

Elastic MapReduce

Operation Guide for EMR on CVM



Copyright Notice

©2013–2024 Tencent Cloud. All rights reserved.

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

Trademark Notice

 Tencent Cloud

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

Service Notice

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

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

Contact Us

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

Contents

Operation Guide for EMR on CVM

Planning Cluster

- Business Evaluation

- Cluster Types

Configuring Cluster

- Creating a Cluster

- Administrative rights

 - CAM Overview

 - Role Authorization

 - Collaborator/Sub-account Permissions

 - CAM-Enabled EMR API Authorization Granularity Details

 - Authorization Granularity Scheme

 - Cluster COS Service Role

- Setting Tags

- Software Configuration

- Mounting CHDFS Instance

- Hive Metadata Database

- Security Group Configuration

Cluster Management

- Instance Information

- Node Specification Management

- Inspect and Update Public IP

- Logging in to Clusters

- Cluster Scale-Out

- Cluster Scale-In

- Release now

- Auto Scaling

 - Overview

 - Custom Scaling Configuration

 - Managed Scaling Configuration

- Repairing Disks

- Disk Update Check

- Scaling up Cloud Disks

- Changing Configurations

- Exporting Software Configuration

- Cluster Scripts

- Cluster Restoration

- Cluster Termination

- Operation Log

- Task Center

Managing Service

- User Management

- Restarting Service

- Starting/Stopping Services

- Role Management

- Client Management

- Configuration Management

 - Configuration Management

 - Configuration Status

 - Configuration Rollback

 - Configuration Group Management

- Services

YARN Resource Scheduling

Overview

Configuring Capacity Scheduler

Tag Management

Viewing Scheduling History

HBase RIT Fixing

Component port

Monitoring and Alarming

Node Status

Service Status

Cluster Event

Logs

Application Analysis

HDFS File Storage Analysis

StarRocks Query Management

Hive Query Management

HBase Table Analysis

Kyuubi Query Management

Cluster Inspection**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 History

Operation Guide for EMR on CVM

Planning Cluster

Business Evaluation

Last updated: 2023-12-21 16:00:14

Select a cluster type

EMR clusters offer a variety of cluster types, allowing you to select the type that best suits your business needs:

- **Hadoop Cluster:** Based on open-source Hadoop and its surrounding ecosystem components, it offers five application scenarios: default scenario, Zookeeper, HBase, Presto, Kudu; providing big data solutions that cater to massive data storage, offline/real-time data analysis, streaming data computation, and machine learning scenarios.
- **Kafka Cluster:** A distributed, partitioned, multi-replica, multi-subscriber message processing system, coordinated based on Zookeeper, primarily suitable for asynchronous processing, message communication, and scenarios involving the reception and distribution of streaming data.
- **StarRocks Cluster:** Utilizing comprehensive vectorization technology, it supports a high-speed unified OLAP analysis database, suitable for a variety of data analysis scenarios including multi-dimensional analysis, real-time analysis, and high concurrency scenarios.

Select a billing model

EMR clusters offer two billing models:

- **Annual and Monthly Subscription Clusters:** The billing model for all nodes in the cluster is based on an annual and monthly subscription, suitable for long-term, stable computational clusters.
- **Pay-as-you-go Clusters:** The billing model for all nodes in the cluster is based on pay-as-you-go, suitable for clusters that exist for short periods or periodically.

Note

When shutting down pay-as-you-go nodes of the EMR cluster in the CVM console, please exercise caution when selecting the shutdown mode, as EMR nodes do not support the no-charge shutdown mode.

Selecting Machine Type Specifications

EMR offers a variety of cloud server types, including EMR Standard, EMR Computational, EMR High IO, EMR Memory, and EMR Big Data types (if you require Blackstone 2.0, you can contact us through [Pre-sales Consultation](#)).

You can select the machine type based on your own business needs and cost considerations.

- If you have specific latency requirements for offline computing, we recommend that you opt for local disk or big data machine types.
- If you require the use of the real-time database HBase, we recommend that you select the EMR High IO type and opt for a local SSD disk to achieve optimal performance.
- Local disk models are not supported for deployment on Master and Common nodes. Please select a non-local disk model.

Recommended Node Specifications

EMR defines five types of nodes, from which you can choose based on the cluster type:

| Cluster Type | Scenarios | Node type | Recommended Specification |
|--------------|---------------|-----------|---|
| Hadoop | Default scene | Master | Master Node: It is advisable to opt for an instance specification with a larger memory, preferably at least 8G. Choosing a cloud disk for storage can enhance the stability of the cluster. |
| | | Core | If the majority of your data resides on COS object storage, the functions of Core nodes and Task nodes are similar, with a size not less than 500G. Core nodes do not possess elasticity features. If your architecture does not utilize COS object storage, then the Core nodes are responsible for the cluster's computational and storage tasks. EMR defaults to triple redundancy, so when estimating the size of the data disk, the space for triple redundancy must be considered. It is recommended to use big data models. |

| | | |
|-----------|--------|--|
| | Task | <p>If your architecture does not employ COS object storage, the use of Task nodes can be dispensed with.</p> <p>If the majority of your data is stored on COS object storage, then Task nodes can be utilized as elastic computing resources, acquired as needed.</p> <p>If your cluster's billing mode is on an annual or monthly basis, the billing mode for Task nodes should be pay-as-you-go.</p> |
| | Common | <p>Common nodes: These nodes are primarily used as zk nodes. A minimum specification of 2C4G with a 100G cloud disk can satisfy the requirements.</p> |
| | Router | <p>Router nodes: Primarily used to alleviate the load on the master node and serve as a task submission machine, it is advisable to choose a model with a larger memory, preferably not less than the Master specification.</p> |
| Zookeeper | Common | <p>Common nodes: Primarily used as zk nodes, a minimum specification of 2C4G with a 100G cloud disk is sufficient to meet the requirements.</p> |
| HBase | Master | <p>Master Node: It is advisable to opt for an instance specification with a larger memory, preferably at least 8G. Choosing a cloud disk for storage can enhance the stability of the cluster.</p> |
| | Core | <p>If the majority of your data is stored on COS object storage, then the functions of Core nodes and Task nodes are similar, with a size not less than 500G.</p> <p>Please note, Core nodes do not possess elastic capabilities.</p> <p>If your architecture does not utilize COS object storage, then the Core nodes are responsible for the cluster's computational and storage tasks.</p> |
| | Task | <p>If your architecture does not employ COS object storage, the use of Task nodes can be dispensed with.</p> <p>If the majority of your data is stored on COS object storage, then Task nodes can be utilized as elastic computing resources, acquired as needed.</p> <p>If your cluster's billing mode is on an annual or monthly basis, and you require the Task node's billing mode to be pay-as-you-go, then you need to set the Task node quantity to zero here. You can then scale up the pay-as-you-go Task nodes as needed via the console or API.</p> |
| | Common | <p>Common nodes: Primarily used as zk nodes, a minimum specification of 2C4G with a 100G cloud disk is sufficient to meet the requirements.</p> |
| kudu | Router | <p>Router nodes: Primarily used to alleviate the load on the master node and serve as a task submission machine, it is advisable to choose a model with a larger memory, preferably not less than the Master specification.</p> |
| | Master | <p>Master Node: It is advisable to opt for an instance specification with a larger memory, preferably at least 8G. Choosing a cloud disk for storage can enhance the stability of the cluster.</p> |
| | Core | <p>If the majority of your data is stored on COS object storage, then the functions of Core nodes and Task nodes are similar, with a size not less than 500G.</p> <p>Please note: Core nodes do not possess elastic capabilities.</p> <p>If your architecture does not utilize COS object storage, then the Core nodes are responsible for the cluster's computational and storage tasks. EMR defaults to triple redundancy, so when estimating the size of the data disk, the space for triple redundancy must be considered. It is recommended to use big data models.</p> |
| | Task | <p>If your architecture does not employ COS object storage, the use of Task nodes can be dispensed with.</p> <p>If the majority of your data is stored on COS object storage, then Task nodes can be utilized as elastic computing resources, acquired as needed.</p> <p>If your cluster's billing mode is on an annual or monthly basis, and you require the Task node's billing mode to be pay-as-you-go, then you need to set the Task node quantity to zero here. You can then scale up the pay-as-you-go Task nodes as needed via the console or API.</p> |
| | Common | <p>Common nodes: Primarily used as zk nodes, a minimum specification of 2C4G with a 100G cloud disk is sufficient to meet the requirements.</p> |

| | | | |
|-----------|---------------|--------|---|
| | presto | Router | Router nodes: Primarily used to alleviate the load on the master node and serve as a task submission machine, it is advisable to choose a model with a larger memory, preferably not less than the Master specification. |
| | | Master | Master Node: It is advisable to opt for an instance specification with a larger memory, preferably at least 8G. Choosing a cloud disk for storage can enhance the stability of the cluster. |
| | | Core | If the majority of your data is stored on COS object storage, then the functions of Core nodes and Task nodes are similar, with a size not less than 500G. Please note: Core nodes do not possess elastic capabilities. If your architecture does not utilize COS object storage, then the Core nodes are responsible for the cluster's computational and storage tasks. EMR defaults to triple redundancy, so when estimating the size of the data disk, the space for triple redundancy must be considered. It is recommended to use big data models. |
| | | Task | If your architecture does not employ COS object storage, the use of Task nodes can be dispensed with. If the majority of your data is stored on COS object storage, then Task nodes can be utilized as elastic computing resources, acquired as needed. If your cluster's billing mode is on an annual or monthly basis, and you require the Task node's billing mode to be pay-as-you-go, then you need to set the Task node quantity to zero here. You can then scale up the pay-as-you-go Task nodes as needed via the console or API. |
| | | Common | Common nodes: Primarily used as zk nodes, a minimum specification of 2C4G with a 100G cloud disk is sufficient to meet the requirements. |
| Kafka | Default scene | Router | Router nodes: Primarily used to alleviate the load on the master node and serve as a task submission machine, it is advisable to choose a model with a larger memory, preferably not less than the Master specification. |
| | | Core | Core Nodes: It is recommended to choose models with higher CPU and memory. As local disks run the risk of data loss in the event of disk failure, it is advisable to opt for cloud disks. |
| StarRocks | Default scene | Common | Common Nodes: It is recommended that the minimum configuration for CPU and memory should not be less than 4C16G. |
| | | Master | Master Nodes: It is recommended to choose instance specifications with larger memory, with a recommended memory size of at least 8G. All metadata on the Master nodes is stored in memory. |
| | Default scene | Core | Core Nodes: It is recommended to choose instance specifications with larger memory, with a recommended memory size of at least 8G. For better IO performance and stability, it is advisable to use cloud SSD disks. |
| | | Router | Router Nodes: These deploy the Frontend module, achieving high availability for read and write operations. Therefore, it is recommended to choose models with larger memory, preferably not less than the Master specifications. |

Note

- Different cluster types have varying requirements for node specifications. Currently, the system will default to recommending configurations that meet the cluster requirements. You can adjust the model specifications according to business needs, with the recommended models serving only as a reference.
- Core Nodes do not possess elasticity features. If your architecture does not utilize COS object storage, then the Core Nodes are responsible for the cluster's computation and storage tasks. EMR defaults to triple backup, so when estimating the size of the data disk, the space for triple backup must be considered. It is recommended to use big data models.

