title	Metric	Unit	Description
	CapacityTotal	GB	Total cluster storage capacity
Chuster storess conseit.	CapacityUsed	GB	Used cluster storage capacity
Cluster storage capacity	CapacityRemaining	GB	Remaining cluster storage capacity
	CapacityUsedNonDFS	GB	Non-HDFS used cluster capacity
Cluster load	TotalLoad	1	Current connections
Total files in cluster	FilesTotal	-	Total number of files
	BlocksTotal	-	Total number of blocks
	PendingReplicationBlocks	-	Number of blocks waiting to be backed up
	UnderReplicatedBlocks	-	Number of blocks with insufficient replicas
	CorruptBlocks	-	Number of corrupted blocks
Blocks	ScheduledReplicationBlocks	-	Number of blocks arranged for backup
	PendingDeletionBlocks	-	Number of blocks waiting to be deleted
	ExcessBlocks	-	Number of excess blocks
	PostponedMisreplicatedBlocks	-	Number of abnormal blocks postponed to be processed
Block capacity	BlockCapacity	-	Block capacity
Cluster data node	NumLiveDataNodes	-	Number of live data nodes
	NumDeadDataNodes	-	Number of data nodes marked as dead



	NumDecomLiveDataNodes	-	Number of decommissioned live nodes
	NumDecomDeadDataNodes	-	Number of decommissioned dead nodes
	NumDecommissioningDataNodes	-	Number of decommissioning nodes
	NumStaleDataNodes	-	Number of DataNodes marked as stale
HDFS storage space utilization	CapacityUsedRate	-	HDFS cluster storage space utilization
Snapshots	Snapshots	-	Number of snapshots
Disk failure	VolumeFailuresTotal	-	Total number of volume failures across all DataNodes

### HDFS - NameNode

Title	Metric	Unit	Description
Data traffic	ReceivedBytes	Bytes/s	Data receiving rate
Data trainc	MetricUnitReceivedBytesBytes/sSentBytesBytes/sRpcQueueTimeNumOps1/sRpcQueueTimeAvgTimemsRpcProcessingTimeAvgTimemsRpcProcessingTimeAvgTime1 per timeRpcAuthenticationFailures1 per timeRpcAuthenticationFailures3 per time	Data sending rate	
QPS	RpcQueueTimeNumOps	1/s	RPC call rate
Request	RpcQueueTimeAvgTime	ms	Average RPC latency
processing latency	RpcProcessingTimeAvgTime	ms	Average RPC request processing time
Authentication	RpcAuthenticationFailures	1 per time	Number of RPC authentica failures
	RpcAuthenticationSuccesses	2 per time	Number of RPC authentica successes
authorization	RpcAuthorizationFailures	3 per time	Number of RPC authorizati failures
	RpcAuthorizationSuccesses	4 per time	Number of RPC authorizati successes
Current connections	NumOpenConnections	-	Number of current connect



Length of RPC processing queue	CallQueueLength	-	Length of current RPC processing queue
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapCommitte configured by JVM
	MemHeapUsedM	MB	Size of HeapMemory curre used by JVM
J VIVI Memory	MemHeapCommittedM	MB	Committed size of JVM HeapMemory
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
	MemMaxM	MB	Maximum size of memory available to JVM runtime
Block reporting latency	BlockReportAvgTime	count/s	Average latency of process DataNode blocks per seco
	ThreadsNew	-	Number of threads in NEW status
	ThreadsRunnable	-	Number of threads in RUNNABLE status
	ThreadsBlocked	-	Number of threads in BLOCKED status
JVM Inreads	ThreadsWaiting	-	Number of threads in WAITING status
	ThreadsTimedWaiting	-	Number of threads in TIME WAITING status
	ThreadsTerminated	-	Number of threads in Terminated status
JVM logs	LogFatal	-	Number of FATAL-level loc
	LogError	-	Number of ERROR-level Ic
	LogWarn	-	Number of WARN-level loc



	LogInfo	-	Number of INFO-level logs
	YGC	-	Young GC count
GC count	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	SO	%	Percentage of used Survive memory
	S1	%	Percentage of used Survive memory
Manager	E	%	Percentage of used Eden memory
proportion	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
	CCS	%	Percentage of used compressed class space memory
Storages marked as content stale	NumStaleStorages	-	Number of DataNode stora marked as content stale
Pending block- related messages for later processing on the standby NameNode	PendingDataNodeMessageCount	count/s	Number of DataNode requir queued on the standby NameNode
Missing blocks	NumberOfMissingBlocks	-	Number of missing blocks
	NumberOfMissingBlocksWithReplicationFactorOne	-	Number of missing blocks 1)



	AllowSnapshotOps	count/s	Number of AllowSnapshot operations executed per second
	DisallowSnapshotOps	count/s	Number of DisallowSnapsh operations executed per second
	CreateSnapshotOps	count/s	Number of CreateSnapsho operations executed per second
Snapshot	DeleteSnapshotOps	count/s	Number of DeleteSnapsho operations executed per second
operation	ListSnapshottableDirOps	count/s	Number of ListSnapshottableDir operations executed per second
	SnapshotDiffReportOps	count/s	Number of SnapshotDiffReportOps operations executed per second
	RenameSnapshotOps	count/s	Number of RenameSnapshotOps operations executed per second
File operation	CreateFileOps	count/s	Number of CreateFile operations executed per second
	GetListingOps	count/s	Number of GetListing operations executed per second
	TotalFileOps	count/s	Number of TotalFileOps operations executed per second
	DeleteFileOps	count/s	Number of DeleteFile operations executed per second

🔗 Tencent Cloud

	FileInfoOps	count/s	Number of FileInfo operation executed per second
	GetAdditionalDatanodeOps	count/s	Number of GetAdditionalDatanode operations executed per second
	CreateSymlinkOps	count/s	Number of CreateSymlink operations executed per second
	GetLinkTargetOps	count/s	Number of GetLinkTarget operations executed per second
	FilesInGetListingOps	count/s	Number of FilesInGetListin operations executed per second
	FilesDeleted	count	Number of deleted or renar files and folders
File statistics	FilesCreated	count	Number of created files and folders
	FilesAppended	count	Number of appended
Transaction operation	TransactionsNumOps	count/s	Number of journal transact operations processed per second
	TransactionsBatchedInSync	count/s	Number of journal transact operations batch processe per second
Image operation	GetEditNumOps	count/s	Number of GetEditNumOp operations executed per second
	GetImageNumOps	count/s	Number of GetImageNum( operations executed per second
	PutImageNumOps	count/s	Number of PutImageNumC operations executed per second



Sync operation	SyncsNumOps	count/s	Number of journal sync operations processed per second
Block	BlockReceivedAndDeletedOps	count/s	Number of BlockReceivedAndDeleted operations executed per second
operation	BlockOpsQueued	count/s	Number of processed DataNode block reporting operations
Cache reporting	CacheReportNumOps	count/s	Number of CacheReport operations processed per second
Block reporting	BlockReportNumQps	count/s	Number of DataNode block reporting operations proces per second
Sync operation latency	SyncsAvgTime	ms	Average latency of process journal sync operations
Cache reporting latency	CacheReportAvgTime	ms	Average latency of cache reporting
	GetEditAvgTime	ms	Average latency of reading Edit files
Image operation latency	GetImageAvgTime	ms	Average latency of reading image files
	PutImageAvgTime	ms	Average latency of writing image files
Transaction operation latency	TransactionsAvgTime	ms	Average latency of process journal transaction operatic
Start time	StartTime	ms	Process start time
Active/Standby status	State	1	NameNode HA status
Active/Standby status	State	1: Active.	NameNode active/standby status

		0: Standby	
	PeakThreadCount	-	Peak number of threads
Threads	ThreadCount	-	Number of threads
	DaemonThreadCount	-	Number of backend thread
Transactions since the last checkpoint	SinceLastCheckpoint	count	Total number of transaction since the last checkpoint
Checkpoint time	LastCheckpoint	time	Time since the last checkp
Length of the queue waiting for file locks	LockQueueLength	count	LockQueueLength length of the queue waiting file locks
	CompleteAvgTime	ms	Average latency of Complete requests
	CreateAvgTime	ms	Average latency of Cre requests
	RenameAvgTime	ms	Average latency of Rename requests
Average RPC time (1)	AddBlockAvgTime	ms	Average latency of AddBlock requests
	GetListingAvgTime	ms	Average latency of GetListing requests
	GetFileInfoAvgTime	ms	Average latency of GetFileInfo requests
	SendHeartbeatAvgTime	ms	Average latency of SendHeartbeat requests
Average RPC time (2)	RegisterDatanodeAvgTime	ms	Average latency of RegisterDatanode request:
	BlockReportAvgTime	ms	Average latency of BlockReport requests
	DeleteAvgTime	ms	Average latency of De



			requests
	RenewLeaseAvgTime	ms	Average latency of RenewLease requests
	BlockReceivedAndDeletedAvgTime	ms	Average latency of BlockReceivedAndDeleted requests
	FsyncAvgTime	ms	Average latency of fsy requests
	VersionRequestAvgTime	ms	Average latency of VersionRequest requests
Average RPC time (3)	ListEncryptionZonesAvgTime	ms	Average latency of ListEncryptionZones reque
	SetPermissionAvgTime	ms	Average latency of SetPermission requests
	SetTimesAvgTime	ms	Average latency of SetTimes requests
	SetSafeModeAvgTime	ms	Average latency of SetSafeMode requests
	MkdirsAvgTime	ms	Average latency of Mk requests
	GetServerDefaultsAvgTime	ms	Average latency of GetServerDefaults request
	GetBlockLocationsAvgTime	ms	Average latency of GetBlockLocations reques
RPC statistics (1)	CompleteNumOps	count/s	Number of Complete c per second
	CreateNumOps	count/s	Number of Create call per second
	RenameNumOps	count/s	Number of Rename ca per second
	AddBlockNumOps	count/s	Number of AddBlock ( per second



	GetListingNumOps	count/s	Number of GetListing per second
	GetFileInfoNumOps	count/s	Number of GetFileInfo calls second
	SendHeartbeatNumOps	count/s	Number of SendHeart calls per second
RPC statistics (2)	RegisterDatanodeNumOps	count/s	Number of RegisterDatanode calls pe second
	BlockReportNumOps	count/s	Number of BlockRepo calls per second
	DeleteNumOps	count/s	Number of Delete calls second
	RenewLeaseNumOps	count/s	Number of RenewLea calls per second
	BlockReceivedAndDeletedNumOps	count/s	Number of BlockReceivedAndDeleted calls per second
	FsyncNumOps	count/s	Number of fsync calls second
	VersionRequestNumOps	count/s	Number of VersionReques calls per second
RPC statistics (3)	ListEncryptionZonesNumOps	count/s	Number of ListEncryptionZones calls p second
	SetPermissionNumOps	count/s	Number of SetPermise calls per second
	SetTimesNumOps	count/s	Number of SetTimes of per second
	SetSafeModeNumOps	count/s	Number of SetSafeMc calls per second
	MkdirsNumOps	count/s	Number of Mkdirs call per second



GetServerDefaultsNumOps	count/s	Number of GetServerDefaults calls pe second
GetBlockLocationsNumOps	count/s	Number of GetBlockLocati calls per second

#### HDFS - DataNode

Title	Metric	Unit	Description
Xceivers	XceiverCount	-	Number of Xceivers
	BytesWrittenMB	Bytes/s	DataNode byte write rate
	BytesReadMB	Bytes/s	DataNode byte read rate
rate	RemoteBytesReadMB	Bytes/s	Remote client byte read rate
	RemoteBytesWrittenMB	Bytes/s	Remote client byte write rate
Client connections	WritesFromRemoteClient	-	Remote client write QPS
	WritesFromLocalClient	-	Local client write QPS
	ReadsFromRemoteClient	-	Remote client read QPS
	ReadsFromLocalClient	-	Local client read QPS
Block verification failure	BlockVerificationFailures	count/s	Number of block verification failures
Disk failure	VolumeFailures	count/s	Number of disk failures
Network error	DatanodeNetworkErrors	count/s	Network error statistics
Heartbeat latency	HeartbeatsAvgTime	ms	Average heartbeat time
Heartbeat QPS	HeartbeatsNumOps	count/s	Heartbeat QPS
Packet transfer RT	SendDataPacketTransferNanosAvgTime	ms	Average time of sending packets
Block operation	ReadBlockOpNumOps	count/s	Block read OPS from DataNode
	WriteBlockOpNumOps	count/s	Block write OPS to



			DataNode
	BlockChecksumOpNumOps	count/s	Checksum OPS by DataNode
	CopyBlockOpNumOps	count/s	Block copying OPS
	ReplaceBlockOpNumOps	count/s	Block replacement OPS
	BlockReportsNumOps	count/s	Block reporting OPS
	IncrementalBlockReportsNumOps	count/s	Incremental block reporting OPS
	CacheReportsNumOps	count/s	Cache reporting OPS
	PacketAckRoundTripTimeNanosNumOps	count/s	Number of ACK round trips processed per second
Fsync operation	FsyncNanosNumOps	count/s	Number of fsync operations processed per second
Flush operation	FlushNanosNumOps	count/s	Number of flush operations processed per second
	ReadBlockOpAvgTime	ms	Average block read time
	WriteBlockOpAvgTime	ms	Average block write time
	BlockChecksumOpAvgTime	ms	Average block check time
	CopyBlockOpAvgTime	ms	Average block copy time
	ReplaceBlockOpAvgTime	ms	Average block replacement time
Block operation latency statistics	BlockReportsAvgTime	ms	Average block reporting time
	IncrementalBlockReportsAvgTime	ms	Average time of incremental block reporting
	CacheReportsAvgTime	ms	Average time of cache reporting
	PacketAckRoundTripTimeNanosAvgTime	ms	Average time of ACK round trip processing
Flush latency	FlushNanosAvgTime	ms	Average flush time



Fsync latency	FsyncNanosAvgTime	ms	Average fsync time
	RamDiskBlocksWrite	blocks/s	Total number of blocks written to memory
	RamDiskBlocksWriteFallback	blocks/s	Total number of blocks failed to be written to memory (failover to disk)
	RamDiskBlocksDeletedBeforeLazyPersisted	blocks/s	Total number of blocks deleted before the application is saved to the disk
RamDisk Blocks	RamDiskBlocksReadHits	blocks/s	Number of blocks read from memory
	RamDiskBlocksEvicted	blocks/s	Total number of blocks cleared in memory
	RamDiskBlocksEvictedWithoutRead	blocks/s	Total number of blocks retrieved from memory
	RamDiskBlocksLazyPersisted	blocks/s	Number of disk writes by lazy writer
	RamDiskBytesLazyPersisted	Bytes/s	Total number of bytes written to disk by lazy writer
RamDisk write speed	RamDiskBytesWrite	Bytes/s	Total number of bytes written to memory
JVM memory	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapCommittedM configured by JVM
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
	MemHeapCommittedM	MB	Committed size of JVM HeapMemory
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
🔗 Tencent Cloud

	MemMaxM	MB	Maximum size of memory available to JVM runtime	
	ThreadsNew	-	Number of threads in NEW status	
	ThreadsRunnable	-	Number of threads in RUNNABLE status	
IV/M threads	ThreadsBlocked	-	Number of threads in BLOCKED status	
J WIN INTEAUS	ThreadsWaiting	-	Number of threads in WAITING status	
	ThreadsTimedWaiting	-	Number of threads in TIMED WAITING status	
	ThreadsTerminated	-	Number of threads in Terminated status	
	LogFatal	-	Number of Fatal logs	
	LogError	-	Number of Error logs	
J VIVI IOUS	LogWarn	-	Number of Warn logs	
	LogInfo	-	Number of Info logs	
	YGC	-	Young GC count	
GC count	FGC	-	Full GC count	
	FGCT	S	Full GC time	
GC time	GCT	S	Garbage collection time	
	YGCT	S	Young GC time	
Memory zone proportion	SO	%	Percentage of used Survivor 0 memory	
	E	%	Percentage of used Eden memory	
	CCS	%	Percentage of used compressed class space memory	

🔗 Tencent Cloud

	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	М	%	Percentage of used Metaspace memory
Data traffia	ReceivedBytes	Bytes/s	Data receiving rate
Data tranic	SentBytes	Bytes/s	Data sending rate
QPS	RpcQueueTimeNumOps	count/s	RPC call rate
Request	RpcQueueTimeAvgTime	ms	Average RPC latency
processing latency	RpcProcessingTimeAvgTime	count/s	Average RPC request processing time
	RpcAuthenticationFailures	count/s	Number of RPC authentication failures
Authentication	RpcAuthenticationSuccesses	count/s	Number of RPC authentication successes
and authorization	RpcAuthorizationFailures	count/s	Number of RPC authorization failures
	RpcAuthorizationSuccesses	count/s	Number of RPC authorization successes
Current connections	NumOpenConnections	-	Number of current connections
Length of RPC processing queue	CallQueueLength	1	Length of current RPC processing queue
	CurrentThreadSystemTime	ms	System time
GPU time	CurrentThreadUserTime	ms	User time
Start time	StartTime	S	Process start time
Throada	PeckThreadCount	-	Peak number of threads
Threads	DaemonThreadCount	-	Number of backend threads
Read/Write	write	ms	Write time



latency	read	ms	Read time
Packet transfer QPS	DataPacketOps	count/s	Packet transfer QPS
Blocks	Related to disk information, such as `/data/qcloud/data/hdfs`	-	Blocks
Used disk capacity	Related to disk information, such as `/data/qcloud/data/hdfs`	GB	Used disk capacity
Free disk capacity	Related to disk information, such as `/data/qcloud/data/hdfs`	GB	Free disk capacity
Reserved disk capacity	Related to disk information, such as `/data/qcloud/data/hdfs`	GB	Reserved disk capacity

### HDFS - JournalNode

Title	Metric	Unit	Description
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapCommittedM configured by JVM
IVM momony	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
3 VM memory	MemHeapCommittedM	MB	Committed size of JVM HeapMemory
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
	MemMaxM	MB	Maximum size of memory available to JVM runtime
JVM threads	ThreadsNew	-	Number of threads in NEW status
	ThreadsRunnable	-	Number of threads in RUNNABLE status
	ThreadsBlocked	-	Number of threads in BLOCKED status



	ThreadsWaiting	-	Number of threads in WAITING status
	ThreadsTimedWaiting	-	Number of threads in TIMED WAITING status
	ThreadsTerminated	-	Number of threads in Terminated status
	LogFatal	-	Number of FATAL-level logs
	LogError	-	Number of ERROR-level logs
JVIVI logs	LogWarn	-	Number of WARN-level logs
	LogInfo	-	Number of INFO-level logs
	YGC	-	Young GC count
GC count	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	SO	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
Monory room avenuetion	CCS	%	Percentage of used compressed class space memory
Memory zone proportion	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
Data traffia	ReceivedBytes	Bytes/s	Data receiving rate
	SentBytes	Bytes/s	Data sending rate
Request processing latency	RpcQueueTimeAvgTime	ms	Average RPC latency