Network and Security

To ensure the network security of the cluster, the EMR cluster will be placed within a VPC, to which we will add a security group policy. To ensure convenient access to the WebUI of Hadoop ecosystem components, we have enabled an external IP for one of the Master nodes, which is billed according to traffic. By default, Router nodes do not have external IP enabled. If needed, you can freely bind an elastic public IP in the CVM console.

Note

- Master Nodes default to enabling external IP when creating a cluster, but users can choose not to enable external IP based on their circumstances.
- Enabling the public network for the cluster's Master Nodes is primarily used for SSH login and viewing component WebUI.
- The Master Nodes will enable external network, which is billed based on traffic, with a bandwidth cap of 5M. After creating the cluster, you can adjust this network in the console.

Cluster Types

Last updated: 2023-12-21 16:00:37

EMR accommodates a variety of cluster types and corresponding application scenarios, defining five kinds of node types. The node types supported, the number of nodes deployed, and the services deployed vary among different cluster types and application scenarios. Depending on the business requirements, one can choose different cluster types and application scenarios to create a cluster.

Cluster Type Description

Hadoop Cluster

| Scenarios | Description | Node Deployment Description |
|---------------|---|---|
| Default scene | Based on open-source Hadoop and its surrounding ecosystem components, we offer big data solutions for scenarios such as massive data storage, offline/real-time data analysis, streaming data computation, and machine learning. | <ul style="list-style-type: none"> • Master Node: Serves as the management node, ensuring the normal scheduling of the cluster; primarily deploys processes such as NameNode, ResourceManager, HMaster; the quantity is 1 in non-HA mode, and 2 in HA mode. Note: When the deployment components include Kudu, the cluster only supports HA mode, with the number of Master nodes being 3. • Core Node: Acts as a computation and storage node, all your data in HDFS is stored in the Core nodes. Therefore, to ensure data safety, it is not permissible to scale down after expanding the Core nodes; primarily deploys processes such as DataNode, NodeManager, RegionServer. The quantity is ≥ 2 in non-HA mode, and ≥ 3 in HA mode. • Task Node: Functions as a pure computation node, does not store data, the data to be computed comes from Core nodes and COS. Therefore, Task nodes are often used as elastic nodes, which can be expanded and scaled down at any time; primarily deploys processes such as NodeManager, PrestoWork; the number of Task nodes can be changed at any time to achieve elastic scaling of the cluster, with a minimum value of 0. • Common Node: Provides data sharing synchronization and high availability fault tolerance services for HA cluster Master nodes; primarily deploys distributed coordinator components, such as ZooKeeper, JournalNode nodes. The quantity is 0 in non-HA mode, and ≥ 3 in HA mode. • Router Node: Utilized to share the load of the Master nodes or to serve as the task submission machine for the cluster, it can be expanded and scaled down at any time; primarily deploys Hadoop packages, with the option to deploy software and processes such as Hive, Hue, Spark; the number of Router nodes can be changed at any time, with a minimum value of 0. |
| ZooKeeper | Applicable for establishing distributed, high-availability coordination services in large-scale clusters. | Common Node: Primarily deploys the distributed coordinator component ZooKeeper, the number of deployed nodes must be odd, with a minimum of 3 Common nodes, and only supports High Availability (HA). |
| HBase | Appropriate for storing massive amounts of unstructured or semi-structured data, offering a highly reliable, high-performance, column-oriented, and scalable distributed storage system for real-time data read and write operations. | <ul style="list-style-type: none"> • Master Node: Serves as the management node, ensuring the normal scheduling of the cluster; primarily deploys processes such as NameNode, ResourceManager, HMaster; the quantity is 1 in non-HA mode, and 2 in HA mode. • Core Node: Acts as a computation and storage node, all your data in HDFS is stored in the Core nodes. Therefore, to ensure data safety, it is not permissible to scale down after expanding the Core nodes; primarily deploys processes such as DataNode, NodeManager, RegionServer. The quantity is ≥ 2 in non-HA mode, and ≥ 3 in HA mode. • Task Node: Acts as a pure computing node, does not store data, the data to be computed comes from Core nodes and COS, hence Task nodes are often used as elastic nodes, capable of being expanded and scaled down at any time; primarily deploys processes such as NodeManager; the number of Task nodes can be changed at any time to achieve elastic scaling of the cluster, with a minimum value of 0. |

| | | |
|--------|---|---|
| | | <ul style="list-style-type: none"> • Common Node: Provides data sharing synchronization and high availability fault tolerance services for HA cluster Master nodes; primarily deploys distributed coordinator components, such as ZooKeeper, JournalNode nodes. The quantity is 0 in non-HA mode, and ≥ 3 in HA mode. • Router Node: Employed to alleviate the load of the Master nodes or to function as the task submission mechanism for the cluster, it can be expanded and scaled down at any time; the quantity of Router nodes can be modified at any time, with a minimum value of 0. |
| Presto | Offers an open-source distributed SQL query engine, suitable for interactive analytical queries, supporting rapid query analysis of massive data. | <ul style="list-style-type: none"> • Master Node: Functions as a management node, ensuring the normal scheduling of the cluster; primarily deploys processes such as NameNode, ResourceManager, etc.; the quantity is 1 in non-HA mode, and 2 in HA mode. • Core Node: Functions as a computing and storage node, all your data in HDFS is stored in Core nodes, therefore, to ensure data safety, it is not permissible to scale down after expanding Core nodes; primarily deploys processes such as DataNode, NodeManager, etc. The quantity is ≥ 2 in non-HA mode, and ≥ 3 in HA mode. • Task Node: Functions as a pure computation node, does not store data, the data to be computed comes from Core nodes and COS. Therefore, Task nodes are often used as elastic nodes, which can be expanded and scaled down at any time; primarily deploys processes such as NodeManager, PrestoWork; the number of Task nodes can be changed at any time to achieve elastic scaling of the cluster, with a minimum value of 0. • Common Node: Provides data sharing synchronization and high availability fault tolerance services for HA cluster Master nodes; primarily deploys distributed coordinator components, such as ZooKeeper, JournalNode nodes. The quantity is 0 in non-HA mode, and ≥ 3 in HA mode. • Router Node: Employed to alleviate the load of the Master nodes or to function as the task submission mechanism for the cluster, it can be expanded and scaled down at any time; the quantity of Router nodes can be modified at any time, with a minimum value of 0. |
| Kudu | Provides a distributed, scalable columnar storage manager, supporting random read/write and OLAP analysis for processing rapidly updating data. | <ul style="list-style-type: none"> • Master Node: Functions as a management node, ensuring the normal scheduling of the cluster; primarily deploys processes such as NameNode, ResourceManager, etc.; the quantity is 1 in non-HA mode, and 2 in HA mode. • Core Node: Serves as a computational and storage node, all your data in HDFS is entirely stored within the Core nodes, hence, to ensure data security, it is not permissible to scale down after expanding Core nodes; the quantity is ≥ 2 in non-HA mode, and ≥ 3 in HA mode. • Task Node: Functions as a pure computational node, does not store data, the data being computed comes from Core nodes and COS, hence, Task nodes are often used as elastic nodes, capable of scaling up and down at any time; the number of Task nodes can be altered at any time to achieve elastic scaling of the cluster, with a minimum value of 0. • Common Node: Provides data sharing synchronization and high-availability fault tolerance services for HA cluster Master nodes; primarily deploys distributed coordination components, such as ZooKeeper, JournalNode, etc.; the quantity is 0 in non-HA mode, and ≥ 3 in HA mode. • Router Node: Employed to alleviate the load of the Master nodes or to function as the task submission mechanism for the cluster, it can be expanded and scaled down at any time; the quantity of Router nodes can be modified at any time, with a minimum value of 0. |

Kafka Cluster

| Scenarios | Description | Node Deployment Description |
|---------------|--|--|
| Default scene | Offers a distributed, partitioned, multi-replica, multi-subscriber messaging system, coordinated based on ZooKeeper, primarily | <ul style="list-style-type: none"> • Core Node: Functions as the Backend module, primarily providing data storage capabilities; deploys processes such as BE, Broker, etc., the quantity is ≥ 1 in non-HA mode, and ≥ 2 in HA mode. • Common Node: Provides data sharing synchronization and high-availability fault tolerance services for HA cluster Core nodes; the quantity is 0 in non-HA mode, and ≥ 3 in HA mode. |

| | | |
|--|---|--|
| | <p>suitable for asynchronous processing, message communication, and scenarios involving the reception and distribution of streaming data.</p> | |
|--|---|--|

StarRocks Cluster

| Scenarios | Description | Node Deployment Description |
|---------------|---|---|
| Default scene | <p>StarRocks employs comprehensive vectorization technology, supporting a high-speed unified OLAP analysis database, suitable for various data analysis scenarios such as multi-dimensional analysis, real-time analysis, and high concurrency.</p> | <ul style="list-style-type: none"> • Master Node: Functions as the Frontend module, while also providing Web UI capabilities; deploys processes such as FE Follower, Broker, etc., the quantity is ≥ 1 in non-HA mode, and ≥ 3 in HA mode. • Core Node: Functions as the Backend module, primarily providing data storage capabilities; deploys processes such as BE, Broker, etc., with a deployment quantity of ≥ 3. • Task Node: Functions as a pure computation node, does not store data, and the data to be computed comes from Core nodes and COS. Therefore, Task nodes are often used as elastic nodes, capable of scaling up and down at any time; primarily deploys Compute Node processes; the number of Task nodes can be changed at any time to achieve elastic scaling of the cluster, with a minimum value of 0. • Router Node: Deploys the Frontend module, achieving high availability for read and write operations; can optionally deploy FE Observer, Broker, and other processes, can scale up by adding Router nodes, but does not support scaling down. |

Configuring Cluster

Creating a Cluster

Last updated: 2023-12-21 16:01:42

Scenario

This document provides a detailed guide on how to create an EMR on CVM cluster via the EMR console, along with the associated configurations.

Instructions

Log in to the [EMR Console](#), on the EMR on CVM cluster list page, click on 'Create Cluster'. During the cluster creation process, you will need to complete four steps: software configuration, region and hardware configuration, basic configuration, and confirmation of configuration information.

Software configuration

| Configuration items | Configuration Notes |
|---------------------------|--|
| Regions | The cluster is deployed in a physical data center, with each region representing an independent physical data center. The internal networks of cloud servers in different regions are not interconnected. Please note: once the cluster is created, the region cannot be changed, so choose carefully. |
| Cluster Type | EMR on CVM supports a variety of cluster types. Please select the appropriate type based on your actual business needs for cluster creation. The default is the Hadoop cluster type. For an introduction to each cluster type, please refer to Cluster Type . |
| Scenarios | The Hadoop cluster type supports five scenarios: Default, Zookeeper, HBase, Presto, and Kudu. Choose the appropriate scenario for deployment based on your actual business needs. |
| Instance Edition | The components and their versions bundled with different product editions vary. |
| Component Deployment | Optional components can be customized and deployed according to individual requirements. |
| Kerberos Mode | By default, this feature is disabled. Once enabled, the open-source components in the cluster will start in Kerberos security mode. For more information, please refer to Introduction to Kerberos . |
| Component dependency mode | By default, this feature is disabled. Once activated, the components of the pre-built cluster are shared for use in the current cluster. For an introduction to the dependency mode, please refer to Component Configuration Sharing . Note: Since the selected components are deployed in an existing cluster, the cluster providing the dependent components cannot be directly destroyed. All clusters dependent on its components must be destroyed first. |
| Software configuration | Optional configurations are available. Prior to the cluster launch, you can specify a json file to modify the component configuration parameters, or access an external cluster. For more details, please refer to Software Configuration . |

Elastic MapReduce [Back to Product Details](#) [Documentation](#) [Billing](#) [Console](#)

1 Software Configuration 2 AZ and Hardware Configuration 3 Basic Configuration 4 Configuration Confirmation

Software Configuration

Region: South China, East China, North China, **Central and Southwest China**, Hong Kong/Macao/Taiwan (China), Asia Pacific, North America and Europe

City: Chengdu

Cluster Type:

- Hadoop**: A big data distributed framework, suitable for offline/real-time big data analysis
- Druid**: A column-oriented storage engine, suitable for high-concurrency real-time analysis
- ClickHouse**: A column-oriented storage and analytics engine, suitable for real-time wide table analysis

Use Cases: **Hadoop-Default**, Zookeeper, Hbase, Presto, Kudu

Product Version: EMR-V2.2.0.tlinux [Product Release Notes](#)

Components to Deploy:

| | | | | | | | | | |
|----------------|-----------------|-----------------|---------------|------------|--------------|-------------|-------------|------------|--------------|
| hdfs-2.6.5 | gpm-2.6.5 | zookeeper-3.5.4 | kylin-1.2.0 | hive-2.3.5 | tez-0.9.2 | hbase-1.4.9 | spark-2.4.3 | livy-0.7.0 | presto-0.228 |
| impala-2.10.0 | kylin-2.5.2 | flink-1.9.2 | storm-1.2.3 | hudi-0.5.1 | ranger-1.2.0 | sqoop-1.4.7 | flume-1.9.0 | hue-4.6.0 | oozie-5.1.0 |
| zeppelin-0.8.2 | superset-0.35.2 | alluxio-1.8.1 | ganglia-3.7.2 | | | | | | |

Advanced settings:

- Kerberos mode: Enable Kerberos authentication
- Component Dependency Mode: Enable

Software Configuration:

Next

Region and Hardware Configuration

| Configuration items | Configuration Notes |
|---------------------------|---|
| Billing Mode | Two billing modes are supported: prepaid (annual/monthly) and pay-as-you-go. With the prepaid mode, you prepay the product fees for N months in advance, which is more cost-effective compared to pay-as-you-go. In the pay-as-you-go mode, you pay based on usage duration. Real-name verification is required for your account, and a fee equivalent to two hours of usage is frozen upon activation (coupons cannot be used as freezing vouchers). The frozen resource fee is refunded upon termination. |
| (Cross) Availability Zone | <ol style="list-style-type: none"> You can choose whether to deploy across zones as needed. Cross-Availability Zone deployment will distribute services and management roles evenly across multiple Availability Zones, providing varying degrees of high availability for services. By default, it is a single Availability Zone, with cross-Availability Zone whitelisting available. Different availability zones within the same region support different machine specifications; it is recommended to choose the latest availability zone. Cloud products in different regions do not communicate with each other via the internal network and cannot be switched after purchase. It is advisable to select an availability zone in a region close to your business data to reduce access latency and increase download speed. |
| Deployment Strategy | When choosing to deploy across availability zones, both average and balanced strategies are supported. In the average strategy, a "primary availability zone + secondary availability zone" deployment scheme is used. If the secondary availability zone experiences an issue, the primary availability zone can continue to provide services normally. Balanced Strategy: A deployment scheme of "Primary Availability Zone + Secondary Availability Zone + Balanced Availability Zone", where the primary and secondary zones are deployed in a mutually supportive manner. If any availability zone experiences an issue, it does not affect the normal service provision of the other zones. |
| Cluster network | To ensure the security of the EMR cluster, we will place each node of the cluster into a private network. You need to set up a private network to ensure the correct creation of the EMR cluster. |
| External Cluster Network | This can be used for SSH login from the external network and for accessing the component WebUI from the external network. After the successful creation of the cluster, adjustments can be made to this network via the console. By default, the external network of the master1 node is enabled. |
| Security Group | The Security Group functions as a firewall, used for setting the network access control of the Cloud Server CVM. If there is no Security Group, EMR will automatically create a new one for you. If there is an |

| | |
|------------------------|---|
| | <p>existing Security Group in use, it can be directly selected. If the number of Security Groups has reached the maximum limit and a new one cannot be created, some of the unused Security Groups can be deleted. You can view the Security Groups that are currently in use.</p> <p>Creating a Security Group: EMR assists users in creating a Security Group, enabling port 22 and 30001, along with the necessary internal communication network segments.</p> <p>Existing EMR Security Group: Select an already created EMR Security Group as the Security Group for the current instance, enabling port 22 and 30001, along with the necessary internal communication network segments.</p> |
| Remote login | Port 22 is commonly used for remote login. When creating a new Security Group, this port will be enabled by default. You can close this port based on business needs, although it is enabled by default. |
| High Availability (HA) | High Availability is enabled by default. Different cluster types and application scenarios have varying node type deployment quantities under HA or non-HA scenarios. For more details, please refer to Cluster Types . |
| Node type | <p>Select the appropriate machine configuration for different node types based on business needs. For more details, please refer to Business Evaluation.</p> <p>Storage fees description</p> <ol style="list-style-type: none"> Currently, we support mounting various types of cloud disks (each type can only be selected once) and multiple cloud disks (up to 20) on Core nodes, Task nodes, and Router nodes. Local disk models are not supported for deployment on Master and Common nodes. Please select a non-local disk model. |
| Placement group | Optional configuration, placement groups are strategies for distributing cloud server instances on underlying hardware. For more details, please refer to Placement Groups . |
| Hive Metadata Database | <p>If the Hive component is selected, the Hive metadata database offers two storage methods: The first default cluster, Hive metadata is stored in the MetaDB purchased independently by the cluster; The second method involves associating with an external Hive metadata database. You can choose to associate with EMR-MetaDB or a self-built MySQL database. The metadata will be stored in the associated database and will not be destroyed with the cluster. For more details, please refer to Hive Metadata Management.</p> <p>Please note: When one or more components such as Hue, Ranger, Oozie, Druid, or Superset are selected, the system will automatically purchase a MetaDB for metadata storage of components other than Hive.</p> |

The screenshot shows the configuration console for Elastic MapReduce. At the top, there are navigation links for 'Documentation', 'Billing', and 'Console'. A progress bar indicates four steps: 1. Software Configuration, 2. AZ and Hardware Configuration (current step), 3. Basic Configuration, and 4. Configuration Confirmation.

Billing Type
Billing mode: Pay-as-you-go

AZ and Network Configuration
AZ: [Visual representation of Availability Zones with a tooltip: "Cloud products in different AZs are not interconnected over private networks. The AZ cannot be changed after purchase. You are advised to select an AZ closest to your business data to reduce access delay and increase download speed."]

Cluster Public Network: Enable public network for cluster nodes

Cluster Network: [Please select] [Please select] - subnet IPs in total, - available.
If existing networks do not meet your needs, you can create a [VPC](#) or [subnet](#) in the console.

Secure Login

Security Group ⓘ Create One Select an Existing One

EMR will create a security group and open ports 22 and 30001 as well as necessary private network IP ranges. The new security group will be named emr-xxxxxxx_yyyyMMdd. Do not modify the security group name.
 Mouse over to view [outbound rules](#) and [inbound rules](#)

Remote Login ⓘ Enable

Node Configuration

High Availability (HA) ⓘ Enable

| Node Type | Node spec | Node count |
|--------------------------|---|------------|
| MasterNode Configuration | Standard SA2: 4-core16G System Disk: SSD Cloud Disk50G*1 Data Disk: SSD Cloud Disk200G*1 Modify | - 2 + |
| CoreNode Configuration | Standard SA2: 4-core8G System Disk: SSD Cloud Disk50G*1 Data Disk: SSD Cloud Disk200G*1 Modify | - 3 + |
| TaskNode Configuration | Standard SA2: 4-core8G System Disk: SSD Cloud Disk50G*1 Data Disk: SSD Cloud Disk200G*1 Modify | - 0 + |
| CommonNode Configuration | Standard SA2: 2-core4G System Disk: SSD Cloud Disk50G*1 Data Disk: SSD Cloud Disk200G*1 Modify | - 3 + |

Basic configurations

| Configuration items | Configuration Notes |
|---------------------|---|
| Project | Assign the current cluster to different project groups. Please note: The project to which the cluster belongs cannot be modified after the cluster is created. |
| Cluster name | Distinguish different EMR clusters by setting unique cluster names. The system generates a modifiable random name for the cluster, with a length limit of 6–36 characters, and only allows the inclusion of Chinese characters, letters, numbers, hyphens, and underscores. |
| Login Options | Currently, EMR offers two methods for logging into cluster services, nodes, and MetaDB: custom password setting and associated key methods. The SSH key is only used for EMR–UI shortcut login. The default username is "root", while the username for the Superset component WebUI shortcut is "admin". |
| Bootstrap actions | Optional configuration, the bootstrap script operation allows you to execute custom scripts during the cluster creation process. This facilitates modifications to the cluster environment, installation of third-party software, and utilization of proprietary data. For more settings, please refer to Bootstrap Actions . |
| Tag | Optional configuration, you can add tags to the cluster or node resources during creation, facilitating the management of cluster and node resources. A maximum of five tags can be bound, and tag keys must be unique. |

Elastic MapReduce [Back to Product Details](#) [Documentation](#) [Billing](#) [Console](#)

1 Software Configuration — 2 AZ and Hardware Configuration — **3 Basic Configuration** — 4 Configuration Confirmation

Basic Configuration

Project:

Cluster Name:

Login Method: Set Password Associate Key Use Guide [If existing keys do not meet your needs, you can Create Now](#)

Set Password:

Password:

[Advanced settings](#)

| Run | Name | Script Location | Parameter | Operation |
|--|------|-----------------|-----------|-----------|
| No data yet | | | | |
| + Add Bootstrap Action | | | | |

Tag: [+ Add](#)

Configuration Cost **0.828 USD/hour** 0.828 USD/hour

Configuration confirmation

| Configuration items | Configuration Notes |
|---------------------|--|
| Configuration list | Confirm whether the deployed information is accurate. |
| Auto-Renewal | Optional feature, the system will check the available balance on the user's account seven days prior to the cluster's expiration, renewing the resources of the cluster set for auto-renewal. The annual and monthly clusters are set to auto-renew by default, but users can manually deselect this option. |
| Terms of Service | Agree to the "Elastic MapReduce Service Level Agreement" and "Refund Agreement" . |

Elastic MapReduce [Back to Product Details](#) [Documentation](#) [Billing](#) [Console](#)

1 Software Configuration — 2 AZ and Hardware Configuration — 3 Basic Configuration — **4 Configuration Confirmation**

Configuration List

| Software Configuration | | | | | | Edit |
|------------------------|--|--------------|---------|-----------------|-------------------|----------------------|
| Region | Guangzhou | Cluster Type | Hadoop | Use Cases | Hadoop-Default | |
| Components to Deploy | hdfs-2.8.5,yarn-2.8.5,zookeeper-3.5.5,knox-1.2.0 | Kerberos | Disable | Product Version | EMR-V2.2.0.tlinux | |

| AZ and Hardware Configuration | | | | | | Edit |
|-------------------------------|--|------------------------|--|------------------------|--|----------------------|
| Billing mode | Pay-as-you-go | AZ | Guangzhou Zone 7 | Cluster Public Network | Enable | |
| Security Group | emr-g9k9z7r_20220421 | High Availability (HA) | Disable | Hive Metadatabase | None | |
| Cluster Network | Roy-0018&ceshi001 | MetaDB | --- | | | |
| MasterNode | Standard SA3: 8-core16G System Disk: SSD Cloud Disk50G*1 Data Disk: SSD Cloud Disk200G*1 Instance count*1 | CoreNode | Standard SA3: 8-core16G System Disk: SSD Cloud Disk50G*1 Data Disk: SSD Cloud Disk200G*1 Instance count*2 | TaskNode | Standard SA3: 8-core16G System Disk: SSD Cloud Disk50G*1 Data Disk: SSD Cloud Disk200G*1 Instance count*0 | |

| Basic Configuration | | | | | | Edit |
|---------------------|-----------------|--------------|--------------|-----------------|---------|----------------------|
| Project | DEFAULT PROJECT | Cluster Name | EMR-51gr0xfc | Disk Encryption | Disable | |

Agreement: I agree to the [Elastic MapReduce Service Level Agreement](#) and [Refund Policy](#)

Upon completion of the above configurations, click 'Purchase' to proceed with payment. Once the payment is successful, the EMR cluster enters the creation process. The newly created cluster can be found in the EMR console approximately 10 minutes later.

Note

- Pay-as-you-go clusters: The creation process begins immediately. Once the cluster is successfully created, its status

changes to 'Running'.

- Annual and monthly subscription clusters: An order is first generated. The creation of the cluster commences only after the order has been paid in full.
- You can view the instance information of each node in the CVM console. To ensure the normal operation of the EMR cluster, please refrain from altering the configuration information of these instances within the CVM console.

Subsequent steps

Upon successful creation of the cluster, you may log in according to your circumstances and perform further configurations and operations on the cluster. For specific operations, please refer to the following documentation:

- Cluster Configuration: [Software Configuration](#), [Mounting CHDFS](#), [Unifying HIVE Metadata](#)
- Cluster Management: [Setting Tags](#), [Setting Boot Actions](#), [Cluster Termination](#)

Administrative rights

CAM Overview

Last updated: 2023-12-21 16:03:31

Cloud Access Management (CAM) Introduction

While utilizing Tencent Cloud's Elastic MapReduce (EMR) services, different departments and personnel roles require varying levels of access, mitigating potential security risks such as information leaks and operational errors. In such instances, you can assign distinct permissions to different individuals through sub-accounts, thereby circumventing the aforementioned issues. By default, sub-accounts do not possess the authority to use EMR or access related resources. Consequently, it is necessary to create policies that specify the resources or permissions required by sub-accounts.

Cloud Access Management (CAM) is a web service provided by Tencent Cloud, designed to assist customers in securely managing access permissions, resource management, and usage rights for their Tencent Cloud accounts. With CAM, you can create, manage, and terminate users (or groups), and control who can use which Tencent Cloud resources through identity management and policy administration.

When utilizing EMR, policies can be associated with an individual user or a group of users, granting or denying them the ability to complete specified tasks with designated resources. For more fundamental information related to CAM policies, please refer to [Policy Syntax](#). For more practical information regarding CAM policies, please consult [Policies](#).

CAM Application Scenarios

1. Assigning Resource Management Permissions to Sub-users

Within CAM, you can create users or roles, assign them individual security credentials (console login passwords, cloud API keys, etc.), or request temporary security credentials for accessing Tencent Cloud resources. You can manage permissions to control which actions users and roles can execute and which resources they can access.

2. Possessing external Tencent Cloud identities and these users require access to Tencent Cloud resources.

Through CAM, you can leverage your existing authentication system to provide employees and services with access to Tencent Cloud services and resources.

3. Need to provide an additional layer of protection for sub-users beyond username and password.

Supports three verification methods: WeChat QR code verification, MFA device verification (divided into hardware MFA device verification and virtual MFA device verification), and mobile phone verification code verification.

Product Support for CAM Granularity Introduction

EMR supports API granularity authentication divided into **resource-level** and **interface-level authentication**. The interface input includes the cluster string ID for interface-level authentication, which is mainly forwarded directly to the CAM side for authentication by the cloud API; resource-level authentication is carried out by the EMR backend for fine-grained authentication. EMR supports three methods of authorization: resource-level authorization, interface-level authorization, and tag-based authorization.

Resource-level and Interface-level Authorization: You can grant sub-accounts the authority to manage individual resources through policy syntax. For details, please refer to the [Authentication Granularity Scheme](#).

Tag-based Authorization: You can assign management rights to sub-accounts for resources under corresponding tags by tagging resources. For details, please refer to the [Authentication Granularity Scheme](#).

Role Authorization

Last updated: 2023-12-21 16:03:46

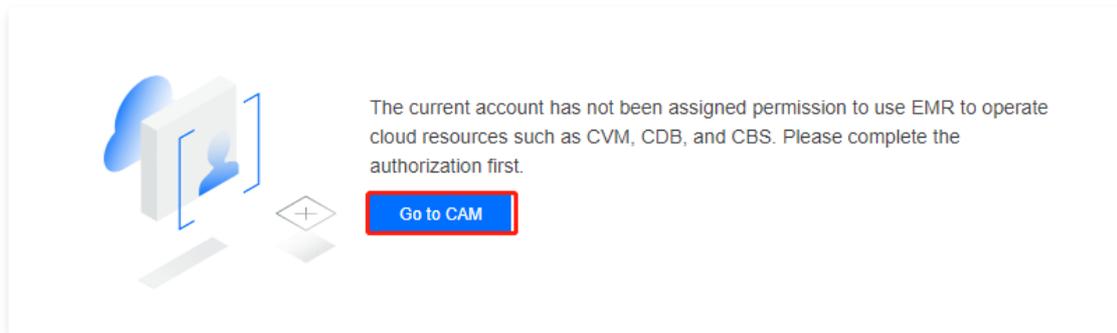
When utilizing the Elastic MapReduce service, users are required to grant the system default role, EMR_QCSRole, to the service account. Only after this role has been successfully assigned, can Elastic MapReduce invoke related services (such as CVM, COS, etc.) to establish clusters and preserve logs.

Note

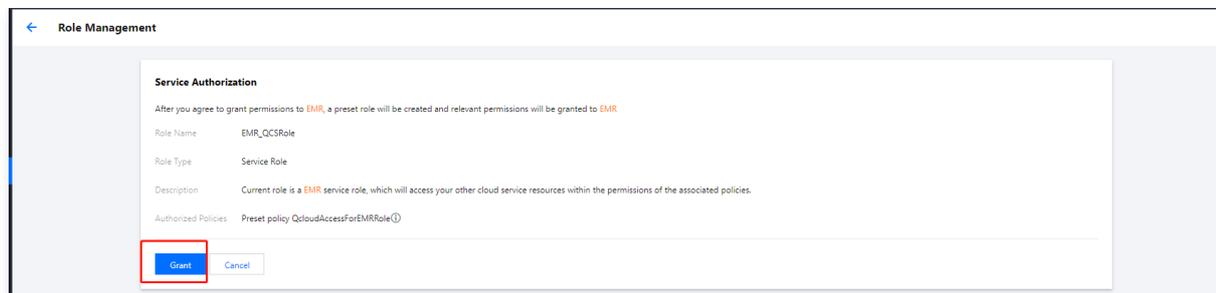
Upon initial activation of the Elastic MapReduce service, the primary account must complete the role authorization process. Failure to do so will result in both the sub-account and the primary account being unable to utilize the Elastic MapReduce service.

Role Authorization Process

1. In the event of a failure in granting the EMR_QCSRole role to the service account when a user is creating a cluster or establishing an on-demand execution plan, the following prompt will appear. Subsequently, click on **Go to CAM** to carry out role authorization.



2. Click on **Grant** to assign the default role, EMR_QCSRole, to the service account of Elastic MapReduce.



3. Upon completion of the authorization, users need to refresh the Elastic MapReduce console or purchase page to resume normal operations. For more detailed policy information related to EMR_QCSRole, please log in to the [Access Management Console](#). For detailed information on the permissions included in EMR_QCSRole, please refer to [Collaborator/Sub-account Permissions](#).

Collaborator/Sub-account Permissions

Last updated: 2023-12-21 16:03:52

Tencent Cloud's Elastic MapReduce (hereinafter referred to as EMR), necessitates access or manipulation of other related cloud products during its operation. To ensure the smooth operation and usage of EMR by collaborators or sub-users, this document provides guidance on the granting of relevant permissions.