Authentication and authorization	RpcAuthenticationFailures count		Number of RPC authentication failures
	RpcAuthenticationSuccesses	count/s	Number of RPC authentication successes
	<b>RpcAuthorizationFailures</b>	count/s	Number of RPC authorization failures
	RpcAuthorizationSuccesses	count/s	Number of RPC authorization successes
Current connections	NumOpenConnections	-	Number of current connections
Length of RPC processing queue	CallQueueLength	1	Length of current RPC processing queue
CDILItimo	CurrentThreadSystemTime	ms	System time
	CurrentThreadUserTime	ms	User time
Start time	StartTime	S	Process start time
Throade	PeckThreadCount	-	Peak number of threads
IIIEaus	DaemonThreadCount	-	Number of backend threads

### HDFS - ZKFC

Title	Metric	Unit	Description	
GC count	YGC	-	Young GC count	
de count	FGC	-	Full GC count	
	FGCT	S	Full GC time	
GC time	GCT	S	Garbage collection time	
	YGCT	S	Young GC time	
Memory zone proportion	S0	%	Percentage of used Survivor 0 memory	
	Е	%	Percentage of used Eden memory	
	CCS	%	Percentage of used compressed class space memory	
	S1	%	Percentage of used Survivor 1 memory	



0	%	Percentage of used Old memory
М	%	Percentage of used Metaspace memory

#### **HDFS-Router**

Title	Metric	Unit	Description
ALTER TABLE request duration	HIVE.HMS.API_ALTER_TABLE	ms	Average duration of ALTER TABLE requests
ALTER TABLE WITH ENV CONTEXT request duration	HIVE.HMS.API_ALTER_TABLE_WITH_ENV_CONTEXT	ms	Average duration of ALTER TABLE WITH ENV CONTEXT requests
CREATE TABLE request duration	HIVE.HMS.API_CREATE_TABLE	ms	Average duration of CREATE TABLE requests
CREATE TABLE WITH ENV CONTEXT request duration	HIVE.HMS.API_CREATE_TABLE_WITH_ENV_CONTEXT	ms	Average duration of CREATE TABLE WITH ENV CONTEXT requests
DROP TABLE request duration	HIVE.HMS.API_DROP_TABLE	ms	Average duration of DROP TABLE requests
DROP TABLE WITH ENV CONTEXT request duration	HIVE.HMS.API_DROP_TABLE_WITH_ENV_CONTEXT	ms	Average duration of DROP TABLE WITH



			ENV CONTEXT requests
GET TABLE request duration	HIVE.HMS.API_GET_TABLE	ms	Average duration of GET TABLE requests
GET TABLES request duration	HIVE.HMS.API_GET_TABLES	ms	Average duration of GET TABLES requests
GET MULTI TABLE request duration	HIVE.HMS.API_GET_MULTI_TABLE	ms	Average duration of GET MULTI TABLE requests
GET TABLE REQ request duration	HIVE.HMS.API_GET_TABLE_REQ	ms	Average duration of GET TABLE REQ requests
GET DATABASE request duration	HIVE.HMS.API_GET_DATABASE	ms	Average duration of GET DATABASE requests
GET DATABASES request duration	HIVE.HMS.API_GET_DATABASES	ms	Average duration of GET DATABASES requests
GET ALL DATABASES request duration	HIVE.HMS.API_GET_ALL_DATABASES	ms	Average duration of GET ALL DATABASES requests
GET ALL FUNCTIONS request	HIVE.HMS.API_GET_ALL_FUNCTIONS	ms	Average duration of



duration			GET ALL FUNCTIONS requests
Current number of active CREATE TABLE requests	HIVE.HMS.ACTIVE_CALLS_API_CREATE_TABLE	-	Current number of active CREATE TABLE requests
Current number of active DROP TABLE requests	HIVE.HMS.ACTIVE_CALLS_API_DROP_TABLE	-	Current number of active DROP TABLE requests
Current number of active ALTER TABLE requests	HIVE.HMS.ACTIVE_CALLS_API_ALTER_TABLE	-	Current number of active ALTER TABLE requests

# **YARN** Monitoring Metrics

Last updated : 2025-02-12 16:52:06

### **YARN - overview**

Title	Metric		Description
	NumActiveNMs	-	Number of live NodeManagers
Nodoc	NumDecommissionedNMs	-	Number of decommissioned NodeManagers
NUCLES	NumLostNMs	-	Number of lost NodeManagers
	NumUnhealthyNMs		Number of unhealthy NodeManagers
	AllocatedVCores	-	Number of allocated VCores in the current queue
	ReservedVCores	-	Number of reserved VCores in the current queue
CPU cores	AvailableVCores	-	Number of available VCores in the current queue
	PendingVCores	-	Number of pending VCores in resource requests in the current queue
Total applications	AppsSubmitted	-	Number of submitted jobs in the current queue
	AppsRunning	-	Number of running jobs in the current queue
	AppsPending	-	Number of pending jobs in the current queue
	AppsCompleted	-	Number of completed jobs in the current queue
	AppsKilled	-	Number of killed jobs in the current queue
	AppsFailed	-	Number of failed jobs in the



			current queue
	ActiveApplications	-	Number of active jobs in the current queue
	running_0	-	Number of running jobs in the current queue that have run for less than 60 minutes
	running_60	-	Number of running jobs in the current queue that have run for 60–300 minutes
	running_300	-	Number of running jobs in the current queue that have run for 300–1,440 minutes
	running_1440	-	Number of running jobs in the current queue that have run for more than 1,440 minutes
	AllocatedMB	MB	Amount of allocated memory in the current queue
	AvailableMB	MB	Amount of available memory in the current queue
Memory size	PendingMB	MB	Amount of pending memory in resource requests in the current queue
	ReservedMB	MB	Amount of reserved memory in the current queue
	AllocatedContainers	-	Number of allocated containers in the current queue
Containers	PendingContainers	-	Number of pending containers in resource requests in the current queue
	ReservedContainers	-	Number of reserved containers in the current queue
Total allocated/released containers	AggregateContainersAllocated	-	Total number of allocated containers in the current queue
	AggregateContainersReleased	-	Total number of released



			containers in the current queue
Users	ActiveUsers	-	Number of active users in the current queue
	allocatedMB	MB	Amount of allocated memory in the cluster
Momory	availableMB	MB	Amount of available memory in the cluster
wemory	reservedMB	MB	Amount of reserved memory in the cluster
	totalMB	MB	Total amount of memory in the cluster
	completed	-	Number of completed jobs in the cluster during the statistical period
	failed	-	Number of failed jobs in the cluster during the statistical period
	killed	-	Number of killed jobs in the cluster during the statistical period
Αρριτατιοτις	pending	-	Number of pending jobs in the cluster during the statistical period
	running	-	Number of running jobs in the cluster during the statistical period
	submitted	-	Number of submitted jobs in the cluster during the statistical period
Containers	containersAllocated	-	Number of allocated containers in the cluster
	containersPending	-	Number of pending containers in the cluster
	containersReserved	-	Number of reserved containers in the cluster
Memory utilization	usageRatio	%	Current memory utilization of the cluster
Cores	allocatedVirtualCores	-	Number of allocated CPU cores in



			the cluster
	availableVirtualCores	-	Number of available CPU cores in the cluster
	reservedVirtualCores	-	Number of reserved CPU cores in the cluster
	totalVirtualCores	-	Total number of CPU cores in the cluster
CPU utilization	usageRatio	%	Current CPU utilization of the cluster
Launched AMs	AMLaunchDelayNumOps	-	Launched AMs
Average time for RM to launch AM	AMLaunchDelayAvgTime	ms	Average time for RM to launch AM
Total registered AMs	AMRegisterDelayNumOps	-	Total registered AMs
Average time for AM to register with RM	AMRegisterDelayAvgTime	ms	Average time for AM to register with RM
Queue CPU utilization	YARN.RM.QUEUE.VCORES.RATIO	-	Utilization of CPU allocated for the current queue
Queue memory utilization	YARN.RM.QUEUE.MEM.RATIO	-	Utilization of memory allocated for the current queue

### YARN - ResourceManager

Title	Metric	Unit	Description
RPC authentications/authorizations	<b>RpcAuthenticationFailures</b>	-	Number of failed RPC authentications
	RpcAuthenticationSuccesses	-	Number of successful RPC authentications
	<b>RpcAuthorizationFailures</b>	-	Number of failed RPC authorizations
	<b>RpcAuthorizationSuccesses</b>	-	Number of



			successful RPC authorizations
Data reasived/cent by DDC	ReceivedBytes	bytes/s	Amount of data received by RPC
Data received/sent by hFC	SentBytes	bytes/s	Amount of data sent by RPC
RPC connections	NumOpenConnections	-	Current number of open connections
	RpcProcessingTimeNumOps	-	Number of RPC requests
RPC requests	RpcQueueTimeNumOps	-	Number of RPC requests
RPC queue length	CallQueueLength	-	Length of the current RPC queue
Average RPC processing	RpcProcessingTimeAvgTime	S	Average RPC request processing time
une	RpcQueueTimeAvgTime	S	Average time of RPC in the queue
	YGC	-	Young GC count
GC count	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
Memory zone proportion	SO	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
	CCS	%	Percentage of used compressed class space memory

	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
	ThreadsNew	-	Number of threads in NEW status
	ThreadsRunnable	-	Number of threads in RUNNABLE status
IVM threads	ThreadsBlocked	-	Number of threads in BLOCKED status
o vivi tili eads	ThreadsWaiting	-	Number of threads in WAITING status
	ThreadsTimedWaiting	-	Number of threads in TIMED WAITING status
	ThreadsTerminated	-	Number of threads in Terminated status
	LogFatal	-	Number of Fatal logs
	LogError	-	Number of Error logs
JVM logs	LogWarn	-	Number of Warn logs
	LogInfo	-	Number of Info logs
JVM memory	MemNonHeapUsedM	MB	Non-heap memory size used by process
	MemNonHeapCommittedM	MB	Non-heap memory size committed to process
	MemHeapUsedM	MB	Heap memory size used by process



	MemHeapCommittedM	MB	Heap memory size committed to process
	MemHeapMaxM	MB	Maximum heap memory size available to process
	MemMaxM	MB	Maximum memory size available to process
CPU utilization	ProcessCpuLoad	%	CPU utilization
Cumulative CPU usage time	ProcessCpuTime	ms	Cumulative CPU usage time
	MaxFileDescriptorCount	-	Maximum number of file descriptors
File descriptors	OpenFileDescriptorCount	-	Number of opened file descriptors
Process execution duration	Uptime	S	Process execution duration
Worker threads	DaemonThreadCount	-	Number of daemon threads in the process
	ThreadCount	-	Number of threads in the process
Node status	haState	1:Active,0:Standby	ResourceManager active/standby status
Active/Standby switch	switchOccurred	-	ResourceManager active/standby switch

### YARN - JobHistoryServer

Title	Metric	Unit	Description
JVM threads	ThreadsNew	-	Number of threads in NEW status



	ThreadsRunnable	-	Number of threads in RUNNABLE status
	ThreadsBlocked	-	Number of threads in BLOCKED status
	ThreadsWaiting	-	Number of threads in WAITING status
	ThreadsTimedWaiting	-	Number of threads in TIMED WAITING status
	ThreadsTerminated	-	Number of threads in Terminated status
	LogFatal	-	Number of FATAL-level logs
	LogError	-	Number of ERROR-level logs
JVIVI logs	LogWarn	-	Number of WARN-level logs
	LogInfo	-	Number of INFO-level logs
	MemNonHeapUsedM	MB	Non-heap memory size used by process
	MemNonHeapCommittedM	MB	Non-heap memory size committed to process
	MemHeapUsedM	MB	Heap memory size used by process
JVW memory	MemHeapCommittedM	MB	Heap memory size committed to process
	MemHeapMaxM	MB	Maximum heap memory size available to process
	MemMaxM	MB	Maximum memory size available to process
	YGC	-	Young GC count
GC count	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
Memory zone	SO	%	Percentage of used Survivor 0 memory
proportion	E	%	Percentage of used Eden memory
	CCS	%	Percentage of used compressed class space memory



	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
CPU utilization	ProcessCpuLoad	%	CPU utilization
Cumulative CPU usage time	ProcessCpuTime	ms	Cumulative CPU usage time
File descriptors	MaxFileDescriptorCount	-	Maximum number of file descriptors
	OpenFileDescriptorCount	-	Number of opened file descriptors
Process execution duration	Uptime	S	Process execution duration
Worker threads	DaemonThreadCount	-	Number of daemon threads in the process
	ThreadCount	-	Number of threads in the process

#### YARN - NodeManager

Title	Metric	Unit	Description
	YGC	-	Young GC count
Ge count	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
Memory zone proportion	SO	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
	CCS	%	Percentage of used compressed class space memory
	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory



	Μ	%	Percentage of used Metaspace memory
	ThreadsNew	-	Number of threads in NEW status
	ThreadsRunnable	-	Number of threads in RUNNABLE status
	ThreadsBlocked	-	Number of threads in BLOCKED status
JVM threads	ThreadsWaiting	-	Number of threads in WAITING status
	ThreadsTimedWaiting	-	Number of threads in TIMED WAITING status
	ThreadsTerminated	-	Number of threads currently in TERMINATED status
	LogFatal	-	Number of FATAL-level logs
IVM logs	LogError	-	Number of ERROR-level logs
o vivi logs	LogWarn	-	Number of WARN-level logs
	LogInfo	-	Number of INFO-level logs
	MemNonHeapUsedM	MB	Non-heap memory size used by process
	MemNonHeapCommittedM	MB	Non-heap memory size committed to process
	MemHeapUsedM	MB	Heap memory size used by process
JVM memory	MemHeapCommittedM	MB	Heap memory size committed to process
	MemHeapMaxM	MB	Maximum heap memory size available to process
	MemMaxM	MB	Maximum memory size available to process
Total containers	ContainersLaunched	-	Number of launched containers
	ContainersCompleted	-	Number of completed containers



	ContainersFailed	-	Number of failed containers
	ContainersKilled	-	Number of killed containers
	ContainersIniting	-	Number of containers being initialized
	ContainersRunning	-	Number of running containers
	AllocatedContainers	-	Number of containers allocated by NodeManager
Average container launch time	ContainerLaunchDurationAvgTime	ms	Average container launch time
Container launches	ContainerLaunchDurationNumOps	-	Container launches
CDU agreg	AvailableVCores	-	Number of VCores available to NodeManager
CFU coles	AllocatedVCores	-	Number of VCores allocated by NodeManager
Momory gize	AllocatedGB	GB	Amount of memory allocated by NodeManager
Memory size	AvailableGB	GB	Amount of memory available to NodeManager
CPU utilization	ProcessCpuLoad	%	CPU utilization
Cumulative CPU usage time	ProcessCpuTime	ms	Cumulative CPU usage time
Eilo docoriptoro	MaxFileDescriptorCount	-	Maximum number of file descriptors
	OpenFileDescriptorCount	-	Number of opened file descriptors
Process execution duration	Uptime	S	Process execution duration
Worker threads	DaemonThreadCount	-	Number of daemon threads in the process
	ThreadCount	-	Number of threads in the process

# **ZooKeeper Monitoring Metrics**

Last updated : 2023-12-27 14:48:59

#### ZooKeeper

Title	Metric	Unit	Description
CC count	YGC	-	Young GC count
GC Count	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	SO	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
Memory zone proportion	CCS	%	Percentage of used compressed class space memory
	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	М	%	Percentage of used Metaspace memory
JVM memory	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
	MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM



	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
	MemHeapInitM	MB	Size of initial JVM HeapMem
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMem
CPU utilization	ProcessCpuLoad	%	CPU utilization
File descriptors	zk_max_file_descriptor_count	-	Maximum number of file descriptors
The descriptors	zk_open_file_descriptor_count	-	Number of opened file descriptors
Cumulative CPU usage time	ProcessCpuTime	ms	Cumulative CPU usage time
Process execution duration	Uptime	S	Process execution duration
Workor throads	DaemonThreadCount	-	Number of daemon threads
Worker Infeads	ThreadCount	-	Total number of threads
Connections	zk_num_alive_connections	-	Number of current connections
	zk_avg_latency	ms	Average ZooKeeper processing latency
Latency	zk_max_latency	ms	Maximum ZooKeeper processing latency
	zk_min_latency	ms	Minimum ZooKeeper processing latency
	zk_watch_count	-	Number of ZooKeeper watches
Zondes	zk_znode_count	-	Number of ZooKeeper znodes
	zk_ephemerals_count	-	Number of ephemeral ZooKeeper nodes
Data size	zk_approximate_data_size	Byte	Volume of data stored in ZooKeeper

Node status	zk_server_state	1: Master. 0: Slave. 2: Standalone	ZooKeeper node type
Received/Sent packets	zk_packets_received	count/s	ZooKeeper packet receiving rate
	zk_packets_sent	count/s	ZooKeeper packet sending rate
Queuing requests	zk_outstanding_requests	-	Number of queuing requests

# **HBase Monitoring Metrics**

Last updated : 2023-12-27 14:49:18

#### **HBase - overview**

Title	Metric	Unit	Description
	ritCount	-	Number of regions in transition
Cluster regions in RIT status	ritCountOverThreshold	-	Number of regions that have been in transition for more than the threshold time
Cluster RIT time	ritOldestAge	ms	Age of the longest region in transition
Average number of regions per RegionServer	averageLoad	-	Average number of regions per RegionServer
Cluster Degion Servere	numRegionServers	-	Number of live RegionServers
Cluster RegionServers	numDeadRegionServers	-	Number of dead RegionServers
Data read/written from/to	receivedBytes	bytes/s	Amount of data received by cluster
HMaster	sentBytes	bytes/s	Amount of data sent by cluster
Total cluster API requests	clusterRequests	count/s	Total number of cluster requests
Cluster assignment	Assign_num_ops	-	Number of region assignments
manager operation	BulkAssign_num_ops	-	Number of bulk region assignments
Cluster load balancing operations	BalancerCluster_num_ops	-	Number of cluster load balancing operations

#### HBase - HMaster

Title	Metric	Unit	Description
GC count	YGC	-	Young GC count
	FGC	-	Full GC count
GC time	FGCT	S	Full GC time



	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	S0	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
Memory zone	CCS	%	Percentage of used compressed class space memory
proportion	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
	LogFatal	-	Number of Fatal logs
	LogError	-	Number of Error logs
JVINI IOUS	LogWarn	-	Number of Warn logs
	LogInfo	-	Number of Info logs
	MemNonHeapUsedM	MB	Non-heap memory size used by process
	MemNonHeapCommittedM	MB	Non-heap memory size committed to process
	MemHeapUsedM	MB	Heap memory size used by process
JVM memory	MemHeapCommittedM	MB	Heap memory size committed to process
	MemHeapMaxM	MB	Maximum heap memory size available to process
	MemMaxM	MB	Maximum memory size available to process
JVM threads	ThreadsNew	-	Number of threads in NEW status
	ThreadsRunnable	-	Number of threads in RUNNABLE status
	ThreadsBlocked	-	Number of threads in BLOCKED status
	ThreadsWaiting	-	Number of threads in WAITING status



	ThreadsTimedWaiting	-	Number of threads in TIMED WAITING status
	ThreadsTerminated	-	Number of threads currently in TERMINATED status
RPC connections	numOpenConnections	-	Number of RPC connections
	FailedSanityCheckException	-	Number of FailedSanityCheckException exceptions
	NotServingRegionException	-	Number of NotServingRegionException exceptions
RPC exceptions	OutOfOrderScannerNextException	-	Number of OutOfOrderScannerNextException exceptions
	RegionMovedException	-	Number of RegionMovedException exceptions
	RegionTooBusyException	-	Number of RegionTooBusyException exceptions
	UnknownScannerException	-	Number of UnknownScannerException exceptions
RPC queue requests	numCallsInPriorityQueue	-	Number of requests in the general queue
	numCallsInReplicationQueue	-	Number of RPC requests in the replication queue
Process start time	masterActiveTime	S	Master active time
Process start time	masterStartTime	S	Master process start time

## HBase - RegionServer

Title	Metric	Unit	Description
GC count	YGC	-	Young GC count
	FGC	-	Full GC count
GC time	FGCT	S	Full GC time
	GCT	S	Garbage collection time



	YGCT	S	Young GC time
	SO	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
Memory zone	CCS	%	Percentage of used compressed class space memory
proportion	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
	LogFatal	-	Number of Fatal logs
IV/M logs	LogError	-	Number of Error logs
J VIVI IOGS	LogWarn	-	Number of Warn logs
	LogInfo	-	Number of Info logs
	MemNonHeapUsedM	MB	Non-heap memory size used by process
	MemNonHeapCommittedM	MB	Non-heap memory size committed to process
	MemHeapUsedM	MB	Heap memory size used by process
JVM memory	MemHeapCommittedM	MB	Heap memory size committed to process
	MemHeapMaxM	MB	Maximum heap memory size available to process
	MemMaxM	MB	Maximum memory size available to process
JVM threads	ThreadsNew	-	Number of threads in NEW status
	ThreadsRunnable	-	Number of threads in RUNNABLE status
	ThreadsBlocked	-	Number of threads in BLOCKED status
	ThreadsWaiting	-	Number of threads in WAITING status



	ThreadsTimedWaiting	-	Number of threads in TIMED WAITING status
	ThreadsTerminated	-	Number of threads currently in TERMINATED status
Regions	regionCount	-	Number of regions
Region localization	percentFilesLocal	%	Percentage of HFiles on the local HDFS data node in the region
Region replica localization	percentFilesLocalSecondaryRegions	%	Percentage of HFiles on the local HDFS data node in the region replica
RPC	authenticationFailures	-	Number of RPC authentication failures
authentications	authenticationSuccesses	-	Number of RPC authentication successes
RPC connections	numOpenConnections	-	Number of RPC connections
	FailedSanityCheckException	-	Number of FailedSanityCheckException exceptions
	NotServingRegionException	-	Number of NotServingRegionException exceptions
RPC exceptions	OutOfOrderScannerNextException	-	Number of OutOfOrderScannerNextException exceptions
	RegionMovedException	-	Number of RegionMovedException exceptions
	RegionTooBusyException	-	Number of RegionTooBusyException exceptions
	UnknownScannerException	-	Number of UnknownScannerException exceptions
RPC handlers	numActiveHandler	-	Number of active RPC handlers
	numActiveWriteHandler	-	Number of active read RPC handlers
	numActiveReadHandler	-	Number of active write RPC handlers



	numActiveScanHandler	-	Number of active scan RPC handlers
	numCallsInPriorityQueue	-	Number of requests in the priority queue
	numCallsInReplicationQueue	-	Number of RPC requests in the replication queue
RPC queue	numCallsInPriorityQueue	-	Number of requests in the general queue
requests	numCallsInWriteQueue	-	Number of RPC calls in the write call queue
	numCallsInReadQueue	-	Number of RPC calls in the read call queue
	numCallsInScanQueue	-	Number of RPC calls in the scan call queue
WAL files	hlogFileCount	-	Number of WAL files
WAL file size	hlogFileSize	Byte	WAL file size
MemStore size	memStoreSize	MB	MemStore size
Stores	storeCount	-	Number of stores
StoreFiles	storeFileCount	-	Number of StoreFiles
StoreFile size	storeFileSize	MB	StoreFile size
Disk write rate	flushedCellsSize	bytes/s	Disk write rate
	Append_mean	ms	Average Append latency
	Replay_mean	ms	Average Replay latency
Average latency	Get_mean	ms	Average GET latency
,	updatesBlockedTime	ms	Number of milliseconds updates have been blocked so the memstore can be flushed
RegionServer disk writes	FlushTime_num_ops	-	Number of MemStore flushes



Requests in	splitQueueLength	-	Length of the split queue
queue	compactionQueueLength	-	Length of the compaction queue
	flushQueueLength	-	Length of the region flush queue
Replay operations	Replay_num_ops	-	Number of Replay operations
	slowAppendCount	-	Number of Append requests that took over 1s to complete
	slowDeleteCount	-	Number of Delete requests that took over 1s to complete
Slow operations	slowGetCount	-	Number of Get requests that took over 1s to complete
	slowIncrementCount	-	Number of Increment requests that took over 1s to complete
	slowPutCount	-	Number of Put requests that took over 1s to complete
Split roquest	splitRequestCount	-	Number of split requested
Split request	splitSuccessCount	-	Number of successfully executed splits
	blockCacheCount	-	Number of blocks in the block cache
Cache blocks	blockCacheHitCount	-	Number of block cache hits
	blockCacheMissCount	-	Number of block cache misses
Cache read hit rate	blockCacheExpressHitPercent	%	Cache read hit rate
Memory size used by the cache block	blockCacheSize	Byte	Memory size used by the cache block
	staticBloomSize	Byte	Uncompressed size of static bloom filters
Index size	staticIndexSize	Byte	Uncompressed size of static indexes
	storeFileIndexSize	Byte	Size of indexes in StoreFiles on disk
Received bytes	receivedBytes	bytes/s	Received bytes



	sentBytes	bytes/s	Sent bytes
	Total	count/s	Total number of requests. When there are scan requests, this value will be smaller than the sum of read and write requests
	Read	count/s	Number of read requests
	Write	count/s	Number of write requests
	Append_num_ops	count/s	Number of Append requests
Read and write	Mutate_num_ops	count/s	Number of Mutate requests
10400510	Delete_num_ops	count/s	Number of Delete requests
	Increment_num_ops	count/s	Number of Increment requests
	Get_num_ops	count/s	Number of Get requests
	Put_num_ops	count/s	Number of Put requests
	ScanTime_num_ops	count/s	Scan requests (time)
	ScanSize_num_ops	count/s	Scan requests (size)
Mutations	mutationsWithoutWALCount	-	Number of mutations
Mutation size	mutationsWithoutWALSize	Byte	Mutation size
Process start time	regionServerStartTime	S	Process start time
Log sync	source.sizeOfLogQueue	-	Total length of synced logs
Sync duration	source.ageOfLastShippedOp	ms	Sync duration
Boquests	ReadRequestCount	count/s	Read requests/s
nequesis	WriteRequestCount	count/s	Write requests/s
Bequeets	Read	count/s	Read requests/s
าายนยรเร	Write	count/s	Write requests/s
Store size	memstoreSize	Byte	MemStore size
SIDIE SIZE	storeFileSize	Byte	StoreFile size

Table-level request latency	getTime_99th_percentile	ms	99th percentile of request processing latency	
	scanTime_99th_percentile	ms	99th percentile of request processing latency	
	putTime_99th_percentile	ms	99th percentile of request processing latency	
	incrementTime_99th_percentile	ms	99th percentile of request processing latency	
	appendTime_99th_percentile	ms	99th percentile of request processing latency	
	deleteTime_99th_percentile	ms	99th percentile of request processing latency	
Request processing	99th_percentile	ms	99th percentile of request processing latency	
latency	99.9th_percentile	ms	99.9% request processing latency	
Request queueing	99th_percentile	ms	99th percentile of request queueing latency	
latency	99.9th_percentile	ms	99.9% request queueing latency	
	max	bytes	Maximum scan size	
Scan size	mean	bytes	Average scan size	
	min	bytes	Minimum scan size	
	max	S	Maximum scan time	
Scan time	mean	S	Average scan time	
	min	S	Minimum scan time	

# **Hive Monitoring Metrics**

Last updated : 2023-12-27 14:49:37

#### Hive - HiveMetaStore

Title	Metric	Unit	Description
	YGC	-	Young GC c
GC count	FGC	-	Full GC cou
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection tir
	YGCT	S	Young GC t
Memory zone proportion	S0	%	Percentage used Surviv memory
	E	%	Percentage used Eden memory
	CCS	%	Percentage used compressec class space memory
	S1	%	Percentage used Surviv memory
	0	%	Percentage used Old m
	Μ	%	Percentage used Metas memory
JVM memory	MemHeapUsedM	MB	Size of HeapMemo

🔗 Tencent Cloud

			currently us JVM
	MemHeapCommittedM	MB	Size of HeapMemo committed t JVM
	MemHeapMaxM	MB	Size of HeapMemo configured t JVM
	MemHeapInitM	MB	Size of initia JVM HeapN
	MemNonHeapUsedM	MB	Size of NonHeapM currently us JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMe currently committed k JVM
	MemNonHeapInitM	MB	Size of initia JVM NonHeapMe
File	OpenFileDescriptorCount	-	Number of opened file descriptors
descriptors	MaxFileDescriptorCount	-	Maximum number of fi descriptors
CPU utilization	ProcessCpuLoad	%	Process CP utilization
	SystemCpuLoad	%	System CPl utilization
Percentage of CPU	CPURate	seconds/second	Percentage CPU usage



usage time			
Worker	DaemonThreadCount	-	Number of daemon three
threads	ThreadCount	-	Total numbe threads
Cumulative CPU usage time	ProcessCpuTime	ms	Cumulative usage time
Process execution duration	Uptime	S	Process execution duration
GC extra sleep time	ExtraSleepTime	ms/s	GC extra sle time
ALTER TABLE request duration	HIVE.HMS.API_ALTER_TABLE	ms	Average du of ALTER T requests
ALTER TABLE WITH ENV CONTEXT request duration	HIVE.HMS.API_ALTER_TABLE_WITH_ENV_CONTEXT	ms	Average du of ALTER T WITH ENV CONTEXT requests
CREATE TABLE request duration	HIVE.HMS.API_CREATE_TABLE	ms	Average du of CREATE TABLE requ
CREATE TABLE WITH ENV CONTEXT request duration	HIVE.HMS.API_CREATE_TABLE_WITH_ENV_CONTEXT	ms	Average dun of CREATE TABLE WIT ENV CONT requests
DROP TABLE request duration	HIVE.HMS.API_DROP_TABLE	ms	Average du of DROP T <i>I</i> requests
		1	1



DROP TABLE WITH ENV CONTEXT request duration	HIVE.HMS.API_DROP_TABLE_WITH_ENV_CONTEXT	ms	Average dui of DROP T/ WITH ENV CONTEXT requests
GET TABLE request duration	HIVE.HMS.API_GET_TABLE	ms	Average du of GET TAE requests
GET TABLES request duration	HIVE.HMS.API_GET_TABLES	ms	Average du of GET TAE requests
GET MULTI TABLE request duration	HIVE.HMS.API_GET_MULTI_TABLE	ms	Average du of GET MUI TABLE requ
GET TABLE REQ request duration	HIVE.HMS.API_GET_TABLE_REQ	ms	Average du of GET TAE REQ reques
GET DATABASE request duration	HIVE.HMS.API_GET_DATABASE	ms	Average du of GET DATABASE requests
GET DATABASES request duration	HIVE.HMS.API_GET_DATABASES	ms	Average du of GET DATABASE requests
GET ALL DATABASES request duration	HIVE.HMS.API_GET_ALL_DATABASES	ms	Average du of GET ALL DATABASE requests
GET ALL FUNCTIONS request duration	HIVE.HMS.API_GET_ALL_FUNCTIONS	ms	Average du of GET ALL FUNCTION requests
Current number of	HIVE.HMS.ACTIVE_CALLS_API_CREATE_TABLE	-	Current num of active

active CREATE TABLE requests			CREATE T/ requests
Current number of active DROP TABLE requests	HIVE.HMS.ACTIVE_CALLS_API_DROP_TABLE	-	Current nurr of active DF TABLE requ
Current number of active ALTER TABLE requests	HIVE.HMS.ACTIVE_CALLS_API_ALTER_TABLE	-	Current num of active AL TABLE requ

#### Hive - HiveServer2

Title	Metric	Unit	Description
CC count	YGC	-	Young GC count
Ge count	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
Memory zone proportion	S0	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
	CCS	%	Percentage of used compressed class space memory


	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
JVM memory	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
	MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
	MemHeapInitM	MB	Size of initial JVM HeapMem
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMem
Percentage of used heap memory	MemHeapUsedRate	%	Proportion of the heap memory



			size currently used by JVM to the heap memory size configured for JVM
CPU utilization	ProcessCpuLoad	%	CPU utilization
Percentage of CPU usage time	CPUUsedRate	seconds/second	Percentage of CPU usage time
File descriptors	MaxFileDescriptorCount	-	Maximum number of file descriptors
	OpenFileDescriptorCount	-	Number of opened file descriptors
Cumulative CPU usage time	ProcessCpuTime	ms	Cumulative CPU usage time
Process execution duration	Uptime	S	Process execution duration
	DaemonThreadCount	-	Number of daemon threads
worker threads	ThreadCount	-	Total number of threads
Driver execution latency	99th_percentile	ms	99th percentile of driver execution latency
	Avg	ms	Average driver execution latency
Open connections	NumOpenConnections	-	Number of open connections
Current size of the pool of async HiveServer2 threads	PoolSize	-	Current size of the pool of async HiveServer2 threads
Current size of the queue	QueueSize	-	Current size of



		the queue for async HiveServer2 operations
Closed	-	Number of closed operations
Finished	-	Number of completed operations
Canceled	-	Number of canceled operations
Error	-	Number of erroneous operations
ExtraSleepTime	ms/s	GC extra sleep time
	Count	Current number of serializePlan requests
	Count	Current number of semanticAnalyze requests
HIVE.H2.ACTIVE.CALLS.API	Count	Current number of runtask requests
	Count	Current number of releaseLocks requests
	Count	Current number of getSplits requests
HIVE.H2.SQL.OPERATION.PENDING	ms	Average duration of SQL tasks in PENDING status
	Closed Finished Canceled Error ExtraSleepTime HIVE.H2.ACTIVE.CALLS.API	Closed-Finished-Finished-Canceled-Error-ExtraSleepTimems/sFunctionCountCount-



Duration of SQL tasks in RUNNING status	HIVE.H2.SQL.OPERATION.RUNNING	ms	Average duration of SQL tasks in RUNNING status
Current number of active users	HIVE.H2.SQL.OPERATION	Count	Current number of active users
Query execution time	HIVE.H2.EXECUTING.QUERIES	ms	Average query execution time
Query submission time	HIVE.H2.SUBMITTED.QUERIES	ms	Query submission time
Number of submitted Hive on MR jobs	HIVE.H2.MR.TASKS	Count	Number of submitted Hive on MR jobs
Number of submitted Hive on Spark jobs	HIVE.H2.SPARK.TASKS	Count	Number of submitted Hive on Tez jobs
Number of submitted Hive on Tez jobs	HIVE.H2.TEZ.TASKS	Count	Number of submitted Hive on Spark jobs
Failed queries	HIVE.H2.FAILED.QUERIES.RATE	Count/min	Number of failed queries per minute
Number of worker threads	HIVE.H2.THREAD.COUNT	-	Number of JVM blocked threads
		-	Number of JVM terminate threads
		-	Number of JVM deadlock threads
		-	Number of JVM new threads
		-	Number of JVM runnable threads
		-	Number of JVM timed_waiting threads

		-	Number of JVM waiting threads
Number of sessions	HIVE.H2.OPEN.SESSIONS	-	Number of open sessions
Current number of active sessions	HIVE.H2.ACTIVE.SESSIONS	-	Number of active sessions

## Hive - HiveWebHcat

Title	Metric	Unit	Description		
	YGC	-	Young GC count		
Ge count	FGC	-	Full GC count		
	FGCT	S	Full GC time		
GC time	GCT	S	Garbage collection time		
	YGCT	S	Young GC time		
	S0	%	Percentage of used Survivor 0 memory		
	Е	%	Percentage of used Eden memory		
Moment zone propertion	CCS	%	Percentage of used compressed class space memory		
Memory zone proportion	S1	%	Percentage of used Survivor 1 memory		
	0	%	Percentage of used Old memory		
	М	%	Percentage of used Metaspace memory		

# Spark Monitoring Metrics

Last updated : 2023-12-27 14:49:54

## Spark - history server

Title	Metric	Unit	Description	
	YGC	-	Young GC count	
Ge count	FGC	-	Full GC count	
	FGCT	S	Full GC time	
GC time	GCT	S	Garbage collection time	
	YGCT	S	Young GC time	
	S0	%	Percentage of used Survivor 0 memory	
	E	%	Percentage of used Eden memory	
Momory zono proportion	CCS	%	Percentage of used compressed class space memory	
Memory zone proportion	S1	%	Percentage of used Survivor 1 memory	
	0	%	Percentage of used Old memory	
	М	%	Percentage of used Metaspace memory	

# **Presto Monitoring Metrics**

Last updated : 2023-12-27 14:50:22

#### **Presto - overview**

Title	Metric	Unit	Description
	Active	-	Number of active nodes
Nodes	Total	-	Total number of nodes
	Failed	-	Number of failed nodes
Query	RunningQueries	-	Total number of running queries
Query	QueuedQueries	-	Total number of waiting queries
	FailedQueries	count/min	Total number of failed queries
	AbandonedQueries	count/min	Total number of abandoned queries
Query frequency	CanceledQueries	count/min	Total number of canceled queries
	CompletedQueries	count/min	Total number of completed queries
	StartedQueries	count/min	Total number of started queries
Data volume input/output per	InputDataSizeOneMinute	GB/min	Data input rate
minute	OutputDataSizeOneMinute	GB/min	Data output rate

#### Presto - worker

Title	Metric	Unit	Description
GC count	YGC	-	Young GC count



	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	SO	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
Memory zone	CCS	%	Percentage of used compressed class space memory
proportion	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
JVM memory	MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
	MemHeapInitM	MB	Size of initial JVM HeapMem
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMem
Data input/output	InputDataSize.OneMinute.Rate	GB/min	Data input rate
rate	OutputDataSize.OneMinute.Rate	GB/min	Data output rate
Processes	PeakThreadCount	-	Peak number of threads
	ThreadCount	-	Number of threads



	DaemonThreadCount	-	Number of backend threads
Process execution duration	Uptime	S	Process execution duration
Process start time	StartTime	S	Process start time
File descriptors	MaxFileDescriptorCount	-	Maximum number of file descriptors
	OpenFileDescriptorCount	-	Number of opened file descriptors

## Presto - coordinator

Title	Metric	Unit	Description
GC count	YGC	-	Young GC count
	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	SO	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
Memory zone	CCS	%	Percentage of used compressed class space memory
ριοροιτιστ	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
JVM memory	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
	MemHeapCommittedM	MB	Size of HeapMemory currently committed



			by JVM
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
	MemHeapInitM	MB	Size of initial JVM HeapMem
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMem
Processes	PeakThreadCount	-	Peak number of threads
	ThreadCount	-	Number of threads
	DaemonThreadCount	-	Number of backend threads
Process execution duration	Uptime	S	Process execution duration
Process start time	StartTime	S	Process start time
File descriptors	MaxFileDescriptorCount	-	Maximum number of file descriptors
	OpenFileDescriptorCount	-	Number of opened file descriptors