Overview of Permission Policies

| Policy Name | Description | Required | Description |
|-------------------------|---|--------------|--|
| QcloudCamRoleFullAccess | Full Read/Write Access Permissions for User and Authority (CAM) Roles | Not required | Utilized for customizing service roles, enabling granular permission control over data access between products. Refer to Cluster COS Service Roles for more details. |
| QcloudEMRFullAccess | Full Read/Write Access Permissions for Elastic MapReduce (EMR) | Not required | Full operational permissions for all features of the EMR product. Refer to Purchasing and Managing EMR Clusters for more details. |
| QcloudEMRReadOnlyAccess | Read-Only Access Permissions for Elastic MapReduce (EMR) | Not required | Viewing Permissions for All Features of the EMR Product |
| QcloudEMRPurchaseAccess | Financial Permissions for the Elastic MapReduce Product | Not required | Refer to Purchasing and Managing EMR Clusters . If there is no need for purchasing or changing configurations, this permission may not be enabled. |

Note

The preset policy QcloudEMRPurchaseAccess allows you to manage the financial permissions for all users to purchase the Elastic MapReduce product. This policy grants users financial permissions for CVM, CDB, and EMR simultaneously. If there is a need to restrict users from purchasing CVM and CDB, the corresponding product order operation permissions may not be granted.

Use Cases

Granting EMR Access to Other Cloud Product Permissions

- Service Role (Mandatory):** When using the EMR product service, it is necessary to access services such as CVM, CBS, and TencentDB on the cloud. During the initial product purchase, the EMR product service role EMR_QCSRole must be assigned and the policy QcloudAccessForEMRRole (used to apply for basic resources such as CVM, CBS, TencentDB, as well as read permissions for COS) and the Elastic MapReduce (EMR) access permissions to cloud resources must be granted.
- Service-Related Role (Optional):** EMR directly writes to or computes data in the Object Storage (COS). To ensure data security, EMR needs to be granted read and write access to COS resources through temporary keys. This requires granting the service-related role EMR_QCSLinkedRoleInApplicationDataAccess and binding the preset policy QcloudAccessForEMRLinkedRoleInApplicationDataAccess.

Special Instructions for COS Bucket Access Authorization:

- As of August 20, 2023, for new users or existing users modifying authorization policies, the default authorized service-related role is EMR_QCSLinkedRoleInApplicationDataAccess.
- Existing users currently have the policy QcloudAccessForEMRRoleInApplicationDataAccess bound to the service role EMR_QCSRole in their authorization strategy.
- When both the service-related role and the service role are authorized, the service-related role is used by default. In the cluster instance information authorization policy, the object storage is shown as authorized, and the cluster COS service role echo display shows the EMR_QCSLinkedRoleInApplicationDataAccess role.

Purchasing and Managing EMR Clusters

In scenarios involving resource purchases such as creating clusters, adding components, changing configurations, and expanding collaborators/sub-users, the QcloudEMRFullAccess and custom TencentDB purchase policies need to be assigned. If there are no resource purchase scenarios such as service configuration management and restarts, only the QcloudEMRFullAccess policy needs to be assigned.

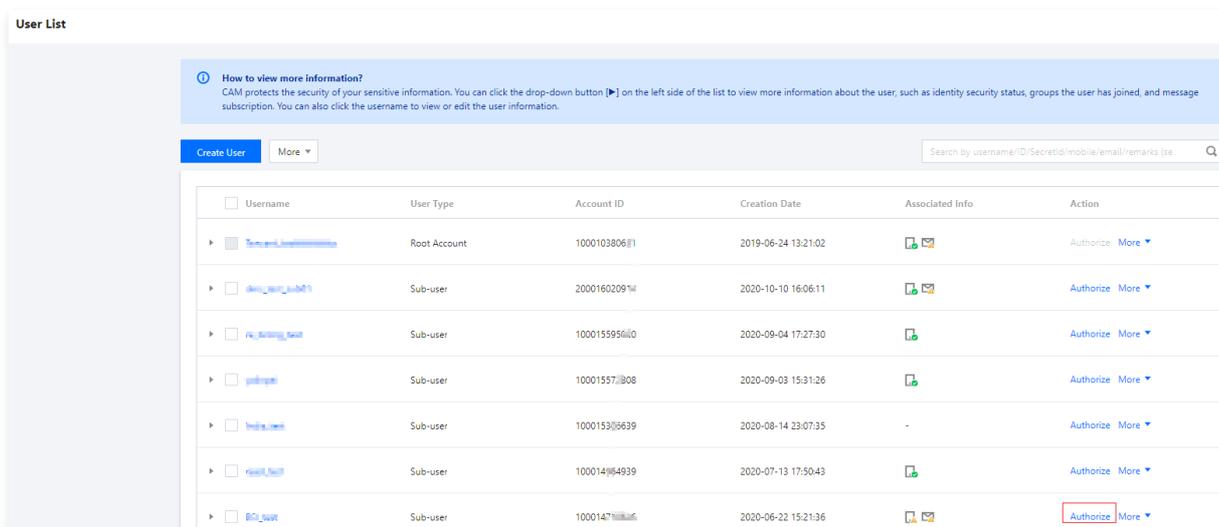
Note

For annual and monthly purchase methods, if financial permissions are not granted, a pending order will be generated and the link will be given to an account with financial permissions for approval. Pay-as-you-go purchase methods do not support order approval and must be granted financial permissions.

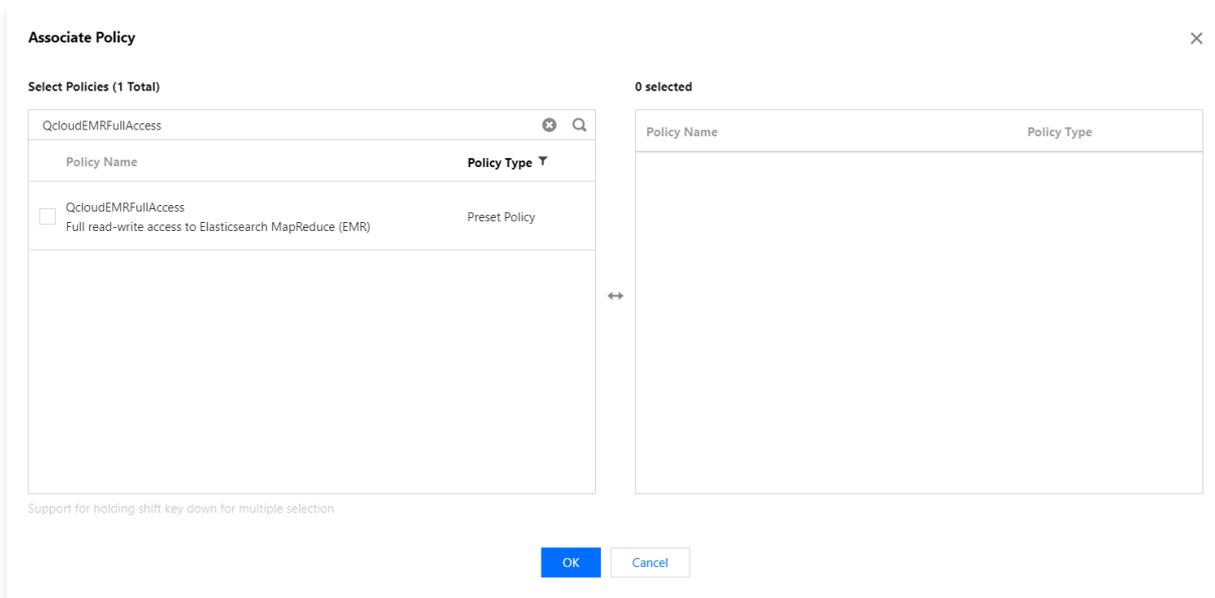
| Policy Category | Policy Name | Policy Description |
|-------------------|-------------------------|---|
| EMR Preset Policy | QcloudEMRFullAccess | Full Read/Write Access to EMR (Choose One) |
| EMR Preset Policy | QcloudEMRReadOnlyAccess | Read-Only Access to EMR (Choose One) |
| EMR Preset Policy | QcloudEMRPurchaseAccess | Financial Permissions for the Elastic MapReduce Product |

The primary account grants the above permissions to sub-users or collaborators, with the specific steps as follows:

1. Log in to the [Access Management Console](#), find the corresponding sub-user or collaborator in **Users > User List**, and click on **Authorize**.



2. In the associated policy, search for the policy in the above table (the QcloudEMRFullAccess policy is used as an example in the figure below). After selecting this policy, click **Confirm** to proceed.
3. Grant the EMR financial policy QcloudEMRPurchaseAccess, following the same steps as in Step 2.

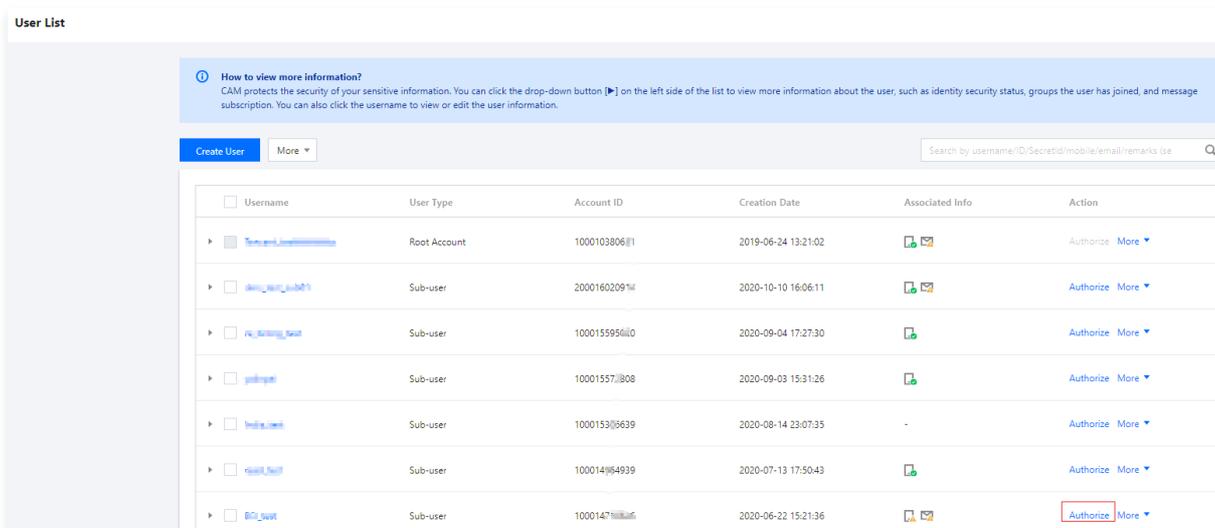


Cluster COS Service Role

EMR primary accounts, collaborators with QcloudCamRoleFullAccess permissions, and sub-users can precisely control COS bucket permissions and other cloud resource permissions. For more details, see [Cluster COS Service Role](#).

The primary account grants QcloudCamRoleFullAccess to sub-users or collaborators, with the specific steps as follows:

1. Log in to the [Access Management Console](#), find the corresponding sub-user or collaborator in **Users > User List**, and click on **Authorize**.