# **Trino Monitoring Metrics**

Last updated : 2023-12-27 14:50:38

# Trino - Overview

Title	Metric	Unit	Description
	Active	-	Number of active nodes
Nodes	Total	-	Total number of nodes
	Failed	-	Number of failed nodes
Query	RunningQueries	-	Total number of running queries
Query	QueuedQueries	-	Total number of waiting queries
	FailedQueries	count/min	Total number of failed queries
	AbandonedQueries	count/min	Total number of abandoned queries
Query frequency	CanceledQueries	count/min	Total number of canceled queries
	CompletedQueries	count/min	Total number of completed queries
	StartedQueries	count/min	Total number of started queries
Data volume input/output per	InputDataSizeOneMinute	GB/min	Data input rate
minute	OutputDataSizeOneMinute	GB/min	Data output rate

### Trino - worker

Title	Metric	Unit	Description
GC count	YGC	-	Young GC count



	FGC	-	Full GC count
GC time	FGCT	S	Full GC time
	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	S0	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
Memory zone	CCS	%	Percentage of used compressed class space memory
proportion	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	M %		Percentage of used Metaspace memory
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
JVM memory	MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
	MemHeapInitM	MB	Size of initial JVM HeapMem
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMem
Data input/output rate	InputDataSize.OneMinute.Rate	GB/min	Data input rate
Data mput/output late	OutputDataSize.OneMinute.Rate	GB/min	Data output rate
Worker threads	PeakThreadCount	-	Peak number of threads



	ThreadCount	-	Number of threads
	DaemonThreadCount	-	Number of backend threads
Process execution duration	Uptime	S	Process execution duration
File descriptors	MaxFileDescriptorCount	-	Maximum number of file descriptors
	OpenFileDescriptorCount	-	Number of opened file descriptors

### Trino - coordinator

Title	Metric	Unit	Description
GC count	YGC	-	Young GC count
GC count	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	SO	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
Memory zone	CCS	%	Percentage of used compressed class space memory
proportion	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
JVM memory	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
	MemHeapCommittedM	MB	Size of HeapMemory currently committed



			by JVM
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
	MemHeapInitM	MB	Size of initial JVM HeapMem
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMem
Worker threads	PeakThreadCount	-	Peak number of threads
	ThreadCount	-	Number of threads
	DaemonThreadCount	-	Number of backend threads
Process execution duration	Uptime	S	Process execution duration
Process start time	StartTime	S	Process start time
File descriptors	MaxFileDescriptorCount	-	Maximum number of file descriptors
	OpenFileDescriptorCount	-	Number of opened file descriptors

# **ClickHouse Monitoring Metrics**

Last updated : 2023-12-27 14:50:56

### **ClickHouse - metrics**

Title	Metric	Unit	Description
Natural connections	tcp	-	Number of TCP server connections
Network connections	http	-	Number of HTTP server connections
ZooKeeper event subscriptions	watches	-	Number of ZooKeeper event subscriptions
Ephemeral nodes held in ZooKeeper	ephemeralNode	-	Number of ephemeral nodes held in ZooKeeper
Active tooks in PookgroundPool	backgroundPoolTask	-	Number of active tasks in BackgroundProcessingPool
Active tasks in backgroundFool	backgroundSchedulePoolTask	-	Number of active tasks in BackgroundSchedulePool
Threads waiting for locks in Context	contextLockWait	-	Number of threads waiting for locks in Context
Throttled INSERT queries	delayedInserts	-	Number of throttled INSERT queries
Requests in fly to data sources of dictionaries of cache type	dictCacheRequests	-	Number of requests in fly to data sources of dictionaries of cache type
Pending files to process for asynchronous insertion into distributed tables	distributedSend	-	Number of pending files to process for asynchronous insertion into distributed tables
Threads	global	-	Number of threads in the global thread pool
	globalActive	-	Number of active threads in the global thread pool



	local	-	Number of threads in the local thread pool
	localActive	-	Number of active threads in the local thread pool
Replicas participating in leader election	leaderElection	-	Number of replicas participating in leader election
Replicated tables	leaderReplica	-	Number of replicated tables that are leaders
nepilcaled lables	readonlyReplica	-	Number of replicated tables in read-only status
	memoryTracking	GB	Total memory size allocated to queries being executed
Total allocated memory size	backgroundProcessingPool	GB	Total memory size allocated to the background processing pool
Total allocated memory size	backgroundSchedulePool	GB	Total memory size allocated to the background scheduling pool
	memoryTrackingForMerges	GB	Total memory size allocated to background merges
Merges being executed in the background	merge	-	Number of merges being executed in the background
Opened files	read	-	Number of opened readable files
Opened lifes	write	-	Number of opened writable files
Table changes	partMutation	-	Number of table changes
Query processing threads	queryThread	-	Number of query processing threads
Stopped or waiting queries	queryPreempted		Number of stopped or waiting queries



System calls	read		Number of read system calls
	write	-	Number of write system calls
	fetch	-	Number of blocks collected from replicas
Blocks	send	-	Number of blocks sent to replicas
	check	-	Number of blocks for consistency check
Server change	revision	-	Number of server changes
Version number	version	1	Version number
	waitingRead	-	Number of threads waiting for read/write locks to read from tables
Threads waiting for read/write locks	waitingWrite	-	Number of threads waiting for read/write locks to write to tables
	activeRead	-	Number of threads holding read locks on tables
	activeWrite	-	Number of threads holding write locks on tables
Rows in buffer tables	storageBufferRows	-	Number of rows in buffer tables
Bytes in buffer tables	storageBufferBytes	MB	Number of bytes in buffer tables

### **ClickHouse - events**

Title	Metric	Unit	Description
Queries	total	-	Total number of queries that may be executed
	select	-	Number of SELECT queries that may be executed



	insert	-	Number of INSERT queries that may be executed
Inserted rows	insertedRows	-	Number of rows inserted into all tables
Inserted bytes	insertedBytes	-	Number of bytes inserted into all tables
Total time waiting for system	read	μs	Total time waiting for read system calls
calls	write	μs	Total time waiting for write system calls
Opened files	fileOpen	-	Number of opened files
Reads/Writes from/to file descriptors	read	-	Number of reads from file descriptors
	write	-	Number of writes to file descriptors
Bytes read/written from/to file	read	MB	Number of bytes read from file descriptors
descriptors	write	MB	Number of bytes written to file descriptors
	realtime	μs	Total time processing threads
Total time processing threads	user	μs	Total time processing threads executing CPU instructions in user space
	system	μs	Total time processing threads executing CPU instructions in OS kernel space
Compiled regular expressions	regexpCreated	-	Number of compiled regular expressions

# ClickHouse - asynchronous metrics

Title	Metric	Unit	Description
Size of cache of marks of StorageMergeTree	markCacheBytes	MB	Size of cache of marks of StorageMergeTree
Cache files of marks of StorageMergeTree	markCacheFiles	-	Number of cache files of marks of StorageMergeTree
Maximum number of active blocks	maxPartCountForPartition	-	Number of largest active data blocks in partition
Databases	databaseCount	-	Number of databases
Data tables	tableCount	-	Number of data tables
Maximum replica latency	absolute	μs	Maximum replica queue latency



	relative	μs	Maximum difference of the absolute latency from other blocks
Size of the queue for the operations to be performed	replicasMaxQueueSize	-	Size of the queue for the operations to be performed

# **Druid Monitoring Metrics**

Last updated : 2023-12-27 14:51:14

### **Druid - Broker**

Metric Name	Unit	Description
YGC	-	Young GC count
FGC	-	Full GC count
FGCT	S	Full GC time
GCT	S	Garbage collection time
YGCT	S	Young GC time
S0	%	Percentage of used Survivor 0 memory
E	%	Percentage of used Eden memory
CCS	%	Percentage of memory used by compressed class space
S1	%	Percentage of used Survivor 1 memory
0	%	Percentage of used Old memory
М	%	Percentage of used Metaspace memory
MemHeapInitM	MB	Size of initial JVM HeapMem
MemNonHeapInitM	MB	Size of initial JVM NonHeapMem
MemHeapMaxM	MB	Size of HeapMemory configured by JVM
MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM
MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
ThreadsNew	-	Number of threads in NEW state
ThreadsRunnable	-	Number of threads in RUNNABLE state
ThreadsBlocked	-	Number of threads in BLOCKED state

ThreadsWaiting	-	Number of threads in WAITING state
ThreadsTimedWaiting	-	Number of threads in TIMED_WAITING state
ThreadsTerminated	-	Number of threads in TERMINATED state
LogFatal	-	Number of FATAL logs
LogError	-	Number of ERROR logs
LogWarn	-	Number of WARN logs
LogInfo	-	Number of INFO logs
jetty.numOpenConnections	-	Number of open Jetty connections
segment.scan.pending	-	Number of segments in queue waiting to be scanned
broker.query.count	-	number of total queries
broker.query.success.count	-	number of queries successfully processed
broker.query.failed.count	-	number of failed queries
broker.query.interrupted.count	-	number of queries interrupted due to cancellation or timeout
normal.count	-	Number of queries with delay shorter than 1 second
abnormal.count	-	Number of queries with delay longer than or equal to 1 second

## **Druid - Coordinator**

Metric Name	Unit	Description
segment.assigned.count	-	Number of segments loaded into Druid cluster
segment.moved.count	-	Number of segments moved in Druid cluster
segment.dropped.count	-	Number of segments deleted in Druid cluster due to expiration
segment.deleted.count	-	Number of segments deleted in Druid cluster due to rule settings
segment.unneeded.count	-	Number of segments deleted in Druid cluster for being marked as "unused"
segment.cost.raw	ms	Used in cost balancing. The raw cost of hosting segments.
segment.cost.normalization	ms	Used in cost balancing. The normalization of hosting segments.

segment.cost.normalized	ms	Used in cost balancing. The normalized cost of hosting segments.
segment.loadQueue.size	Bytes	Size in bytes of segments to load.
segment.loadQueue.failed	-	Number of segments that failed to load.
segment.loadQueue.count	-	Number of segments to load.
segment.dropQueue.count	-	Number of segments to drop.
segment.overshadowed.count	-	Number of overshadowed segments.
tier.historical.count	-	Number of available historical nodes in each tier.
tier.replication.factor	-	Configured maximum replication factor in each tier.
tier.required.capacity	Bytes	Total capacity in bytes required in each tier.
tier.total.capacity	Bytes	Total capacity in bytes available in each tier.
compact.task.count	-	Number of compaction tasks
YGC	-	Young GC count
FGC	-	Full GC count
FGCT	S	Full GC time
GCT	S	Garbage collection time
YGCT	S	Young GC time
SO	%	Percentage of used Survivor 0 memory
E	%	Percentage of used Eden memory
CCS	%	Percentage of memory used by compressed class space
S1	%	Percentage of used Survivor 1 memory
0	%	Percentage of used Old memory
Μ	%	Percentage of used Metaspace memory
MemHeapInitM	MB	Size of initial JVM HeapMem
MemNonHeapInitM	MB	Size of initial JVM NonHeapMem
MemHeapMaxM	MB	Size of HeapMemory configured by JVM



MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM
MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
ThreadsNew	-	Number of threads in NEW state
ThreadsRunnable	-	Number of threads in RUNNABLE state
ThreadsBlocked	-	Number of threads in BLOCKED state
ThreadsWaiting	-	Number of threads in WAITING state
ThreadsTimedWaiting	-	Number of threads in TIMED_WAITING state
ThreadsTerminated	-	Number of threads in TERMINATED state
LogFatal	-	Number of FATAL logs
LogError	-	Number of ERROR logs
LogWarn	-	Number of WARN logs
LogInfo	-	Number of INFO logs
segment.size	Bytes	Total size of used segments in a data source. Emitted only for data sources to which at least one used segment belongs.
segment.count	-	Number of used segments belonging to a data source. Emitted only for data sources to which at least one used segment belongs.
segment.unavailable.count	-	Number of segments (not including replicas) left to load until segments that should be loaded in the cluster are available for queries.
segment.underReplicated.count	-	Number of segments (including replicas) left to load until segments that should be loaded in the cluster are available for queries.
jetty.numOpenConnections	-	Number of open Jetty connections

## **Druid - Historical**

Metric Name	Unit	Description

🔗 Tencent Cloud

YGC	-	Young GC count
FGC	-	Full GC count
FGCT	S	Full GC time
GCT	S	Garbage collection time
YGCT	S	Young GC time
SO	%	Percentage of used Survivor 0 memory
E	%	Percentage of used Eden memory
CCS	%	Percentage of memory used by compressed class space
S1	%	Percentage of used Survivor 1 memory
0	%	Percentage of used Old memory
Μ	%	Percentage of used Metaspace memory
MemHeapInitM	MB	Size of initial JVM HeapMem
MemNonHeapInitM	MB	Size of initial JVM NonHeapMem
MemHeapMaxM	MB	Size of HeapMemory configured by JVM
MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM
MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
ThreadsNew	-	Number of threads in NEW state
ThreadsRunnable	-	Number of threads in RUNNABLE state
ThreadsBlocked	-	Number of threads in BLOCKED state
ThreadsWaiting	-	Number of threads in WAITING state
ThreadsTimedWaiting	-	Number of threads in TIMED_WAITING state
ThreadsTerminated	-	Number of threads in TERMINATED state
LogFatal	-	Number of FATAL logs

LogError	-	Number of ERROR logs
LogWarn	-	Number of WARN logs
LogInfo	-	Number of INFO logs
jetty.numOpenConnections	-	Number of open Jetty connections
segment.scan.pending	-	Number of segments in queue waiting to be scanned
segment.max	Bytes	Maximum byte limit available for segments
segment.pendingdelete	Bytes	On-disk size in bytes of segments that are waiting to be cleared out
historical.query.count	-	Total number of historical queries
historical.query.success.count	-	Total number of historical query successes
historical.query.failed.count	-	Total number of historical query failures
historical.query.interrupted.count	-	Total number of historical query interruptions
normal.count	-	Number of queries with delay shorter than 1 second
abnormal.count	-	Number of queries with delay longer than or equal to 1 second

## Druid - MiddleManager

Metric Name	Unit	Description
YGC	-	Young GC count
FGC	-	Full GC count
FGCT	S	Full GC time
GCT	S	Garbage collection time
YGCT	S	Young GC time
S0	%	Percentage of used Survivor 0 memory
E	%	Percentage of used Eden memory
CCS	%	Percentage of memory used by compressed class space
S1	%	Percentage of used Survivor 1 memory

🔗 Tencent Cloud

0	%	Percentage of used Old memory
Μ	%	Percentage of used Metaspace memory
MemHeapInitM	MB	Size of initial JVM HeapMem
MemNonHeapInitM	MB	Size of initial JVM NonHeapMem
MemHeapMaxM	MB	Size of HeapMemory configured by JVM
MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM
MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
ThreadsNew	-	Number of threads in NEW state
ThreadsRunnable	-	Number of threads in RUNNABLE state
ThreadsBlocked	-	Number of threads in BLOCKED state
ThreadsWaiting	-	Number of threads in WAITING state
ThreadsTimedWaiting	-	Number of threads in TIMED_WAITING state
ThreadsTerminated	-	Number of threads in TERMINATED state
LogFatal	-	Number of FATAL logs
LogError	-	Number of ERROR logs
LogWarn	-	Number of WARN logs
LogInfo	-	Number of INFO logs
jetty.numOpenConnections	-	Number of open Jetty connections

## Druid - Overlord

Metric Name	Unit	Description	
YGC	-	Young GC count	
FGC	-	Full GC count	
FGCT	S	Full GC time	

🔗 Tencent Cloud

GCT	S	Garbage collection time
YGCT	S	Young GC time
SO	%	Percentage of used Survivor 0 memory
E	%	Percentage of used Eden memory
CCS	%	Percentage of memory used by compressed class space
S1	%	Percentage of used Survivor 1 memory
0	%	Percentage of used Old memory
Μ	%	Percentage of used Metaspace memory
MemHeapInitM	MB	Size of initial JVM HeapMem
MemNonHeapInitM	MB	Size of initial JVM NonHeapMem
MemHeapMaxM	MB	Size of HeapMemory configured by JVM
MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM
MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
ThreadsNew	-	Number of threads in NEW state
ThreadsRunnable	-	Number of threads in RUNNABLE state
ThreadsBlocked	-	Number of threads in BLOCKED state
ThreadsWaiting	-	Number of threads in WAITING state
ThreadsTimedWaiting	-	Number of threads in TIMED_WAITING state
ThreadsTerminated	-	Number of threads in TERMINATED state
LogFatal	-	Number of FATAL logs
LogError	-	Number of ERROR logs
LogWarn	-	Number of WARN logs
LogInfo	-	Number of INFO logs



\_

Number of open Jetty connections

# Druid - Router

Metric Name	Unit	Description
YGC	-	Young GC count
FGC	-	Full GC count
FGCT	S	Full GC time
GCT	S	Garbage collection time
YGCT	S	Young GC time
SO	%	Percentage of used Survivor 0 memory
E	%	Percentage of used Eden memory
CCS	%	Percentage of memory used by compressed class space
S1	%	Percentage of used Survivor 1 memory
0	%	Percentage of used Old memory
Μ	%	Percentage of used Metaspace memory
MemHeapInitM	MB	Size of initial JVM HeapMem
MemNonHeapInitM	MB	Size of initial JVM NonHeapMem
MemHeapMaxM	MB	Size of HeapMemory configured by JVM
MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM
MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
ThreadsNew	-	Number of threads in NEW state
ThreadsRunnable	-	Number of threads in RUNNABLE state
ThreadsBlocked	-	Number of threads in BLOCKED state
ThreadsWaiting	-	Number of threads in WAITING state

ThreadsTimedWaiting	-	Number of threads in TIMED_WAITING state
ThreadsTerminated	-	Number of threads in TERMINATED state
LogFatal	-	Number of FATAL logs
LogError	-	Number of ERROR logs
LogWarn	-	Number of WARN logs
LogInfo	-	Number of INFO logs
jetty.numOpenConnections	-	Number of open Jetty connections

# Kudu Monitoring Metrics

Last updated : 2023-12-27 14:51:37

## Kudu - overview

Title	Metric	Unit	Description	
Tablets	TabletRunning	-	Total number of tablets currently running on all tablet servers	
Difference in the number of tablet replicas	ClusterReplicaSkew	-	Difference between the number of replicas on the tablet server hosting the most replicas and the number of replicas on the tablet server hosting the fewest replicas	
TServer threads	ThreadsRunning	-	Number of threads currently running on all tablet servers	
Master threads	ThreadsRunning	-	Number of threads currently running on all masters	
TServer logs	ErrorMessages	-	Number of ERROR-level log messages emitted in all processes	
Masterlage	ErrorMessages	-	Number of ERROR-level log messages emitted in all processes	
Master logs	WarningMessages	-	Number of WARNING-level log messages emitted in all processes	
Oversized write requests	OversizedWriteRequests	-	Number of oversized write requests to the system catalog tablet rejected by the master since start	

### Kudu - server

Title	Metric	Unit	Description
Block cache hit	BlockCacheHit	-	Number of block cache hits. When confirming the cache efficiency, use the value of this metric instead of that of cache_hits
	BlockCacheMiss	-	Number of block cache misses. When confirming the cache efficiency, use the



			value of this metric instead of that of cache_misses
Block cache utilization	BlockCacheUsage	bytes	Memory used by block cache
File cache hit	FileCacheHit	-	Number of file descriptor cache hits. When confirming the cache efficiency, use the value of this metric instead of that of cache_hits
	FileCacheMiss	-	Number of file descriptor cache misses. When confirming the cache efficiency, use the value of this metric instead of that of cache_misses
File cache utilization	FileCacheUsage	-	Number of entries file cache
	ActiveScanners	-	Number of currently active scanners
Scanner	ExpiredScanners	-	Number of scanners that have expired due to inactivity since service start
Block manager blocks	BlockUnderManagement	-	Number of currently managed data blocks
	BlockOpenReading	-	Number of data blocks currently opened for read
	BlockOpenWriting	-	Number of data blocks currently opened for write
Block manager bytes	BytesUnderManagement	bytes	Number of bytes of currently managed data blocks
Block manager	ContainersUnderManagement	-	Number of log block containers
containers	FullContainersUnderManagement	-	Number of full log block containers
Tablet leaders	NumRaftLeaders	-	Number of tablet replicas that are Raft leaders
Tablet sessions	OpenClientSessions	-	Number of currently opened tablet copy client sessions on this server
	OpemSourceSessions	-	Number of currently opened tablet copy source sessions on this server



Tablets	TabletBootstrapping	-	Number of currently bootstrapping tablets
	TabletFailed	-	Number of failed tablets
	TabletInitialized	-	Number of currently initialized tablets
	TabletNotInitialized	-	Number of currently uninitialized tablets
	TabletRunning	-	Number of currently running tablets/Number of currently running threads
	TabletShutdown	-	Number of currently shut down tablets
	TabletStopped	-	Number of currently stopped tablets
	TabletStopping	-	Number of currently stopping tablets
CPI I time	CpuStime	ms	Total system CPU time of process
Gr O time	CpuUtime	ms	Total user CPU time of process
Data path	DataDirsFailed	-	Number of data directories whose disks are currently in failed status
	DataDirsFull	-	Number of data directories whose disks are currently full
Thread	ThreadsRunning	-	Number of currently running threads
	InvoluntarySwitches	-	Total involuntary context switches
Context	VoluntarySwitches	-	Total voluntary context switches
Spinlock	SpinlockContentionTime	μs	Amount of time consumed by contention on internal spinlocks since server start
Log information	ErrorMessages	-	Number of ERROR-level log messages emitted by the application
	WarningMessages	-	Number of WARNING-level log messages emitted by the application
Operations in	TotalCount	-	Total number
queue	Min	-	Minimum number of tasks waiting in



			the queue
	Max	-	Maximum number of tasks waiting in the queue
	Mean	-	Average number of tasks waiting in the queue
	Percentile_99_9	-	99.9th percentile of the number of tasks waiting in the queue
	TotalCount	μs	Total number of operations
	Min	μs	Minimum run time
Operation execution duration	Max	μs	Maximum run time
	Mean	μs	Average run time
	Percentile_99_9	μs	99.9th percentile of the run time
	TotalCount	μs	Total number of operations
	Min	μs	Minimum wait time
Queuing wait time	Max	μs	Maximum wait time
	Mean	μs	Average wait time
	Percentile_99_9	μs	99.9th percentile of the wait time
Allocated bytes	AllocatedBytes	bytes	Number of bytes used by applications. This usually does not match the memory usage reported by the operating system because it does not include TCMalloc overhead or memory fragments
Hybrid clock error	HybridClockError	μs	Server clock maximum error; returns 2^64-1 when unable to read the base clock
Hybrid clock timestamp	HybridClockTimestamp	μs	Hybrid clock timestamp; returns 2^64-1 when unable to read the base clock
TCMalloc memory	HeapSize	bytes	Bytes of system memory reserved by TCMalloc
	CurrentThreadCacheBytes	bytes	A measure of some of the memory



			TCMalloc is using (for small objects)
	TotalThreadCacheBytes	bytes	A limit to how much memory TCMalloc dedicates for small objects
TCMalloc	FreeBytes	bytes	Number of bytes of free mapped pages in the page heap
PageHeap	UnMappedBytes	bytes	Number of bytes of free unmapped pages in the page heap
	ConnectionsAccepted	-	Number of incoming TCP connections made to the RPC server
RPC request	QueueOverflow	-	Number of RPCs dropped because the service queue was full
	TimesOutInQueue	-	Number of RPCs that timed out while waiting in the service queue and thus were not processed
	TotalCount	μs	Total number of operations
	Min	μs	Minimum processing time
RPC FetchData	Max	μs	Maximum processing time
	Mean	μs	Average processing time
	Percentile_99_9	μs	99.9th percentile of the processing time
	TotalCount	μs	Total number of operations
	Min	μs	Minimum processing time
RPC AlterSchema	Max	μs	Maximum processing time
	Mean	μs	Average processing time
	Percentile_99_9	μs	99.9th percentile of the processing time
RPC CreateTablet	TotalCount	μs	Total number of operations
	Min	μs	Minimum processing time
	Max	μs	Maximum processing time
	Mean	μs	Average processing time



	Percentile_99_9	μs	99.9th percentile of the processing time
	TotalCount	μs	Total number of operations
	Min	μs	Minimum processing time
RPC DeleteTablet	Max	μs	Maximum processing time
	Mean	μs	Average processing time
	Percentile_99_9	μs	99.9th percentile of the processing time
	TotalCount	μs	Total number of operations
	Min	μs	Minimum processing time
RPC Quiesce	Max	μs	Maximum processing time
	Mean	μs	Average processing time
	Percentile_99_9	μs	99.9th percentile of the processing time
	TotalCount	μs	Total number of operations
	Min	μs	Minimum processing time
RPC scan	Max	μs	Maximum processing time
	Mean	μs	Average processing time
	Percentile_99_9	μs	99.9th percentile of the processing time
	TotalCount	μs	Total number of operations
	Min	μs	Minimum processing time
RPC ScannerKeepAlive	Max	μs	Maximum processing time
	Mean	μs	Average processing time
	Percentile_99_9	μs	99.9th percentile of the processing time
RPC write	TotalCount	μs	Total number of operations
	Min	μs	Minimum processing time
	Max	μs	Maximum processing time
	Mean	μs	Average processing time



	Percentile_99_9	μs	99.9th percentile of the processing time
Write requests rejected due to queue overloading	QueueOverloadRejections	count	Number of write requests rejected due to queue overloading
Scan rate	ScannedFromDiskRate	bytes/s	Amount of data scanned per second
	ScannerReturnedRate	bytes/s	Amount of data returned per second
Scanner bytes	ScannedFromDisk	bytes	Total amount of data scanned from disk
	ScannerReturned	bytes	Total amount of returned data
Total row operations	RowsInserted	count	Number of rows inserted into the node
	RowsDeleted	count	Number of rows deleted from the node
	RowsUpserted	count	Number of rows upserted into the node
	RowsUpdated	count	Number of rows updated on the node
Row operation rate	RowsInsertedRate	count/s	Number of rows inserted into the node per second
	RowsDeletedRate	count/s	Number of rows deleted from the node per second
	RowsUpsertedRate	count/s	Number of rows upserted into the node per second
	RowsUpdatedRate	count/s	Number of rows updated on the node per second

## Kudu - master

Title	Metric	Unit	Description
Block cache hit	BlockCacheHit	-	Number of block cache hits. When confirming the cache efficiency, use the value of this metric instead of that of cache_hits
	BlockCacheMiss	-	Number of block cache misses. When confirming the cache efficiency, use the value of this metric instead of that of cache_misses


Block cache utilization	BlockCacheUsage	bytes	Memory used by block cache
File seeks kit	FileCacheHit	-	Number of file descriptor cache hits. When confirming the cache efficiency, use the value of this metric instead of that of cache_hits
File cache hit	FileCacheMiss	-	Number of file descriptor cache misses. When confirming the cache efficiency, use the value of this metric instead of that of cache_misses
File cache utilization	FileCacheUsage	-	Number of entries file cache
	BlockUnderManagement	-	Number of currently managed data blocks
Block manager blocks	BlockOpenReading	-	Number of data blocks currently opened for read
	BlockOpenWriting	-	Number of data blocks currently opened for write
Block manager bytes	BytesUnderManagement	bytes	Number of bytes of currently managed data blocks
Block manager	ContainersUnderManagement	-	Number of log block containers
containers	FullContainersUnderManagement	-	Number of full log block containers
CDLL time	CpuStime	ms	Total system CPU time of process
CFU lime	CpuUtime	ms	Total user CPU time of process
Thread	ThreadsRunning	-	Number of currently running threads
Data path	DataDirsFailed	-	Number of data directories whose disks are currently in failed status
	DataDirsFull	-	Number of data directories whose disks are currently full
Allocated bytes	AllocatedBytes	bytes	Number of bytes used by applications. This usually does not match the memory usage reported by the operating system



			because it does not include TCMalloc overhead or memory fragments	
	ErrorMessages	-	Number of ERROR-level log messages emitted by the application	
Log information	WarningMessages	-	Number of WARNING-level log messages emitted by the application	
Contaut	InvoluntarySwitches	-	Total involuntary context switches	
Context	VoluntarySwitches	-	Total voluntary context switches	
	TotalCount	-	Total number	
	Min	-	Minimum number of tasks waiting in the queue	
Operations in queue	Max	-	Maximum number of tasks waiting in the queue	
	Mean	-	Average number of tasks waiting in the queue	
	Percentile_99_9	-	99.9th percentile of the number of tasks waiting in the queue	
	TotalCount	μs	Total number of operations	
	Min	μs	Minimum wait time	
Queuing wait time	Max	μs	Maximum wait time	
	Mean	μs	Average wait time	
	Percentile_99_9	μs	99.9th percentile of the wait time	
	TotalCount	μs	Total number of operations	
	Min	μs	Minimum run time	
Operation execution duration	Max	μs	Maximum run time	
	Mean	μs	Average run time	
	Percentile_99_9	μs	99.9th percentile of the run time	
Spinlock	SpinlockContentionTime	μs	Amount of time consumed by contention on internal spinlocks since server start	



Oversized write requests	OversizedWriteRequests	-	Number of oversized write requests to the system catalog tablet rejected since start
Hybrid clock error	HybridClockError	μs	Server clock maximum error; returns 2^64-1 when unable to read the base clock
Hybrid clock timestamp	HybridClockTimestamp	μs	Hybrid clock timestamp; returns 2^64-1 when unable to read the base clock
Difference in the number of tablet replicas	ClusterReplicaSkew	-	Difference between the number of replicas on the tablet server hosting the most replicas and the number of replicas on the tablet server hosting the fewest replicas
Tablet leaders	NumRaftLeaders	-	Number of tablet replicas that are Raft leaders
Tablet sessions	OpemSourceSessions	-	Number of currently opened tablet copy source sessions on this server
TCMalloc memory	HeapSize	bytes	Bytes of system memory reserved by TCMalloc
	CurrentThreadCacheBytes	bytes	A measure of some of the memory TCMalloc is using (for small objects)
	TotalThreadCacheBytes	bytes	A limit to how much memory TCMalloc dedicates for small objects
TCMalloc page	FreeBytes	bytes	Number of bytes of free mapped pages in the page heap
heap	UnMappedBytes	bytes	Number of bytes of free unmapped pages in the page heap
RPC request	ConnectionsAccepted	-	Number of incoming TCP connections made to the RPC server
	QueueOverflow	-	Number of RPCs dropped because the service queue was full
	TimesOutInQueue	-	Number of RPCs that timed out while waiting in the service queue and thus were not processed



RPC	TotalCount	μs	Total number of operations
NUILEAUEIEIECIIOII	Min	μs	Minimum processing time
	Max	μs	Maximum processing time
	Mean	μs	Average processing time
	Percentile_99_9	μs	99.9th percentile of the processing time
	TotalCount	μs	Total number of operations
	Min	μs	Minimum processing time
RPC ConnectToMaster	Max	μs	Maximum processing time
	Mean	μs	Average processing time
	Percentile_99_9	μs	99.9th percentile of the processing time
	TotalCount	μs	Total number of operations
	Min	μs	Minimum processing time
RPC Ping	Max	μs	Maximum processing time
	Mean	μs	Average processing time
	Percentile_99_9	μs	99.9th percentile of the processing time
	TotalCount	μs	Total number of operations
	Min	μs	Minimum processing time
RPC TSHeartbeat	Max	μs	Maximum processing time
	Mean	μs	Average processing time
	Percentile_99_9	μs	99.9th percentile of the processing time
	TotalCount	μs	Total number of operations
	Min	μs	Minimum processing time
RPC FetchData	Max	μs	Maximum processing time
	Mean	μs	Average processing time
	Percentile_99_9	μs	99.9th percentile of the processing time



# **Alluxio Monitoring Metrics**

Last updated : 2023-12-27 14:51:59

## Alluxio - cluster

Title	Metric	Unit	Description
Total data reads and	BytesReadAlluxio	Bytes	Total number of bytes read from Alluxio storage reported by all workers
	BytesReadUfsAll	Bytes	Total number of bytes read by all workers from all Alluxio UFSes
writes	BytesWrittenAlluxio	Bytes	Total number of bytes written to Alluxio storage in all workers
	BytesWrittenUfsAll	Bytes	Total number of bytes written to all Alluxio UFSes by all workers
	BytesReadAlluxioThroughput	Bytes	Data read throughput from Alluxio storage by all workers
Data read and write	BytesReadUfsThroughput	Bytes	Read throughput from all Alluxio UFSes by all workers
throughput	BytesWrittenAlluxioThroughput	Bytes	Data write throughput to Alluxio storage by all workers
	BytesWrittenUfsThroughput	Bytes	Write throughput to all Alluxio UFSes by all workers
Worker capacity on tiers	CapacityFree	Bytes	Total number of free bytes on all tiers of all workers
	CapacityTotal	Bytes	Total capacity on all tiers of all workers
	CapacityUsed	Bytes	Total number of used bytes on all tiers of all workers
Total workers	Workers	-	Total number of active workers inside the cluster

#### Alluxio - master

Title	Metric	Unit	Description



CompleteFile	CompleteFileOps	-	Total number of CompleteFile operations
operation	FilesCompleted	-	Total number of successful CompleteFile operations
CreateDirectory	CreateDirectoryOps	-	Total number of CreateDirectory operations
operation	DirectoriesCreated	-	Total number of successful CreateDirectory operations
	CreateFileOps	-	Total number of CreateFile operations
CreateFile operation	FilesCreated	-	Total number of successful CreateFile operations
Delete operation	DeletePathOps	-	Total number of Delete operations
Delete operation	PathsDeleted	-	Total number of successful Delete operations
	FreeFileOps	-	Total number of FreeFile operations
FreeFile operation	FilesFreed	-	Total number of successful FreeFile operations
	GetFileBlockInfoOps	-	Total number of GetFileBlockInfo operations
operation	FileBlockInfosGot	-	Total number of successful GetFileBlockInfo operations
	GetFileInfoOps	-	Total number of GetFileInfo operations
GetFileInfo operation	FileInfosGot	-	Total number of successful GetFileInfo operations
GotNowBlock	GetNewBlockOps	-	Total number of GetNewBlock operations
operation	NewBlocksGot	-	Total number of successful GetNewBlock operations
Mount operation	MountOps	-	Total number of Mount operations
Mount operation	PathsMounted	-	Total number of successful Mount operations
	UnmountOps	-	Total number of Unmount operations
Unmount operation	PathsUnmounted	-	Total number of successful Unmount operations
Rename operation	RenamePathOps	-	Total number of Rename operations



	PathsRenamed	-	Total number of successful Rename operations	
SetAcl operation	SetAclOps	-	Total number of SetAcl operations	
SetAttribute operation	SetAttributeOps	-	Total number of SetAttribute operations	
Total filos	FilesPersisted	-	Total number of successfully persisted files	
Total mes	FilesPinned	-	Total number of currently pinned files	
Total file directories	TotalPaths	-	Total number of files and directories in the Alluxio namespace	
GC count	YGC	-	Young GC count	
Ge count	FGC	-	Full GC count	
	FGCT	S	Full GC time	
GC time	GCT	S	Garbage collection time	
	YGCT	S	Young GC time	
	S0	%	Percentage of used Survivor 0 memory	
	E	%	Percentage of used Eden memory	
Memory zone	CCS	%	Percentage of used compressed class space memory	
ρισροπιση	S1	%	Percentage of used Survivor 1 memory	
	0	%	Percentage of used Old memory	
	Μ	%	Percentage of used Metaspace memory	
JVM memory	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM	
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM	
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM	
	MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM	
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM	



MemHeapInitM	MB	Size of initial JVM HeapMem
MemNonHeapInitM	MB	Size of initial JVM NonHeapMem

### Alluxio - worker

Title	Metric	Unit	Description
Async cache request	AsyncCacheDuplicateRequests	-	Total number of duplicate async cache requests received by worker
	AsyncCacheRequests	-	Total number of async cache requests received by worker
	AsyncCacheFailedBlocks	-	Total number of failed async cache blocks in worker
	AsyncCacheRemoteBlocks	-	Total number of blocks that need to be asynchronously cached from remote sources
blocks	AsyncCacheSucceededBlocks	-	Total number of blocks that were asynchronously cached successfully in the worker
	AsyncCacheUfsBlocks	-	Total number of blocks that need to be asynchronously cached from local sources
	BlocksAccessed	-	Total number of blocks accessed in worker
	BlocksCached	-	Total number of blocks used for caching data in worker
	BlocksCancelled	-	Total number of temporary blocks canceled in worker
Blocks	BlocksDeleted	-	Total number of blocks deleted from worker by external requests
	BlocksEvicted	-	Total number of blocks evicted from worker
	BlocksLost	-	Total number of blocks lost in worker
	BlocksPromoted	-	Total number of blocks moved to a new tier in worker
Worker	CapacityFree	Bytes	Total free bytes on all tiers of worker
capacity on	CapacityTotal	Bytes	Total capacity on all tiers of worker



tiers	CapacityUsed	Bytes	Total used bytes on all tiers of worker
GC count	YGC	-	Young GC count
	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	SO	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
Memory zone	CCS	%	Percentage of used compressed class space memory
рюронног	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
JVM memory	MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
	MemHeapInitM	MB	Size of initial JVM HeapMem
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMem

# **PrestoSQL** Monitoring Metrics

Last updated : 2023-12-27 14:52:15

#### Note:

Currently, only PrestoSQL v322, PrestoSQL v350, and later support PrestoSQL metrics.