2. In the associated policy, search for the QcloudCamRoleFullAccess policy. After selecting this policy, click **Confirm** to proceed.

Associate Policy ✕

Select Policies (1 Total) 1 selected

qcloudEMRF ✕ 🔍

| Policy Name | Policy Type |
|---|---------------|
| <input checked="" type="checkbox"/> QcloudEMRFullAccess Full read-write access to Elasticsearch MapRe... | Preset Policy |

| Policy Name | Policy Type |
|--|--|
| QcloudEMRFullAccess Full read-write access to Elasticsearch Ma... | Preset Policy ✕ |

↔

Press Shift to select multiple items

Confirm Cancel

CAM-Enabled EMR API Authorization Granularity Details

Last updated: 2023-12-21 16:04:20

List of APIs supporting resource-level authorization

EMR supports resource-level authorization, allowing you to designate sub-accounts with interface permissions for specific resources.

Note

In the event of operational errors arising from the addition of subsequent interfaces, if permission errors are encountered, the corresponding interface permissions that are missing can be supplemented in the policy based on the error prompts.

APIs supporting resource-level authorization include:

| API Name | API Description | Resource Type/Resource Six-Part Format Example |
|------------------------------|--|---|
| DescribeKeyTabFile | Exporting Keytab File (User Management) | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeShowUserManagerTab | Displaying the User Management Tab Page (User Management) | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeResourceSchedule | Retrieving Data from the Yarn Resource Scheduling Page | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeCamUserList | Querying Cam User List | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeClusterServiceInfo | Inquiring Service Information | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeUserManagerUserList | Querying User List (User Management) | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ModifyUserManagerPwd | Modifying User Password (User Management) | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ModifyResourcePools | Refreshing Dynamic Resource Pool | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DeleteUserManagerUserList | Deleting User List (User Management) | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ModifyResourceScheduler | Modified the Yarn Resource Scheduler, click to deploy for changes to take effect. | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ModifyResourceScheduleSwitch | Upon activating the switch, it is necessary to synchronize the resource scheduler's configuration files to the resource scheduler in order to view the resource scheduler's page and perform related operations on the page. | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| AddUserManagerUserList | Adding New User List (User Management) | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ModifyUserManagerInfo | Updating User Information (User Management) | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ModifyResourceScheduleConfig | Modifying Resource Configuration of Yarn Resource Scheduler | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ModifyResourceScheduleCancel | Cancel Saving Resource Configuration | emr-instance qcs::emr:\${region}:uin/\${uin}:emr- |

| | | |
|--------------------------------|--|---|
| infigForRollback | of Yarn Resource Scheduler | instance/\$emrInstanceId |
| InquirePriceRefundEmr | Inquiry on Refund Amount for Node Termination | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeModifyPayModeNodes | Querying Resources for Cluster Conversion from Pay-as-you-go to Annual or Monthly Subscription | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ModifySecurityGroup | Modifying Cluster Security Group | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeInstanceRenewNodes | Querying Node Renewal for EMR Cluster | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| InquirePriceRenewEmr | Inquiry on Cluster Renewal Pricing | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeInstancesList | Querying EMR Cluster Instance List | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ModifyFlowStatus | Modifying Process Status | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ModifyMasterIp | Updating ERM Cluster IP | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| CheckFlowCanBeCancelled | Confirming if Process Supports Cancellation | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeModifySpec | Querying Configuration Specification Changes | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| AddShellScriptTask | Generating Cluster Script Tasks | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeShellScriptTaskList | Querying Cluster Script Task List | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeShellScriptNodes | Querying Node List of a Single Cluster Script Task | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeNodeList | Viewing Node Information | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeNodeList | Viewing Node Information | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeShellScriptNodeDetails | Querying Detailed Execution Results of a Single Node in a Cluster Script Task | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DeleteShellScriptList | Deleting Cluster Script Task Records | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeMasterIp | Retrieving ERM Cluster Instance IP | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeEKSIstances | Querying TKE Cluster Interface Information | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| AddServiceConfFile | Adding User-Defined Configuration Files | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DeleteServiceConfFile | Deleting User-Defined Configuration Files | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ModifyBootScript | Modifying Bootstrap Scripts | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeInstanceAlias | Retrieving Aliases | emr-instance qcs::emr:\${region}:uin/\${uin}:emr- |

| | | |
|-----------------------------|--|---|
| | | instance/\$emrInstanceId |
| DescribeBootScript | Retrieving Bootstrap Scripts | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeSubJobFlowStatus | Describing EMR Subtask Processes | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ListConfLogs | Retrieving Configuration Distribution Logs | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| GenerateScaleoutGoodsDetail | Generating Expansion Orders | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| GenerateRenewGoodsDetail | Generating Renewal Orders | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| GenerateModifyGoodsDetail | Generating Reconfiguration Orders | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ModifyAutoScaleGlobalConf | Updating Auto-Scaling Global Configuration | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeFlowStatusDetail | Querying the Detailed Status of EMR Task Execution | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeFileIps | Querying the IP List Where the Specified Configuration File is Located | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeExportConfsList | Querying the List of Supported Configuration Files for Export | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeFlowStatus | Querying the Process Status of EMR Instances | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeFlowNum | Querying the Number of EMR Cluster Processes | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeClusterNodes | Querying the Hardware Node Information of the EMR Cluster | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ClearMetadata | Purging EMR Cluster Metadata Information and Dismantling the Cluster | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeAutoScaleGlobalConf | Retrieving Global Configuration for Automatic Scaling | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeAutoScaleSpecs | Retrieving Specifications for Automatic Scaling | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ModifyAutoScaleSpecs | Modifying Specifications for Automatic Scaling | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| AddMetricScaleStrategy | Implementing Metric Load Scaling Rules | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DeleteAutoScaleSpec | Deleting Specifications for Automatic Scaling | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeAutoScaleRecords | Retrieving Historical Records of Automatic Scaling | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DeleteAutoScaleStrategy | Deleting Rules for Automatic Scaling | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ModifyStrategyPriority | Adjusting Rule Prioritization | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeAutoScaleMetaRan | Retrieving Metadata for Automatic | emr-instance qcs::emr:\${region}:uin/\${uin}:emr- |

| | | |
|--------------------------------------|---|---|
| ge | Scaling | instance/\$emrInstanceId |
| DescribeAutoScaleStrategies | Retrieving Rules for Automatic Scaling | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ModifyAutoScaleStrategy | Modifying Rules for Automatic Scaling | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| AddAutoScaleSpec | Adding Specifications for Automatic Scaling | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| TerminateAutoScaleNodes | Terminating All Nodes for Automatic Scaling | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeCbsEncrypt | Describing CBS Encryption | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| ModifyConfigGroup | Modify configuration group | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DeleteConfigGroup | Delete configuration group | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| AddConfigGroup | Add configuration group | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeConfigGroup | Describing Configuration Group | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| SynchronizeGroupConfCheck | Synchronize Configuration Check | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| UnbindInstanceAndNodesTags | Unbind Cluster Tag | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeNodeResourceConfigFast | Quickly Acquire Node Specification Configuration | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| InstallCos | Activate COS Subsequently | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeScaleoutableService | Describing Scalable Services | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeNodeResourceConfig | Obtain Node Specification Configuration | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DeleteNodeResourceConfig | Delete Node Specification Configuration | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| AddNodeResourceConfig | Increase Node Specification Configuration | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| SetNodeResourceConfigDefault | Establish Default Attributes for Node Specification Configuration | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeHbaseTableMetricDataOverview | Retrieve Interface Overview for HBase Table-Level Monitoring Data | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| DescribeInstanceNode | [Tag Console] Fetch Node Resources | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| UpdateWebproxyPassword | Update Proxy Component Password | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| TerminateTasks | Downscale Task Nodes | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |
| TerminateNodes | Decommission Node | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceId |

| | | |
|------------------------------|---|---|
| TerminateInstance | Decommission EMR Instance | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| StopService | Halt Component Services | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| StopMonitor | Cease Component Monitoring | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| StartService | Initiate Component Services | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| StartMonitor | Initiate Component Monitoring | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ScaleOutRouter | Expand Router Nodes | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeDestroyInfo | Inquire EMR Cluster/Node Termination Information | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ScaleOutInstance | Scaling an instance | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| RollBackConf | Roll back configuration | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| RestartService | Restart Component Services | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ModifyServiceParams | Modify Service Component Parameters | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ModifyResource | Reconfigure Instance | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| AssignInstancesProject | Assigning EMR Cluster to a Specific Project | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ModifyOptionalSpecStatus | Modify Alternate Specification Status | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ModifyOptionalSpec | Modify Alternate Specifications | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeSelectedOptionalSpec | Query Selected Alternate Specifications | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ModifyInstanceBasic | Renames cluster | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| InstallSoftware | Installing the add-on | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| InquiryPriceAddRouter | Retrieve the Price for Adding a Router | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| InquiryPriceRenewInstance | Inquiry for Renewal Pricing | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| InquiryPriceUpdateInstance | Inquiry for Configuration Changes Pricing | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| InquiryPriceScaleOutInstance | Inquiry for Expansion Pricing | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeClusterNodes | Querying the Hardware Node Information of the EMR Cluster | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |

| | | |
|-------------------------------|---|---|
| DescribeServiceNodeInfos | Retrieve Information on EMR Cluster Service Processes | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeServiceGroups | Retrieve Information on EMR Cluster Service Groups | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeServiceConfs | Retrieve All Configuration Information of EMR Cluster Services | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeScaleoutGoodsDetail | Retrieve Expansion Order Information | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeRouterGoodsDetail | Retrieve Router Node Order Information | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeRenewGoodsDetail | Retrieve Renewal Order Information | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeOptionalSpec | Retrieve Alternative Specifications for EMR Instance Nodes | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeModifyGoodsDetail | Retrieve Configuration Change Order Information | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeMetricsDimension | Retrieve Monitoring Dimension Values for Different Monitoring Levels | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeMetricMeta | Retrieve EMR Cluster Monitoring Metadata | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeInstanceOplog | Retrieve Operation Logs for EMR Instances | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeExecCustomScript | Retrieve User-Defined Script Information | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeInstanceAlerts | Retrieve EMR Cluster Alert Information | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeInstances | Retrieve EMR Instance Information | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeInstallSoftwareInfo | Retrieve Information on Installed and Uninstalled Components of EMR Cluster | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeExportConfs | Retrieve Export Configuration | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| CheckDiskInfo | Verify if Console Disk Metadata has been Updated | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| SyncDiskInfo | Console Disk Metadata Update | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeServiceConfsNew | Acquire Component Configuration Information (Configuration Management Page) | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeConfFileList | Acquire Configuration File List (Configuration Management Page) | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeServiceConfCategories | Acquire Component Configuration Categories (Configuration Management Page) | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeServiceConfDiff | Configuration Comparison (Configuration Management Page) | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeConfigGroupList | Inquire Configuration Group | emr-instance qcs::emr:\${region}:uin/\${uin}:emr- |