#### **PrestoSQL** - overview

Title	Metric	Unit	Description
	Active	-	Number of active nodes
Nodes	Total	-	Total number of nodes
	Failed	-	Number of failed nodes
Quory	RunningQueries	-	Total number of running queries
Query	QueuedQueries	-	Total number of waiting queries
	FailedQueries	count/min	Total number of failed queries
	AbandonedQueries	count/min	Total number of abandoned queries
Query frequency	CanceledQueries	count/min	Total number of canceled queries
	CompletedQueries	count/min	Total number of completed queries
	StartedQueries	count/min	Total number of started queries
Data volume input/output per	InputDataSizeOneMinute	GB/min	Data input rate
minute	OutputDataSizeOneMinute	GB/min	Data output rate

#### **PrestoSQL** - worker

Title	Metric	Unit	Description



GC count	YGC	-	Young GC count
	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	SO	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
Memory zone	CCS	%	Percentage of used compressed class space memory
proportion	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
JVM memory	MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
	MemHeapInitM	MB	Size of initial JVM HeapMem
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMem
Data input/output	InputDataSize.OneMinute.Rate	GB/min	Data input rate
rate	OutputDataSize.OneMinute.Rate	GB/min	Data output rate



Processes	PeakThreadCount	-	Peak number of threads
	ThreadCount	-	Total number of threads
	DaemonThreadCount	-	Number of daemon threads
Process execution duration	Uptime	S	Process execution duration
Process start time	StartTime	S	Process start time
Filo doscriptors	MaxFileDescriptorCount	-	Maximum number of file descriptors
The descriptors	OpenFileDescriptorCount	-	Number of opened file descriptors

## PrestoSQL - coordinator

Title	Metric	Unit	Description
CC count	YGC	-	Young GC count
Ge count	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	SO	%	Percentage of used Survivor 0 memory
Memory zone	E	%	Percentage of used Eden memory
	CCS	%	Percentage of used compressed class space memory
proportion	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
JVM memory	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM



	MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
	MemHeapInitM	MB	Size of initial JVM HeapMem
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMem
Processes	PeakThreadCount	-	Peak number of threads
	ThreadCount	-	Number of threads
	DaemonThreadCount	-	Number of backend threads
Process execution duration	Uptime	S	Process execution duration
Process start time	StartTime	S	Process start time
File de cariateur	MaxFileDescriptorCount	-	Maximum number of file descriptors
	OpenFileDescriptorCount	-	Number of opened file descriptors

# Impala Monitoring Metrics

Last updated : 2023-12-27 14:52:35

#### Note:

Currently, only Impala v3.4.0 and later support Impala metrics.

### Impala - catalog

Title	Metric	Unit	Description
Resident set size	RSS	bytes	Resident set size
	MemHeapInitM	MB	Peak size of initial JVM HeapMemory
	MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
JVM memory	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMemory
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
Interval between heartbeats from the daemon process to StateStore	Last	S	Latest interval between heartbeats from the daemon process to StateStore
	Max	S	Maximum interval between heartbeats from the daemon process to StateStore
	Mean	S	Average interval between

🔗 Tencent Cloud

			heartbeats from the daemon process to StateStore
	Min	S	Minimum interval between heartbeats from the daemon process to StateStore
	Stddev	S	Standard deviation in the interval between heartbeats from the daemon process to StateStore
	Used	bytes	Number of bytes used by the program
TCMalloc memory	PageheapFreeBytes	bytes	Number of bytes of free mapped pages in the page heap
	PageheapUnmappedBytes	bytes	Number of bytes of free unmapped pages in the page heap
	PhysicalBytesReserved	bytes	Amount of physical memory used by the process
	TotalBytesReserved	bytes	Number of bytes of system memory reserved by TCMalloc
Active connections	Thrift_Server_Connections_Used	-	Number of active connections
Process execution duration	Uptime	S	Process execution duration
File descriptors	MaxFileDescriptorCount	-	Maximum number of file descriptors
	OpenFileDescriptorCount	-	Number of opened file descriptors
Threads	ThreadCount	-	Total number of threads



	DaemonThreadCount	-	Number of daemon threads
CPU utilization	SystemCpuLoad	-	System CPU Utilization

## Impala - StateStore

Title	Metric	Unit	Description
Resident set size	RSS	bytes	Resident set size
TCMalloc memory	Used	bytes	Number of bytes used by the program
	PageheapFreeBytes	bytes	Number of bytes of free mapped pages in the page heap
	PageheapUnmappedBytes	bytes	Number of bytes of free unmapped pages in the page heap
	PhysicalBytesReserved	bytes	Amount of physical memory used by the process
	TotalBytesReserved	bytes	Number of bytes of system memory reserved by TCMalloc
Connections	Used	-	Number of active connections
Running threads	Count	-	Number of running threads
StateStore subscribers	Count	-	Number of StateStore subscribers

# Impala - daemon

Title	Metric	Unit	Description
JVM memory	MemHeapInitM	MB	Peak size of initial JVM HeapMemory
	MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
	MemHeapUsedM	MB	Size of HeapMemory

🔗 Tencent Cloud

			currently used by JVM
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMemory
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	Used	bytes	Number of bytes used by the program
TCMalloc memory	PageheapFreeBytes	bytes	Number of bytes of free mapped pages in the page heap
	PageheapUnmappedBytes	bytes	Number of bytes of free unmapped pages in the page heap
	PhysicalBytesReserved	bytes	Amount of physical memory used by the process
	TotalBytesReserved	bytes	Number of bytes of system memory reserved by TCMalloc
Threads	ThreadCount	-	Total number of threads
meaus	DaemonThreadCount	-	Number of daemon threads
Process execution duration	Uptime	S	Process execution duration
File descriptors	MaxFileDescriptorCount	-	Maximum number of file descriptors
	OpenFileDescriptorCount	-	Number of opened file descriptors

CPU utilization	SystemCpuLoad	-	System CPU Utilization
Beeswax API client connections	Use	-	Number of active Beeswax API connections
	Conn_In_Use	-	Number of active Beeswax API connections to this Impala daemon
	TotalConns	-	Total number of active Beeswax API connections to this Impala daemon
	ConnSetupQueueSize	-	Number of Beeswax API connections that this Impala daemon has received and are waiting to establish
	Use	-	Number of active HS2 API connections
	Conn_In_Use	-	Number of active HS2 API connections
HS2 API client connections	TotalConns	-	Total number of HS2 API connections this Impala daemon has established during the lifecycle
	ConnSetupQueueSize	-	Number of HS2 API connections that this Impala daemon has received and are waiting to establish
	RunningThreads	-	Number of running threads
Thread manager	TotalCreatedThreads	-	Number of threads created during the lifecycle



Memory manager limit	Limit	Bytes	Memory beyond the memory limit (default value: -1)	
Amount of memory beyond the memory limit (default value: -1)	OverLimit	Bytes	Number of threads created during the lifecycle	
	P20			
	P50			
HS2 API client wait time for	P70		HS2 API client wait	
connection establishment	P90	us	establishment	
	P95			
	P99.9			
	P20			
	P50		Beeswax API client wait time for service thread establishment	
Beeswax API client wait time for	P70			
service thread establishment	P90	us		
	P95			
	P99.9			
Timed-out Beeswax API connections	TimeOutCnncRequests	-	Number of timed-out Beeswax API connections	
Time spent parsing requests from request pool (milliseconds)	Total	ms	Time spent parsing requests from request pool (milliseconds)	
Cache misses in external data source cache class	Misses	-	Number of cache misses in external data source cache class	
Impala backend server connection requests that timed out waiting for setup	ConnSetupQueueSize	-	Number of Impala backend server connection requests	



			that timed out waiting for setup
Impala backend connection requests that timed out waiting for setup	TimeOutCnncRequests	-	Number of Impala backend connection requests that timed out waiting for setup
Total Impala backend client connections to this Impala daemon	TotalConnections	-	Total number of Impala backend client connections to this Impala daemon
	P20	_	
	P50		
	P70		
Time spent by Impala backend client	P90	110	Time spent by Impala backend client waiting for connection establishment
waiting for connection establishment	P95	us	
	P99.9	-	
	Count		
	Sum		
	P20		
	P50	-	
	P70		
Time spent by Impala backend client	P90		Time spent by Impala
waiting for service thread	P95	us	for service thread
	P99.9	-	
	Count		
	Sum		
HS2 API client wait time for service	P20	us	HS2 API client wait
thread establishment	P50		time for service thread establishment



	P70		
	P90		
	P95		
	P99.9		
	Count		
	Sum		
	P20		
	P50		
	P70		
HS2 HTTP API client wait time for	P90		HS2 HTTP API client
service thread	P95	us	thread
	P99.9		
	Count		
	Sum		
DataStreamService: Rejected service queue overflows	RpcsQueueOverflow	-	DataStreamService: Number of rejected service queue overflows
ControlStreamService: Rejected service queue overflows	RpcsQueueOverflow	-	ControlStreamService: Number of rejected service queue overflows
DataStreamService: Used bytes	PeakUsageBytes	Bytes	Peak number of bytes used by Memtracker DataStreamService
	CurrentUsageBytes	Bytes	Current number of bytes used by Memtracker DataStreamService
ControlService: Used bytes	PeakUsageBytes	Bytes	Peak number of bytes used by Memtracker



			ControlService	
	CurrentUsageBytes	Bytes	Current number of bytes used by Memtracker ControlService	
Resident set size (RSS) for this process	RSS	Bytes	Resident set size (RSS) for this process	
Total backends registered in StateStore	Total	-	Total number of backends registered in StateStore	
	P20			
	P50	-		
	P70	-	Our make and delay	
	P90			
Query release delay	P95		Query release delay	
	P99.9			
	Count			
	Sum			
Opened HDFS filed	NumFilesOpenForInsert	-	Number of opened HDFS files	
Scan range read during the lifecycle of the process	ScanRangesTotal	-	Scan range read during the lifecycle of the process	
Opened Beeswax sessions	NumOpenBeeswaxSessions	-	Number of opened Beeswax sessions	
Total query fragments processed during the lifecycle of the process	NumFragments	-	Total number of query fragments processed during the lifecycle of the process	
Total scan ranges read during the lifecycle of the process without volume metadata	ScanRangesNumMissingVolumId	-	Total number of scan ranges read during the lifecycle of the process	



			without volume metadata
Hedged read attempts	HedgedReadOps	-	Number of hedged read attempts
Total queries processed during the lifecycle of the process	NumQueries	-	Total number of queries processed during the lifecycle of the process
Total rows supporting caching HS2 FETCH_FIRST	ResultSetCacheTotalNumRows	-	Total number of rows supporting caching HS2 FETCH_FIRST
Total queries registered on this Impala server	NumQueriesRegistered	-	Total number of queries registered on this Impala server
Total backend queries	NumQueriesExecuted	-	Total number of backend queries
Sessions terminated due to inactivity	NumSessionsExpired	-	Number of sessions terminated due to inactivity
Queries terminated due to inactivity	NumQueriesExpired	-	Number of queries terminated due to inactivity
Opened HS2 sessions	NumOpenHS2Sessions	-	Number of opened HS2 sessions
Tables in catalog	NumTables	-	Number of tables in catalog
Databases in catalog	NumDatabases	-	Number of databases in catalog
Bytes written to the disk by the I/O manager	BytesWritten	-	Number of bytes written to the disk by the I/O manager
Files opened by the I/O manager	NumOpenFiles	-	Number of files opened by the I/O manager
Used HDFS file handles	NumFileHandlesOutstanding	Bytes	Number of used



			HDFS file handles
Read local bytes	LocalBytesRead	Bytes	Number of local bytes read by the I/O manager

# **Ranger Monitoring Metrics**

Last updated : 2023-12-27 14:52:55

### Ranger - admin

Title	Metric	Unit	Description
GC count	YGC	-	Young GC count
Ge count	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	SO	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
Memory zone proportion	CCS	%	Percentage of used compressed class space memory
	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
	MemHeapCommittedM	MB	Size of HeapMemory committed by JVM
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
JVM memory	MemHeapInitM	MB	Size of initial JVM HeapMem
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMem



CPU utilization	ProcessCpuLoad	%	CPU utilization
File bandlag	MaxFileDescriptorCount	-	Maximum number of file descriptors
File nancies	OpenFileDescriptorCount	-	Number of opened file descriptors
Cumulative CPU usage time	ProcessCpuTime	ms	Cumulative CPU usage time
Process execution duration	Uptime	S	Process execution duration
	ThreadCount	-	Number of threads
Threads	PeakThreadCount	-	Peak number of threads
	DaemonThreadCount	-	Number of backend threads

# Ranger - UserSync

Title	Metric	Unit	Description
CC count	YGC	-	Young GC count
de count	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	SO	%	Percentage of used Survivor 0 memory
Memory zone proportion	E	%	Percentage of used Eden memory
	CCS	%	Percentage of used compressed class space memory
	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
JVM memory	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
	MemHeapCommittedM	MB	Size of HeapMemory committed by JVM



	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
	MemHeapInitM	MB	Size of initial JVM HeapMem
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMem
CPU utilization	ProcessCpuLoad	%	CPU utilization
File bandles	MaxFileDescriptorCount	-	Maximum number of file descriptors
The nanules	OpenFileDescriptorCount	-	Number of opened file descriptors
Cumulative CPU usage time	ProcessCpuTime	ms	Cumulative CPU usage time
Process execution duration	Uptime	S	Process execution duration
	ThreadCount	-	Number of threads
Threads	PeakThreadCount	-	Peak number of threads
	DaemonThreadCount	-	Number of backend threads

# **COSRanger Monitoring Metrics**

Last updated : 2023-12-27 14:53:49

#### Note:

Metrics such as verification failure statistics, authentication failure statistics, and authentication success statistics currently don't have specific data, which will be available in the future.

### **COSRanger - COSRangerServer**

Title	Metric	Unit	Description
GC count	YGC	-	Young GC count
	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	SO	%	Percentage of used Survivor 0 memory
Memory zone proportion	E	%	Percentage of used Eden memory
	CCS	%	Percentage of used compressed class space memory
	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
JVM memory	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
	MemHeapCommittedM	MB	Size of HeapMemory committed by JVM

	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
	MemHeapInitM	MB	Size of initial JVM HeapMem
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMem
CPU utilization	ProcessCpuLoad	%	CPU utilization
File handles	MaxFileDescriptorCount	-	Maximum number of file descriptors
	OpenFileDescriptorCount	-	Number of opened file descriptors
CPU usage time	ProcessCpuTime	ms	Cumulative CPU usage time
Process execution duration	Uptime	S	Process execution duration
	ThreadCount	-	Number of threads
Worker threads	PeakThreadCount	-	Peak number of threads
	DaemonThreadCount	-	Number of backend threads
-	Leader	-	Whether it is the COSRanger master node
	PermissionAllowCnt	count	Total number of permission allows
Check statistics	AuthDenyCnt	count	Total number of authentication failures
	PermissionDenyCnt	count	Total number of permission denies
Authentication success statistics	Qps	count	Number of queries per second
	Total_5m	count	Total number of requests per five minutes



	Total_1m	count	Total number of requests per minute
	Qps_5m	count	Average number of requests per five minutes
	Qps_1m	count	Average number of requests per minute
Authentication failure statistics	Qps	count	Number of queries per second
	Total_5m	count	Total number of requests per five minutes
	Total_1m	count	Total number of requests per minute
	Qps_5m	count	Average number of requests per five minutes
	Qps_1m	count	Average number of requests per minute
Permission deny statistics	Qps	count	Number of queries per second
	Total_5m	count	Total number of requests per five minutes
	Total_1m	count	Total number of requests per minute
	Qps_5m	count	Average number of requests per five minutes
	Qps_1m	count	Average number of requests per minute
Permission allow statistics	Qps	count	Number of queries per second
	Total_5m	count	Total number of requests per five minutes
	Total_1m	count	Total number of requests per minute
	Qps_5m	count	Average number of requests per five minutes
	Qps_1m	count	Average number of requests



			per minute
accessStat_DELETE operation statistics	Qps	count	Number of queries per second
	Total_5m	count	Total number of requests per five minutes
	Total_1m	count	Total number of requests per minute
	Qps_5m	count	Average number of requests per five minutes
	Qps_1m	count	Average number of requests per minute
	Qps	count	Number of queries per second
accessStat_LIST operation statistics	Total_5m	count	Total number of requests per five minutes
	Total_1m	count	Total number of requests per minute
	Qps_5m	count	Average number of requests per five minutes
	Qps_1m	count	Average number of requests per minute
accessStat_READ operation statistics	Qps	count	Number of queries per second
	Total_5m	count	Total number of requests per five minutes
	Total_1m	count	Total number of requests per minute
	Qps_5m	count	Average number of requests per five minutes
	Qps_1m	count	Average number of requests per minute
accessStat_WRITE operation statistics	Qps	count	Number of queries per second
	Total_5m	count	Total number of requests per five minutes



	Total_1m	count	Total number of requests per minute
	Qps_5m	count	Average number of requests per five minutes
	Qps_1m	count	Average number of requests per minute
	Qps	count	Number of queries per second
rpc_getRangerAuthPolicy call count statistics	Total_5m	count	Total number of requests per five minutes
	Total_1m	count	Total number of requests per minute
	Qps_5m	count	Average number of requests per five minutes
	Qps_1m	count	Average number of requests per minute
rpc_checkPermission call count statistics	Qps	count	Number of queries per second
	Total_5m	count	Total number of requests per five minutes
	Total_1m	count	Total number of requests per minute
	Qps_5m	count	Average number of requests per five minutes
	Qps_1m	count	Average number of requests per minute
rpc_getDelegationToken call count statistics	Qps	count	Number of queries per second
	Total_5m	count	Total number of requests per five minutes
	Total_1m	count	Total number of requests per minute
	Qps_5m	count	Average number of requests per five minutes
	Qps_1m	count	Average number of requests



			per minute
rpc_renewDelegationToken call count statistics	Qps	count	Number of queries per second
	Total_5m	count	Total number of requests per five minutes
	Total_1m	count	Total number of requests per minute
	Qps_5m	count	Average number of requests per five minutes
	Qps_1m	count	Average number of requests per minute
	Qps	count	Number of queries per second
rpc_cancelDelegationToken call count statistics	Total_5m	count	Total number of requests per five minutes
	Total_1m	count	Total number of requests per minute
	Qps_5m	count	Average number of requests per five minutes
	Qps_1m	count	Average number of requests per minute
rpc_getSTS call count statistics	Qps	count	Number of queries per second
	Total_5m	count	Total number of requests per five minutes
	Total_1m	count	Total number of requests per minute
	Qps_5m	count	Average number of requests per five minutes
	Qps_1m	count	Average number of requests per minute
cosRpc_getSTS call duration	Cost_Avg	μs	Average duration in the current second
	Cost_Avg_1m	μs	Average duration per minute