| | | |
|----------------------------------|--|---|
| | Information of Node Type | instance/\$emrInstanceid |
| DescribeHbaseStatus | DescribeHbaseStatus Information Display | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ModifyHbaseRit | Repair Hbase rit | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| TerminateSparkApp | Terminate Spark Job | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeAccessKey | Retrieve Cluster AccessKey | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeAttachableDisks | Inquire Cloud Disk Mounting Support for Nodes | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| InquirePriceRenewDisks | Cloud Disk Renewal Inquiry Pricing | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| AttachDisks | Cloud Disk Mounting | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeNodeDataDisks | Inquire Node Data Disk Information | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| InquirePriceResizeDisks | Cloud Disk Expansion Inquiry Pricing | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ResizeDataDisks | Scale out | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ModifyAutoRenewFlag | Auto-Renewal | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ScaleOutCluster | Cluster Node Expansion | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| StartStopServiceOrMonitor | Utilized for Initiating or Halting Monitoring or Services | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeClusterClients | Client Information Viewing Page | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeDelayedServiceConfig | Retrieving Expired Configurations (Configuration Management Page) | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeFailedServiceConfig | Retrieving Failed Configurations (Configuration Management Page) | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeServiceConfDimensionDiff | Retrieving Comprehensive Configuration Difference Comparison Results (Configuration Management Page) | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ModifyEmrManagerAgent | EMR Manager Agent Update | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ModifyServiceConfDiff | Differential Configuration Overwrite (Configuration Management) | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ModifyYarnLabelState | Issuing Commands for Adding, Deleting, and Binding Tags | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeYarnScheduleHistory | Viewing the Scheduling History of Yarn Resource Allocation | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeYarnDeployMessage | Deployment Activation - Notification Information | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |

| | | |
|--------------------------------|---|---|
| ModifyYarnDeploy | Apply | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeParentLabels | Retrieving the Tag List of the Parent Queue | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeYarnLastestLabels | Retrieving the Most Recent Tag Information | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ModifyYarnLabels | Synchronizing Yarn Node Tags | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ModifyYarnQueue | Modifying Resource Pool in Resource Scheduling | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ModifyOldLabelConfig | Cancel Saving the Edited Content of Yarn Tag Management | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeNodeLabelWebUrl | Obtaining the Web URL of Yarn's Node Label | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeNodeManagerHosts | Resource Scheduling – Tag Management – NodeManager IP List | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ModifyGlobalScaleConf | Cluster Scaling Configuration, Determining Whether to Enable Scaling and the Type of Scaling | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeGlobalScaleConf | Obtaining the Cluster's Scaling Configuration, Including Whether Scaling is Enabled and the Type of Scaling | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| ModifyManagedScaleGlobalConf | Updating the Managed Scaling Global Configuration | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeManagedScaleGlobalConf | Retrieving the Managed Scaling Global Configuration | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeServiceComponentInfos | Describing the Role Information of the Container Cluster | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DescribeRssClusterList | The Current Spark Cluster is Identical to the Rss Cluster of the Same EKS | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| CreateSparkLinkRss | Associating Spark with RSS | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |
| DeleteSparkLinkRss | Dissociating from RSS | emr-instance qcs::emr:\${region}:uin/\${uin}:emr-instance/\$emrInstanceid |

List of APIs Supporting Interface-Level Authorization

| API Name | API Description |
|-----------------------------|------------------------------------|
| RunJobFlow | Establishing Job Execution |
| DescribeJobFlow | Inquiring Job Execution |
| DescribeK8sResourcePrice | Inquiring K8s Resource Pricing |
| DescribePodSpecs | Depicting Pod Specifications |
| GenerateCreateGoodsDetail | Generating Cluster Creation Orders |
| DescribeLogSearchFileNames | Retrieving Log Search File List |
| DescribeInstanceServiceRole | Retrieving Service Role Table Name |

| Names | |
|-----------------------------------|--|
| DescribeCompareMetricsList | Retrieving Comparative Metrics List |
| DescribeHeatMapMetricList | Returning Aggregated Dimension Metrics List for Cluster Hosts |
| DescribeHBaseRegionList | Retrieving HBase Region List |
| DescribeClusterOverview | Querying Cluster Overview |
| DescribeEMRNodeOverview | Querying Deployment Status of Various Processes on Node |
| DescribeNodeHardwareInfo | Querying Basic Configuration of the Host |
| DescribeTopNMeta | Retrieving TopN Metadata Information |
| DescribeMetricDataAutoGranularity | Retrieving Monitoring Data with Automatic Time Granularity Setting |
| DescribeLogSearchTabs | Retrieving Log Search Tab Label Switch |
| DescribeEMRHostOverview | Querying Process Status on Host Detail Page |
| DescribeLogSearchRecords | Retrieving Log Search Content |
| DescribeTopNByProcess | Retrieving TopN Processes |
| DescribeDiskInfo | Returning Detailed Disk Information |
| DescribeTopNByHost | Querying TopN in Host Dimension on Overview Page |
| DescribeHeatMapDistribution | Returning Cluster Heatmap Data |
| DescribeInstanceServiceRoleTable | Retrieving Service Role Table Data |
| DescribeNodeAlias | Retrieving EMR Node Aliases |
| DescribeInstanceNodes | Retrieving Cluster Host Information |
| DescribeKeyPairs | Inquiring Key Pair Information |
| DescribeHbaseTableMetricData | Inquiring HBase Table-Level Monitoring Data |
| DescribeKeyPairs | Inquiring Key Pair Information |
| DescribeEmrMetaDB | Retrieving Hive Unified Metadata Information |
| ModifyEmrRole | Updating EmrRole Information |
| DescribeDisasterRecoverGroup | Retrieving Dispersed Placement Group Information |
| DescribeTags | Retrieving All Cluster Tags |
| InquiryPriceCreateInstance | Initiating Instance Pricing Inquiry |
| DescribeAccountBalance | Inquiring Account Balance |
| DescribeSecurityGroup | Retrieving EMR Security Group Information |
| DescribeCvmSpec | Inquiring Cloud Server Specifications |
| DescribeCdbPrice | Inquiring CDB Pricing |
| CreateInstance | Establishing EMR Instance |
| GetMetricDataForMcController | Console Detail Page Monitoring Information Interface |
| DescribeVpcList | Inquiring VPC List Information |

| | |
|--|---|
| DescribeSceneProductInfo | Acquiring Purchase Page Cluster Scenarios, Types, Products, and Component Information |
| DescribeRegionAndZoneSaleInfo | Acquiring Console and Purchase Page's Regions and Availability Zones |
| DescribeCgwProjects | Acquiring CGW Project List |
| DescribeServiceUpgradeVersion | Viewing Supported Service Minor Version Upgrades |
| ModifyServiceVersion | Service Minor Version Upgrade |
| DescribeServiceRoleInstanceConstraints | Retrieving Service Role Instance Removal Constraint Information |
| CheckSupportServiceRoleInstance | Determining if the Service Supports Addition of Role Instances |
| DescribeServiceRoleNames | Retrieving the Dropdown List of Role Names (or Role Types) for this Service |
| DescribeServiceRoleInstanceDeployableNodes | Querying the Deployment Node List of the Service Role |
| DescribeServiceDeletableRoleInstances | Querying the List of Role Instances that can be Removed from the Service |
| DeleteServiceRoleInstance | Remove role instance |
| AddServiceRoleInstance | Add role instance |
| ModifyResourcesTags | Enforce Tag Modification |
| CreateCluster | Creating a Cluster |
| RepairDisk | Upon completion of disk repair for customer's faulty disk on the CVM side, this interface can be utilized to finalize the business recovery. |
| CreateOutsideClient | For machines outside the EMR cluster, component packages and deployment scripts are provided to facilitate access to EMR and task submission. |
| ModifyCvmReplaceFlag | Modify the EMR cluster to support the attribute of abnormal node compensation switch. |

For detailed authorization schemes at the resource and interface levels, please refer to: [Granularity of Authorization Scheme](#).

Authorization Granularity Scheme

Last updated: 2023-12-21 16:04:31

Six-Segment resource description

qcs: An abbreviation for Qcloud Service, indicating that it is a cloud resource of Tencent Cloud. This field is mandatory.

project_id: Describes project information, compatible with early CAM logic, no need to fill in.

service_type: Product abbreviation, such as CVM.

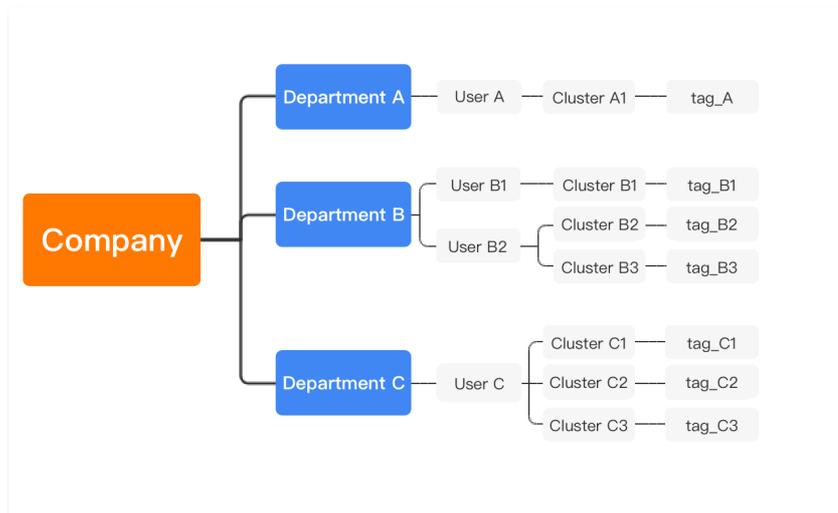
region: Regional information, such as bj (refer to [CAM documentation](#), if different, please provide the abbreviation of the regional list).

account: The main account information of the resource owner, such as uin/164256472.

resource: Specific resource details of each product, such as instance/instance_id1 or instance/*.

Guidelines for Tag Authentication Operations

The following text will introduce the operation process of EMR tag authentication. Before using tags to authenticate resources on EMR, you need to plan the tags for your company's departments or organizations; for example, in the figure below, plan the corresponding usernames, clusters, tags, etc. for departments A/B/C. For tag-related operations, please refer to: [Tags](#).