	Cost_Avg_5m	μs	Average duration per five minutes
	Cost_Max	μs	Maximum duration in the current second
	Cost_Max_1m	μs	Maximum duration per minute
	Cost_Max_5m	μs	Maximum duration per five minutes
	Cost_Min	μs	Minimum duration in the current second
	Cost_Min_1m	μs	Minimum duration per minute
	Cost_Min_5m	μs	Minimum duration per five minutes
cosRpc_renewDelegationToken call duration	Cost_Avg	μs	Average duration in the current second
	Cost_Avg_1m	μs	Average duration per minute
	Cost_Avg_5m	μs	Average duration per five minutes
	Cost_Max	μs	Maximum duration in the current second
	Cost_Max_1m	μs	Maximum duration per minute
	Cost_Max_5m	μs	Maximum duration per five minutes
	Cost_Min	μs	Minimum duration in the current second
	Cost_Min_1m	μs	Minimum duration per minute
	Cost_Min_5m	μs	Minimum duration per five minutes
cosRpc_cancelDelegationToken call duration	Cost_Avg	μs	Average duration in the current second
	Cost_Avg_1m	μs	Average duration per minute
	Cost_Avg_5m	μs	Average duration per five



			minutes
	Cost_Max	μs	Maximum duration in the current second
	Cost_Max_1m	μs	Maximum duration per minute
	Cost_Max_5m	μs	Maximum duration per five minutes
	Cost_Min	μs	Minimum duration in the current second
	Cost_Min_1m	μs	Minimum duration per minute
	Cost_Min_5m	μs	Minimum duration per five minutes
	Cost_Avg	μs	Average duration in the current second
	Cost_Avg_1m	μs	Average duration per minute
	Cost_Avg_5m	μs	Average duration per five minutes
	Cost_Max	μs	Maximum duration in the current second
cosRpc_getDelegationToken call duration	Cost_Max_1m	μs	Maximum duration per minute
	Cost_Max_5m	μs	Maximum duration per five minutes
	Cost_Min	μs	Minimum duration in the current second
	Cost_Min_1m	μs	Minimum duration per minute
	Cost_Min_5m	μs	Minimum duration per five minutes
cosRpc_checkPermission call duration	Cost_Avg	μs	Average duration in the current second
	Cost_Avg_1m	μs	Average duration per minute
	Cost_Avg_5m	μs	Average duration per five minutes


	Cost_Max	μs	Maximum duration in the current second
	Cost_Max_1m	μs	Maximum duration per minute
	Cost_Max_5m	μs	Maximum duration per five minutes
	Cost_Min	μs	Minimum duration in the current second
	Cost_Min_1m	μs	Minimum duration per minute
	Cost_Min_5m	μs	Minimum duration per five minutes
	Cost_Avg	μs	Average duration in the current second
	Cost_Avg_1m	μs	Average duration per minute
	Cost_Avg_5m	μs	Average duration per five minutes
	Cost_Max	μs	Maximum duration in the current second
cosRpc_getRangerAuthPolicy call duration	Cost_Max_1m	μs	Maximum duration per minute
	Cost_Max_5m	μs	Maximum duration per five minutes
	Cost_Min	μs	Minimum duration in the current second
	Cost_Min_1m	μs	Minimum duration per minute
	Cost_Min_5m	μs	Minimum duration per five minutes

# **Doris Monitoring Metrics**

Last updated : 2023-12-27 14:54:08

#### Doris - FE

Title	Metric	Unit	Description
	FeNodeNum	count	Total number of FE nodes
Node information	BeAliveNum	count	Number of alive BE nodes
	BkDeadNum	count	Number of dead broker nodes
Connections	Num	count	Number of FE node JVM connections
JVM threads	Total	count	Total number of threads in the JVM of the FE node, including daemon threads and non- daemon threads
	Peak	count	Peak number of threads in the JVM of the FE node
GC count	YoungGC	count	FE node JVM Young GC count
GC count	OldGC	count	FE node JVM Old GC count
GC time	YoungGC	S	FE node JVM Young GC time
	OldGC	S	FE node JVM Old GC time
	Quantile75	ms	75th percentile of the FE query latency
EE quory latonoy	Quantile95	ms	95th percentile of the FE query latency
T L query latency	Quantile99	ms	99th percentile of the FE query latency
	Quantile999	ms	99.9th percentile of the FE query latency
Maximum compaction score of tablet	MAX	score	Maximum compaction score of FE tablet
Scheduled tablets	ScheduledTablet	count	Number of scheduled tablets in FE
Request response	QPS	count	Number of queries per second
	RPS	count	Number of requests that can be processed per



			second
Query failure rate	ErrRate	%	Query error rate
	SqlModelHitQuery	count	Number of queries that hit the SQL model
	PartitionModelHitQuery	count	Number of queries that hit the partition model
	SqlModelQuery	count	Number of queries in the SQL cache model
Cache query	PartitionModelQuery	count	Number of queries in the partition cache model
	CachePartitionHit	count	Number of partitions that hit the cache in the query
	CachePartitionScan	count	Number of all partitions scanned in the query
Poutino load rows	TotalRows	count	Number of FE routine load rows
Houtine load rows	ErrorRows	count	Number of FE routine load error rows
	Reject	count	Number of rejected transactions on FE
Transaction status	Begin	count	Number of transactions that have begun on FE
statistics	Success	count	Number of successful transactions on FE
	Failed	count	Number of failed transactions on FE
Images	Write	count	Number of image writes on FE
	Push	count	Number of image pushes on FE
ALTED ich statistics	RollupRunning	count	Number of running alter jobs of ROLLUP type
ALTEN JOD Statistics	SchemaChangeRunning	count	Number of running alter jobs of SCHEMA_CHANGE type
UNKNOWN load job statistics	UNKNOWN	count	Number of load jobs of UNKNOWN type in UNKNOWN status
	PENDING	count	Number of load jobs of UNKNOWN type in PENDING status
	ETL	count	Number of load jobs of UNKNOWN type in ETL status
	LOADING	count	Number of load jobs of UNKNOWN type in



			LOADING status
	COMMITTED	count	Number of load jobs of UNKNOWN type in COMMITTED status
	FINISHED	count	Number of load jobs of UNKNOWN type in FINISHED status
	CANCELLED	count	Number of load jobs of UNKNOWN type in CANCELLED status
	UNKNOWN	count	Number of load jobs of SPARK type in UNKNOWN status
	PENDING	count	Number of load jobs of SPARK type in PENDING status
	ETL	count	Number of load jobs of SPARK type in ETL status
SPARK load job statistics	LOADING	count	Number of load jobs of SPARK type in LOADING status
	COMMITTED	count	Number of load jobs of SPARK type in COMMITTED status
	FINISHED	count	Number of load jobs of SPARK type in FINISHED status
	CANCELLED	count	Number of load jobs of SPARK type in CANCELLED status
DELETE load job statistics	UNKNOWN	count	Number of load jobs of DELETE type in UNKNOWN status
	PENDING	count	Number of load jobs of DELETE type in PENDING status
	ETL	count	Number of load jobs of DELETE type in ETL status
	LOADING	count	Number of load jobs of DELETE type in LOADING status
	COMMITTED	count	Number of load jobs of DELETE type in COMMITTED status
	FINISHED	count	Number of load jobs of DELETE type in



		FINISHED status	
	CANCELLED	count	Number of load jobs of DELETE type in CANCELLED status
	UNKNOWN	count	Number of load jobs of INSERT type in UNKNOWN status
	PENDING	count	Number of load jobs of INSERT type in PENDING status
	ETL	count	Number of load jobs of INSERT type in ETL status
INSERT load job statistics	LOADING	count	Number of load jobs of INSERT type in LOADING status
	COMMITTED	count	Number of load jobs of INSERT type in COMMITTED status
	FINISHED	count	Number of load jobs of INSERT type in FINISHED status
	CANCELLED	count	Number of load jobs of INSERT type in CANCELLED status
	UNKNOWN	count	Number of load jobs of BROKER type in UNKNOWN status
	PENDING	count	Number of load jobs of BROKER type in PENDING status
	ETL	count	Number of load jobs of BROKER type in ETL status
BROKER load job statistics	LOADING	count	Number of load jobs of BROKER type in LOADING status
	COMMITTED	count	Number of load jobs of BROKER type in COMMITTED status
	FINISHED	count	Number of load jobs of BROKER type in FINISHED status
	CANCELLED	count	Number of load jobs of BROKER type in CANCELLED status
MINI load job	UNKNOWN	count	Number of load jobs of MINI type in



statistics			UNKNOWN status
	PENDING	count	Number of load jobs of MINI type in PENDING status
	ETL	count	Number of load jobs of MINI type in ETL status
	LOADING	count	Number of load jobs of MINI type in LOADING status
	COMMITTED	count	Number of load jobs of MINI type in COMMITTED status
	FINISHED	count	Number of load jobs of MINI type in FINISHED status
	CANCELLED	count	Number of load jobs of MINI type in CANCELLED status
HADOOP load job statistics	UNKNOWN	count	Number of load jobs of HADOOP type in UNKNOWN status
	PENDING	count	Number of load jobs of HADOOP type in PENDING status
	ETL	count	Number of load jobs of HADOOP type in ETL status
	LOADING	count	Number of load jobs of HADOOP type in LOADING status
	COMMITTED	count	Number of load jobs of HADOOP type in COMMITTED status
	FINISHED	count	Number of load jobs of HADOOP type in FINISHED status
	CANCELLED	count	Number of load jobs of HADOOP type in CANCELLED status

#### Doris - BE

Title	Metric	Unit	Description
Thrift utilization	Broker	count	Number of thrifts used by broker
	Backend	count	Number of thrifts used by BE



	Extdatasource	count	Number of thrifts used by extdatasource
	Frontend	count	Number of thrifts used by FE
	RequestsTotal	count	Number of streaming load requests
Streaming load statistics	CurrentProcessing	count	Number of existing streaming load processes
	PipeCount	count	Number of streaming load pipes
Streaming load time	Duration	ms	Streaming load duration
Streaming load data volume	LoadTotal	bytes	Size of data imported by streaming load
	PlanFragment	count	Number of plan fragments
Fragment statistics	Endpoint	count	Number of DataStreams
	RequestsTotal	count	Number of fragment requests
Fragment request duration	Duration	μs	Fragment request duration
PE momory	Total	bytes	Size of BE memory pool
DE memory	Allocated	bytes	Size of allocated BE memory
Maximum compaction	CumulativeMax	score	Maximum cumulative compaction score of tablet
	BaseMax	score	Maximum base compaction score of tablet
Data handled by	Cumulative	bytes	Amount of data handled by cumulative compaction
compaction	Base	bytes	Amount of data handled by base compaction
Delta data handled by	Cumulative	bytes	Amount of delta data handled by cumulative compaction
compaction	Base	bytes	Amount of delta data handled by base compaction
MemPools used by compaction	CurrentConsumption	count	Sum of MemPools used by compaction (all compaction threads)



Process file handles	Used	count	Number of file handles used by BE process
	SoftLimit	count	Soft limit on file handles for BE process
	HardLimit	count	Hard limit on file handles for BE process
Running threads in process	NUM	count	Number of threads running in BE process
Engine request statistics	FailedBaseCompaction	count	Number of failed engine requests of base_compaction type
	FailedCultCompt	count	Number of failed engine requests of cumulative_compaction type
	TotalBaseCompaction	count	Total number of engine requests of base_compaction type
	TotalCultCompt	count	Total number of engine requests of cumulative_compaction type

#### Doris - BK

Title	Metric	Unit	Description
	YGC	-	Young GC count
GC count	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
Memory zone	SO	%	Percentage of used Survivor 0 memory
proportion	E	%	Percentage of used Eden memory
	CCS	%	Percentage of used compressed class space memory
	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory



	Μ	%	Percentage of used Metaspace memory
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
	MemHeapCommittedM	MB	Size of HeapMemory committed by JVM
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
	MemHeapInitM	MB	Size of initial JVM HeapMem
JVM memory	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMem
CPU utilization	ProcessCpuLoad	%	CPU utilization
File handlag	MaxFileDescriptorCount	-	Maximum number of file descriptors
File fiancies	OpenFileDescriptorCount	-	Number of opened file descriptors
CPU usage time	ProcessCpuTime	ms	Cumulative CPU usage time
Process execution duration	Uptime	S	Process execution duration
Worker threads	ThreadCount	-	Number of threads
	PeakThreadCount	-	Peak number of threads
	DaemonThreadCount	-	Number of backend threads

# Kylin Monitoring Metrics

Last updated : 2023-12-27 14:54:25

### Kylin - Kylin

Title	Metric	Unit	Description
GC count	YGC	-	Young GC count
	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	SO	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
Memory zone proportion	CCS	%	Percentage of used compressed class space memory
	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
	MemHeapCommittedM	MB	Size of HeapMemory committed by JVM
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
JVM memory	MemHeapInitM	MB	Size of initial JVM HeapMemory
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory committed by JVM
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMemory



File descriptors	OpenFileDescriptorCount	-	Number of opened file descriptors
	MaxFileDescriptorCount	-	Maximum number of file descriptors
CPU utilization	ProcessCpuLoad	%	Process CPU utilization
	SystemCpuLoad	%	System CPU Utilization
Worker threads	DaemonThreadCount	-	Number of daemon threads
	ThreadCount	-	Total number of threads
Cumulative CPU usage time	ProcessCpuTime	ms	Cumulative CPU usage time
Process execution duration	Uptime	S	Process execution duration

# **Zeppelin Monitoring Metrics**

Last updated : 2023-12-27 14:54:44

### Zeppelin - Zeppelin

Title	Metric	Unit	Description	
GC count	YGC	-	Young GC count	
	FGC	-	Full GC count	
	FGCT	S	Full GC time	
GC time	GCT	S	Garbage collection time	
	YGCT	S	Young GC time	
	S0	%	Percentage of used Survivor 0 memory	
	E	%	Percentage of used Eden memory	
Momony zono proportion	CCS	%	Percentage of used compressed class space memory	
Memory zone proportion	S1	%	Percentage of used Survivor 1 memory	
	0	%	Percentage of used Old memory	
	М	%	Percentage of used Metaspace memory	

# **Oozie Monitoring Metrics**

Last updated : 2023-12-27 14:55:01

#### Oozie - Oozie

Title	Metric	Unit	Description	
GC count	YGC	-	Young GC count	
	FGC	-	Full GC count	
	FGCT	S	Full GC time	
GC time	GCT	S	Garbage collection time	
	YGCT	S	Young GC time	
	S0	%	Percentage of used Survivor 0 memory	
	E	%	Percentage of used Eden memory	
Momony zono proportion	CCS	%	Percentage of used compressed class space memory	
Memory zone proportion	S1	%	Percentage of used Survivor 1 memory	
	0	%	Percentage of used Old memory	
	М	%	Percentage of used Metaspace memory	

# **Storm Monitoring Metrics**

Last updated : 2023-12-27 14:55:18

#### Storm - Nimbus

Title	Metric	Unit	Description
CC count	YGC	-	Young GC count
Ge count	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	SO	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
Memory zone proportion	CCS	%	Percentage of used compressed class space memory
	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
	MemHeapCommittedM	MB	Size of HeapMemory committed by JVM
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
JVM memory	MemHeapInitM	MB	Size of initial JVM HeapMemory
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory committed by JVM
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMemory



File descriptors	OpenFileDescriptorCount	-	Number of opened file descriptors
	MaxFileDescriptorCount	-	Maximum number of file descriptors
CPU utilization	ProcessCpuLoad	%	Process CPU utilization
Worker threads	DaemonThreadCount	-	Number of daemon threads
	PeakThreadCount	-	Peak number of threads
	ThreadCount	-	Total number of threads
Cumulative CPU usage time	ProcessCpuTime	ms	Cumulative CPU usage time
Process execution duration	Uptime	S	Process execution duration

### Storm - supervisor

Title	Metric	Unit	Description
	YGC	-	Young GC count
Ge count	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	SO	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
Memory zone proportion	CCS	%	Percentage of used compressed class space memory
	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
JVM memory	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
	MemHeapCommittedM	MB	Size of HeapMemory committed by JVM



	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
	MemHeapInitM	MB	Size of initial JVM HeapMemory
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory committed by JVM
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMemory
	OpenFileDescriptorCount	-	Number of opened file descriptors
File descriptors	MaxFileDescriptorCount	-	Maximum number of file descriptors
CPU utilization	ProcessCpuLoad	%	Process CPU utilization
	DaemonThreadCount	-	Number of daemon threads
Worker threads	PeakThreadCount	-	Peak number of threads
	ThreadCount	-	Total number of threads
Cumulative CPU usage time	ProcessCpuTime	ms	Cumulative CPU usage time
Process execution duration	Uptime	S	Process execution duration

# Livy Monitoring Metrics

Last updated : 2023-12-27 14:55:36

### Livy - LivyServer

Title	Metric	Unit	Description
GC count	YGC	-	Young GC count
Ge count	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	SO	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
Memory zone proportion	CCS	%	Percentage of used compressed class space memory
	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
	MemHeapCommittedM	MB	Size of HeapMemory committed by JVM
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
JVM memory	MemHeapInitM	MB	Size of initial JVM HeapMemory
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory committed by JVM
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMemory



File descriptors	OpenFileDescriptorCount	-	Number of opened file descriptors
	MaxFileDescriptorCount	-	Maximum number of file descriptors
CPU utilization	ProcessCpuLoad	%	Process CPU utilization
Worker threads	DaemonThreadCount	-	Number of daemon threads
	PeakThreadCount	-	Peak number of threads
	ThreadCount	-	Total number of threads
Cumulative CPU usage time	ProcessCpuTime	ms	Cumulative CPU usage time
Process execution duration	Uptime	S	Process execution duration

# Kyuubi Monitoring Metrics

Last updated : 2023-12-27 14:55:56

### Kyuubi - KyuubiServer

Title	Metric	Unit	Description
GC count	YGC	-	Young GC count
	FGC	-	Full GC count
	FGCT	S	Full GC time
GC time	GCT	S	Garbage collection time
	YGCT	S	Young GC time
	SO	%	Percentage of used Survivor 0 memory
	E	%	Percentage of used Eden memory
Memory zone proportion	CCS	%	Percentage of used compressed class space memory
	S1	%	Percentage of used Survivor 1 memory
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
	MemHeapCommittedM	MB	Size of HeapMemory committed by JVM
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
JVM memory	MemHeapInitM	MB	Size of initial JVM HeapMemory
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory committed by JVM
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMemory



File descriptors	OpenFileDescriptorCount	-	Number of opened file descriptors
	MaxFileDescriptorCount	-	Maximum number of file descriptors
CPU utilization	ProcessCpuLoad	%	Process CPU utilization
Worker threads	DaemonThreadCount	-	Number of daemon threads
	PeakThreadCount	-	Peak number of threads
	ThreadCount	-	Total number of threads
Cumulative CPU usage time	ProcessCpuTime	ms	Cumulative CPU usage time
Process execution duration	Uptime	S	Process execution duration

# StarRocks Monitoring Metrics

Last updated : 2023-12-27 15:00:50

### StarRocks - BE

Title	Metric	Unit	Description
	Cumulative	rowsets	Amount of delta data handled by cumulative compaction
Delta data handled by compaction	Base	rowsets	Amount of delta data handled by base compaction
	Update	rowsets	Amount of delta data handled by update compaction
	Cumulative	bytes	Amount of data handled by cumulative compaction
Data handled by compaction	Base	bytes	Amount of data handled by base compaction
	Update	bytes	Amount of data handled by update compaction
Maximum compaction	CumulativeMax	score	Maximum cumulative compaction score of tablet
score of tablet	BaseMax	score	Maximum base compaction score of tablet
Engine request failure statistics (1)	base_compaction	count	Number of failed engine requests of base_compaction type
	clone	count	Number of failed engine requests of clone type
	create_rollup	count	Number of failed engine requests of create_rollup type
	create_tablet	count	Number of failed engine requests of create_tablet type
	cumulative_compaction	count	Number of failed engine requests of



			cumulative_compaction type
	delete	count	Number of failed engine requests of delete type
	finish_task	count	Number of failed engine requests of finish_task type
	publish	count	Number of failed engine requests of publish type
	report_all_tablets	count	Number of failed engine requests of report_all_tablets type
Engine request failure	report_disk	count	Number of failed engine requests of report_disk type
statistics (2)	report_tablet	count	Number of failed engine requests of report_tablet type
	report_task	count	Number of failed engine requests of report_task type
	schema_change	count	Number of failed engine requests of schema_change type
Engine request statistics (1)	base_compaction	count	Number of failed engine requests of base_compaction type
	clone	count	Number of failed engine requests of clone type
	create_rollup	count	Number of failed engine requests of create_rollup type
	create_tablet	count	Number of failed engine requests of create_tablet type
	cumulative_compaction	count	Number of failed engine requests of cumulative_compaction type
	delete	count	Number of failed engine requests of delete type
	drop_tablet	count	Number of failed engine requests of drop_tablet type
	finish_task	count	Number of failed engine requests of



			finish_task type
	publish	count	Number of failed engine requests of publish type
	report_all_tablets	count	Number of failed engine requests of report_all_tablets type
	report_disk	count	Number of failed engine requests of report_disk type
Engine request statistics (2)	report_tablet	count	Number of failed engine requests of report_tablet type
	report_task	count	Number of failed engine requests of report_task type
	schema_change	count	Number of failed engine requests of schema_change type
	storage_migrate	count	Number of failed engine requests of storage_migrate type
Fragment statistics	PlanFragment	count	Number of plan fragments
	Endpoint	count	Number of DataStreams
Fragment request duration	Duration	μs	Fragment request duration
	begin	count	Number of transaction requests of begin type
Transaction request	commit	count	Number of transaction requests of commit type
statistics	exec	count	Number of transaction requests of exec type
	rollback	count	Number of transaction requests of rollback type
Data imported by streaming load	LoadTotal	bytes	Amount of data imported by streaming load
Streaming load statistics	CurrentProcessing	count	Number of existing streaming load processes
	PipeCount	count	Number of streaming load pipes



Streaming load duration	Duration	ms	Streaming load duration
	Total	bytes	Size of BE memory pool
DE memory	Allocated	bytes	Size of allocated BE memory
Process file handles	Used	count	Number of file handles used by BE process
	SoftLimit	count	Soft limit on file handles for BE process
	HardLimit	count	Hard limit on file handles for BE process
Running threads in process	Thread	count	Number of threads running in BE process
	Broker	count	Number of thrifts used by broker
Used thrifts	Backend	count	Number of thrifts used by BE
	Frontend	count	Number of thrifts used by FE
Tablet write statistics	Writer	count	Statistics of tablet writes on BE

### StarRocks - FE

Title	Metric	Unit	Description
ALTER job statistics	RollupRunning	count	Number of running alter jobs of rollup type
	SchemaChangeRunning	count	Number of running alter jobs of schema_change type
Image statistics	Write	count	Number of image writes on FE
inage statistics	Push	count	Number of image pushes on FE
Scheduled tablets	ScheduledTablet	count	Number of scheduled tablets on FE
Transaction status	Reject	count	Number of rejected transactions on FE
statistics	Begin	count	Number of started transactions on FE
	Success	count	Number of successful transactions on FE



	Failed	count	Number of failed transactions on FE	
	max	bytes	Maximum amount of heap memory	
Heap memory in JVM	committed	bytes	Amount of committed heap memory	
	used	bytes	Amount of used heap memory	
Non-heap memory in JVM	committed	bytes	Amount of committed non-heap memory	
	used	bytes	Amount of used non-heap memory	
	used	bytes	Amount of used old memory	
Old memory in JVM	peak_used	bytes	Maximum amount of used old memory	
	max	bytes	Maximum amount of old memory	
	used	bytes	Amount of used young memory	
Young memory in JVM	peak_used	bytes	Maximum amount of used young memory	
	max	bytes	Maximum amount of young memory	
Routine load queue size	report queue	count	Size of FE report queue	
Routino load rows	TotalRows	count	Number of FE routine load rows	
Noutine load rows	ErrorRows	count	Number of FE routine load error rows	
Routine load size	Receive	bytes	Size of FE routine load	
Maximum compaction score of tablet	MAX	score	Maximum compaction score of FE tablet	
EditLog write latency	Quantile75	ms	75th percentile of the FE EditLog write latency	
	Quantile95	ms	95th percentile of the FE EditLog write latency	
	Quantile98	ms	98th percentile of the FE EditLog write latency	
	Quantile99	ms	99th percentile of the FE EditLog write latency	



	Quantile999	ms	99.9th percentile of the FE EditLog write latency
GC count	YoungGC	count	FE node JVM Young GC count
	OldGC	count	FE node JVM Old GC count
GC time	YoungGC	Sec	FE node JVM Young GC time
do une	OldGC	Sec	FE node JVM Old GC time
IVM threads	Total	count	Total number of threads in the JVM of the FE node
o vivi tili edds	Peak	count	Peak number of threads in the JVM of the FE node
	UNKNOWN	count	Number of load jobs of BROKER type in UNKNOWN status
	PENDING	count	Number of load jobs of BROKER type in PENDING status
BROKER load job statistics	ETL	count	Number of load jobs of BROKER type in ETL status
	LOADING	count	Number of load jobs of BROKER type in LOADING status
	COMMITTED	count	Number of load jobs of BROKER type in COMMITTED status
	FINISHED	count	Number of load jobs of BROKER type in FINISHED status
	CANCELLED	count	Number of load jobs of BROKER type in CANCELLED status
DELETE load job statistics	UNKNOWN	count	Number of load jobs of DELETE type in UNKNOWN status
	PENDING	count	Number of load jobs of DELETE type in PENDING status
	ETL	count	Number of load jobs of DELETE type in ETL status
	LOADING	count	Number of load jobs of DELETE type in LOADING status



	COMMITTED	count	Number of load jobs of DELETE type in COMMITTED status	
	FINISHED	count	Number of load jobs of DELETE type in FINISHED status	
	CANCELLED	count	Number of load jobs of DELETE type in CANCELLED status	
	UNKNOWN	count	Number of load jobs of HADOOP type in UNKNOWN status	
	PENDING	count	Number of load jobs of HADOOP type in PENDING status	
	ETL	count	Number of load jobs of HADOOP type in ETL status	
HADOOP load job statistics	LOADING	count	Number of load jobs of HADOOP type in LOADING status	
	COMMITTED	count	Number of load jobs of HADOOP type in COMMITTED status	
	FINISHED	count	Number of load jobs of HADOOP type in FINISHED status	
	CANCELLED	count	Number of load jobs of HADOOP type in CANCELLED status	
INSERT load job statistics	UNKNOWN	count	Number of load jobs of INSERT type in UNKNOWN status	
	PENDING	count	Number of load jobs of INSERT type in PENDING status	
	ETL	count	Number of load jobs of INSERT type in ETL status	
	LOADING	count	Number of load jobs of INSERT type in LOADING status	
	COMMITTED	count	Number of load jobs of INSERT type in COMMITTED status	
	FINISHED	count	Number of load jobs of INSERT type in FINISHED status	



	CANCELLED	count	Number of load jobs of INSERT type in CANCELLED status
	NEED_SCHEDULE	count	Number of routine load jobs in NEED_SCHEDULE status
	RUNNING	count	Number of routine load jobs in RUNNING status
Routine load job statistics	PAUSED	count	Number of routine load jobs in PAUSED status
	STOPPED	count	Number of routine load jobs in STOPPED status
	CANCELLED	count	Number of routine load jobs in CANCELLED status
SPARK load job statistics	UNKNOWN	count	Number of load jobs of SPARK type in UNKNOWN status
	PENDING	count	Number of load jobs of SPARK type in PENDING status
	ETL	count	Number of load jobs of SPARK type in ETL status
	LOADING	count	Number of load jobs of SPARK type in LOADING status
	COMMITTED	count	Number of load jobs of SPARK type in COMMITTED status
	FINISHED	count	Number of load jobs of SPARK type in FINISHED status
	CANCELLED	count	Number of load jobs of SPARK type in CANCELLED status
FE MASTER	FE Master	count	Whether it is the FE master. Valid values: 1: Master; 0: Follower.
Node information	FeNodeNum	count	Total number of FE nodes
	BeNodeNum	count	Total number of BE nodes
	BeAliveNum	count	Number of live BE nodes
	BeDecommissionedNum	count	Number of live BE nodes



	BkDeadNum	count	Number of dead broker nodes	
	QPS	count/s	Number of queries per second	
Request response	RPS	count/s	Number of requests that can be processed per second	
	total	count	Total number of FE queries	
	err	count	Total number of FE query errors	
FE query statistics	timeout	count	Number of timed-out FE queries	
	success	count	Total number of successful FE queries	
	slow	count	Total number of slow FE queries	
Query failure rate	ErrRate	%	Query error rate	
FE query latency	Quantile75	ms	75th percentile of the FE query latency	
	Quantile95	ms	95th percentile of the FE query latency	
	Quantile99	ms	99th percentile of the FE query latency	
	Quantile999	ms	99.9th percentile of the FE query latency	
Connections	Num	count	Number of FE node connections	

### StarRocks - Broker

Title	Metric	Unit	Description	
CPU utilization	ProcessCpuLoad	%	Process CPU utilization	
CPU usage time	ProcessCpuTime	ms	Cumulative CPU usage time	
GC count	YGC	-	Young GC count	
	FGC	-	Full GC count	
GC time	GCT	S	Garbage collection time	
	FGCT	S	Full GC time	



	YGCT	S	Young GC time	
	0	%	Percentage of used Old memory	
	Μ	%	Percentage of used Metaspace memory	
Memory zone	CCS	%	Percentage of used compressed class space memory	
ρισμοιτιστι	SO	%	Percentage of used Survivor 0 memory	
	S1	%	Percentage of used Survivor 1 memory	
	E	%	Percentage of used Eden memory	
	MemHeapInitM	MB	Size of initial JVM HeapMemory	
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMemory	
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM	
JVM memory	MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM	
	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM	
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM	
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM	
Filo bandlos	OpenFileDescriptorCount	count	Number of opened file descriptors	
The nancies	MaxFileDescriptorCount	count	Maximum number of file descriptors	
Process execution duration	Uptime	S	Process execution duration	
Worker threads	PeakThreadCount	count	Peak number of threads	
	ThreadCount	count	Total number of threads	
	DaemonThreadCount	count	Number of Daemon threads	

# Kafka Monitoring Metrics

Last updated : 2023-12-27 15:01:37

Title	Metric	Unit	Description
CPU utilization	ProcessCpuLoad	%	Process CPU utilization
CPU usage time	ProcessCpuTime	ms	Cumulative CPU usage time
CC count	YGC	-	Young GC count
Ge count	FGC	-	Full GC count
	GCT	S	Garbage collection time
GC time	FGCT	S	Full GC time
	YGCT	S	Young GC time
	0	%	Percentage of used Old memory
	Μ	%	Percentage of used Metaspace memory
	CCS	%	Percentage of used compressed class space memory
Memory zone proportion	SO	%	Percentage of used Survivor 0 memory
	S1	%	Percentage of used Survivor 1 memory
	E	%	Percentage of used Eden memory
JVM memory	MemHeapInitM	MB	Size of initial JVM HeapMemory
	MemNonHeapInitM	MB	Size of initial JVM NonHeapMemory
	MemHeapMaxM	MB	Size of HeapMemory configured by JVM
	MemHeapCommittedM	MB	Size of HeapMemory currently committed by JVM



	MemHeapUsedM	MB	Size of HeapMemory currently used by JVM
	MemNonHeapCommittedM	MB	Size of NonHeapMemory currently committed by JVM
	MemNonHeapUsedM	MB	Size of NonHeapMemory currently used by JVM
File handles	OpenFileDescriptorCount	count/s	Number of opened file descriptors
	MaxFileDescriptorCount	count/s	Maximum number of file descriptors
Process execution duration	Uptime	S	Process execution duration
	PeakThreadCount	count/s	Peak number of threads
Worker threads	ThreadCount	count/s	Total number of threads
	DaemonThreadCount	count/s	Number of Daemon threads
Production traffic of Broker	OneMinuteRate	bytes/s	Rate of message production by the Broker per minute
Consumption traffic of Broker	OneMinuteRate	bytes/s	Rate of message consumption by the Broker per minute
Rejected consumption traffic	OneMinuteRate	bytes/s	Rate of Topic request rejection per minute
Failed Fetch requests	OneMinuteRate	count/s	Rate of Fetch request failures per minute
Failed Produce requests	OneMinuteRate	count/s	Rate of Produce request failures per minute
Produced messages	OneMinuteRate	count/s	Rate of message production per minute
Read traffic from other Brokers	OneMinuteRate	bytes/s	Amount of traffic read from other Brokers per minute
Read traffic to other Brokers	OneMinuteRate	bytes/s	Amount of traffic read to other Brokers per minute
Fetch requests	OneMinuteRate	count/s	Rate of Fetch requests per



			minute
Produce requests	OneMinuteRate	count/s	Rate of Produce requests per minute
ControllerBroker	IsControllerBroker	-	The metric value is 1 on the Broker where Controller is located and is 0 on other Brokers.
LeaderElection rate	OneMinuteRate	count/s	Rate of LeaderElections per minute
LeaderElection latencies	99thPercentile	ms	The 99th percentile of LeaderElection latencies
	999thPercentile		The 99.9th percentile of LeaderElection latencies
	Mean		Mean value of LeaderElection latencies
UncleanLeaderElections rate	OneMinuteRate	count/s	Rate of UncleanLeaderElections per minute
GlobalPartition count	GlobalPartitionCount	count/s	Number of global partitions observed by the controller
OfflinePartitions count	OfflinePartitionCount	count/s	Number of offline partitions observed by the controller
GlobalTopic count	GlobalTopicCount	count/s	Number of global topics observed by the controller
Offline log directory count	OfflineLogDirectory	count/s	Number of offline log directories
LogFlush rate	OneMinuteRate	calls/s	Rate of message log flush per minute
LogFlush latencies	99thPercentile	ms	The 99th percentile of LogFlush latencies
	999thPercentile		The 99.9th percentile of LogFlush latencies
	Mean		Mean value of LogFlush latencies



Average network processor idleness rate	NetworkProcessorAvgIdlePercent	-	Average idleness rate of threads in a network thread pool
ISR expansion rate	OneMinuteRate	count/s	In-sync replica (ISR) expansion rate per minute
ISR shrinkage rate	OneMinuteRate	count/s	ISR shrinkage rate per minute
Poplicas count	LeaderReplicaCount	count/s	Number of offline replicas
Replicas count	OfflineReplicaCount	count/s	Number of leader replicas
	PartitionCount	count	Number of partitions
Partitions count	UnderMinIsrPartitionCount		Number of partitions under the minimum ISR count
	UnderReplicatedPartitions		Number of UnderReplicatedPartitions
FetchConsumer request latencies	99thPercentile	ms	The 99.9th percentile of FetchConsumer request latencies
	999thPercentile		The 99.9th percentile of FetchConsumer request latencies
	Mean		Mean value of FetchConsumer request latencies
FetchFollower request latencies	99thPercentile	ms	The 99.9th percentile of FetchFollower request latencies
	999thPercentile		The 99.9th percentile of FetchFollower request latencies
	Mean		Mean value of FetchFollower request latencies
Produce request latency	99thPercentile		The 99.9th percentile of Produce request latencies
	999thPercentile	ms	The 99.9th percentile of Produce request latencies
	Mean		Mean value of Produce request latencies



Size of the request queue	RequestQueueSize	size	Size of the request queue
Purgatory size	Fetch	size	Number of requests waiting in fetch purgatory
	Produce	-	Number of requests waiting in producer purgatory
Average request handler idleness rate	OneMinuteRate	-	Idleness rate of request handlers per minute
ZooKeeper request latencies	99thPercentile	ms	The 99th percentile of ZooKeeper request latencies
	999thPercentile		The 99.9th percentile of ZooKeeper request latencies
	Mean		Mean value of ZooKeeper request latencies

# Alarm Configurations

Last updated : 2023-12-27 15:01:54

### Overview

Elastic MapReduce (EMR) has been connected to Cloud Monitor (CM). You can configure alarm policies for EMR nodes and service monitoring metrics in the CM console.

#### Note:

EMR has been connected to CM's default alarms. CM will automatically create a default alarm policy. For information about metrics/events or alarm rules supported by the EMR default policy, see Default Alarm Policy.

You can manually create an alarm policy and set it as the default alarm policy. After the default policy is set, newly purchased instances will be automatically associated with the default policy without requiring manual addition.

### Directions

1. Log in to the Cloud Monitor console and click **Alarm Configuration** > Alarm Policy on the left sidebar.

2. On the Alarm Policy page, click Create.

3. In the pop-up window, configure the basic information, alarm policy, and notification template as instructed below.

Configuration Type	Configuration Item	Description		
Basic information	Policy name	Custom policy name		
	Remarks	Custom policy remarks		
	Monitoring type	Supported cloud service monitoring type		
	Policy type	Select the desired policy type for monitoring Tencent Cloud services.		
	Project	You can filter to find alarm policies under a project in the alarm policy list.		
Alarm rule configuration	Alarm object	If you select an instance ID, the alarm policy will be associated with the selected instance. If you select an instance group, the alarm policy will be associated with the selected instance group. If you select all objects, the alarm policy will be associated all instances that the current account has permission on.		


	Trigger condition	Manual configuration (metric alarm)	Trigger condition: consists of metric, comparison, threshold, statistical period, and the number of consecutive periods. You can expand the trigger condition to view the metric trend, and based on which, set a proper threshold.				
		Manual configuration (event alarm)	Create an event alarm policy to get notifications in case of service resources or underlying infrastructure exceptions.				
		Select template	Select a configured template in the drop-down list. For more information, see Configuring Trigger Condition Template.				
	Configure alarm notification (optional)	Notification template	You can use the preset notification template (alarms will be notified to the root account administrator via SMS and email. This template is selected by default.) or customize one. Up to three notification templates can be bound to each alarm policy. For more information about notification template configuration, see Creating Notification Template.				

4. Click Complete.

## Alarm Records

Last updated : 2023-12-27 15:02:12

## Feature Overview

The alarm records feature allows you to view alarms in the past six months. On the **Alarm Records** page, you can also quickly access alarm policies by policy type and subscribe to them.

## Directions

1. Log in to the EMR console and click the **ID/Name** of the target cluster in the cluster list to enter the cluster details page.

2. Select **Cluster Monitoring > Alarm Records**. On the displayed page, you can perform fuzzy search for alarms by alarm content, filter by time range and alarm status, and sort by start time.

<del>&lt;</del>	Alarm Records									Help Documentatio
Basic Information	Alarm Configuration	Today Yesterday	Last 7 days	2021-11-24 00:00:00 to 2021-11-24 19:31:07					Fuzzy search	of alarm content Q
Bootstrap Actions	Start Time ↓	Policy Type	Alarm Object	Alarm Content		Duration	Alarm Sta T	Policy Name	End Time	Alarm T
Cluster Service					No data yet					
Cluster Resource										
Cluster Monitoring										
Cluster Overview										
* DashBoard BETA										
Cluster Events										
Log Search										
Cluster Inspection										
Alarm Records										
Auto Scaling										
Cluster Scripts										
User Management										
Operation Logs										

3. Click **Alarm Configuration** to go to the alarm configuration page in the Cloud Monitor console to configure alarm policies. For details, see Alarm Configurations.