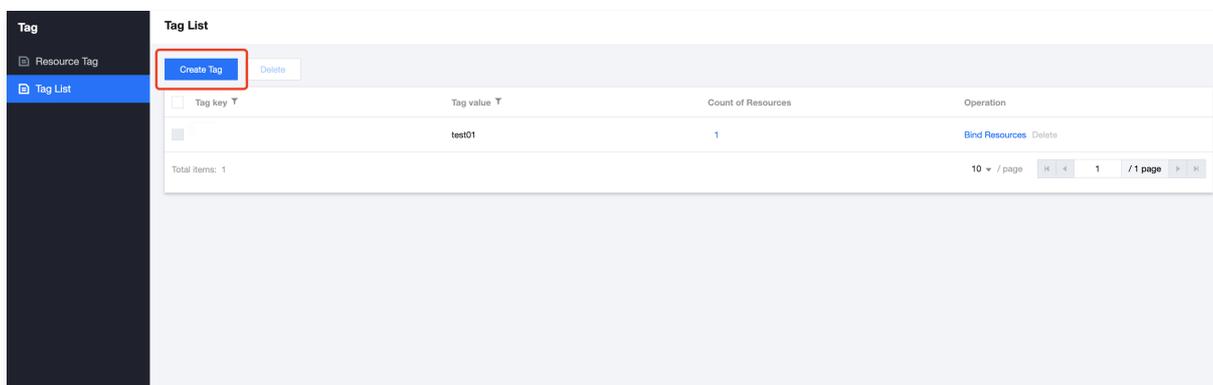


Upon completion of the planning, to enable Department A to manage Cluster A1, the following operations need to be executed in sequence.

1. Create a new tag and assign it to the cluster

- Navigate to the [Tag List](#) and click on **Create New Tag**.

Note
Here and in the subsequent processes, we will use tag_A as an example.



- Enter the tag key and tag value, then click **OK** to successfully create it.

Create Tag ✕

i • Enter a new tag key and tag value to create a new tag. Select an existing tag key to add a tag value to the key

• A tag key has up to 1000 tag values, and a single creation can enter up to 10 tag values

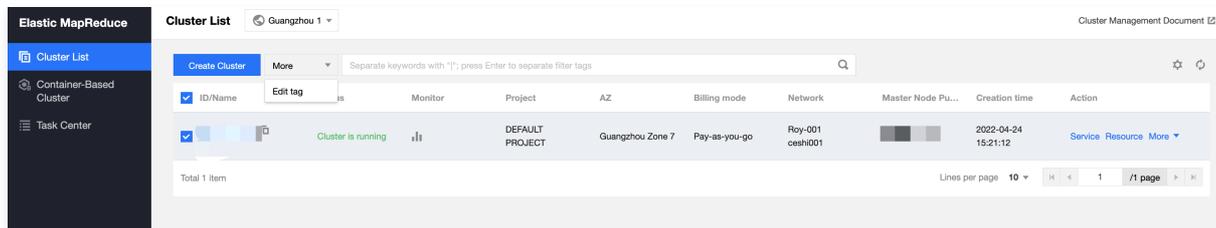
Tag key

Tag value

Delete

[Add tag key](#)

- On the [Cluster List](#) page, select the cluster you wish to tag, click on **More** above, then click on **Edit tag**.



- In the Edit Tags section, select the previously established tag_A, then click **OK**.

Edit Tags ✕

The tag is used to manage resources by category from different dimensions. If the existing tag does not meet your requirements, please go to [Manage Tags](#)

1 resource selected

✕

[+ Add](#)

2. Create a Custom Policy

Currently, some EMR interfaces do not support tag authentication, hence, it is necessary to set two policies. Policy_Dept_A1 is the tag policy, and Policy_Dept_A2 is the policy for interfaces that do not support tags.

- Create a Custom Policy Policy_Dept_A1**
 - Navigate to the [Access Management Policy](#) page, then click on **Create Custom Policy**.
 - Select **Authorize by Tag**.
 - In the Tag Policy Generator, assign the user to select Dept_A. Leave the user group blank. Choose the planned corresponding tags for the tag key and tag value. Add services and operations as per requirements, then click on **Next**.
 - Create a new Policy_Dept_A1 policy. In the policy content, modify the "action" and "resource" fields as follows, then click on

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"
      ]
    }
  }
}
```

Note

The content of the "resource" field should be filled in according to the actual resources. For the rules of filling in, refer to [Resource Description Method](#).

• Create a new custom policy, Policy_Dept_A2

- Navigate to the **Access Management** policy page, then click on **Create New Custom Policy**.
- Select **Create by Policy Syntax**.
- When selecting a policy template, choose the **Blank Template**, then click **Next**.
- Create a new Policy_Dept_A2 policy. After clearing the original content in the policy, paste the following content and click **Finish**.

```
{
  "version": "2.0",
  "statement": [
    {
      "effect": "allow",
      "action": [
        "emr:RunJobFlow","emr:DescribeJobFlow","emr:FindClusterListForMC","emr:EmrScaleoutClusterWithAuth","emr:CreateClusterForMC","emr:CheckSoftRelation","emr:DescribeK8sResourcePrice","emr:DescribePodSpecs","emr:DescribePodSpec","emr:DescribeTkeWhiteList","emr:GenerateCreateGoodsDetail","emr:DescribeLogSearchFileNames","emr:DescribeEMRInstancesExtra","emr:DescribeClusterHardwareInfo","emr:DescribeInstanceServiceRoleNames","emr:DescribeServiceComponents","emr:DescribeCompareMetricsList","emr:DescribeHeatMapMetricList","emr:DescribeHBaseRegionList","emr:DescribeAutoScaleWhiteList","emr:DescribeOptionalSpecWhiteList","emr:DescribeServiceOverview","emr:DescribeEMRInstances","emr:DescribeEMRNodeType","emr:DescribeEMRNodes","emr:DescribeStepByTimeRange","emr:DescribeClusterOverview","emr:DescribeEMRNodeOverview","emr:DescribeNodeHardwareInfo","emr:DescribeTopNMeta","emr:DescribeMetricDataAutoGranularity","emr:DescribeLogSearchTabs","emr:DescribeEMRHostOverview","emr:DescribeLogSearchRecords","emr:DescribeTopNByProcess","emr:DescribeDiskInfo","emr:DescribeMetricInfo","emr:DescribeInstanceServiceAbstract","emr:ModifyMetricMetaPerInstance","emr:DescribeOverviewData","emr:DescribeClusterHardwareType","emr:DescribeClusterEmergencyCallback","emr:DescribeTopNByHost","emr:DescribeHeatMapDistribution","emr:DescribeInstanceServiceRoleTable","emr:DescribeEmrSaleInfoExtend","emr:DescribeNodeAlias","emr:DescribeInstanceNodes","emr:DescribeKeyPairs","emr:DescribeHbaseTableMetricData","emr:DescribeHbaseMetricMeta","emr:DescribeEmrMetaDB","emr:CheckMetaDBNet","emr:ModifyEmrRole","emr:DescribeEmrRole","emr:DescribeDisasterRecoverGroup","emr:DescribeTags","emr:InquiryPriceCreateInstance","emr:GetCreateGoodsDetail","emr:DescribeAccountBalance","emr:DescribeSecurityGroup","emr:DescribeMetricData","emr:DescribeEmrSaleInfo","emr:DescribeCvmSpec","emr:DescribeCdbPrice","emr:CreateInstance","emr:CheckCustomConfig","emr:CheckCosInfo","emr:GetCreateGoodsDetailList","emr:GetScaleoutGoodsDetailList","emr:GetRenewGoodsDetailList","emr:UpdateWebProxyPassForMcController","emr:SubmitServiceParamsForMC","emr:ReleaseClusterForMC","emr:RestartServiceForMC","emr:GetNodeListForMcController","emr:GetMetricDataForMcController","emr:GetServiceGroupForMcController","emr:EmrDestroyTaskNodeWithAuth",
        "emr:DescribeModifySpec","emr:DescribeMasterIp","emr:DescribeMetricsDimension"
      ],
    }
  ],
}
```

```

    "resource": "*"
  }
]
}

```

3. Assign and verify sub-user permissions.

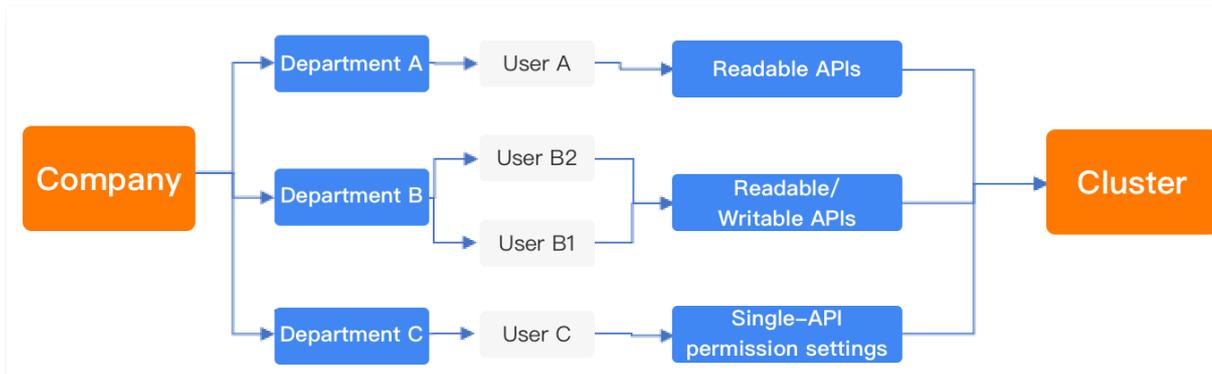
- On the user list page, locate the sub-user to be authorized and click the **Authorize** button on the right; for creating a new sub-user, please refer to: [Create New Sub-User](#).
- Select Policy_Dept_A1 and Policy_Dept_A2, then click **Finish**.
- Log in as a sub-user for verification.
At this point, when logging into EMR, the sub-user can only see and manage the test_A cluster through the tag_A.

△ Note

As the interfaces in the Policy_Dept_A2 policy can only be listed in full for the time being, it is not ruled out that operational errors may occur when new interfaces are added. If you encounter permission errors, please [submit a ticket](#) or [contact us](#).

Guidelines for Resource-Level and Interface-Level Authentication Operations.

The following will introduce the operation process of EMR resource-level authentication. When using resource-level authentication for a specific function of EMR, you need to first plan for your company's departments or organizations; for example, in the figure below, plan the corresponding usernames, interfaces, clusters, etc. for departments A/B/C.



After the planning is complete, to ensure that Department A only has access to the DescribeRegionAndZoneSaleInfo interface of the cluster, the following operations need to be completed in order:

1. Create a New Custom Policy

- Navigate to the [Access Management Policy](#) page, then click on **Create Custom Policy**.
- Select **Create by Policy Syntax**.
- When selecting a policy template, choose the **Blank Template**, then click **Next**.
- Create a new custom policy named DescribeRegionAndZoneSaleInfo-test. After clearing the original content in the policy, configure the access permission deny policy for the DescribeRegionAndZoneSaleInfo interface in EMR within the policy. Once created, save it.

```

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

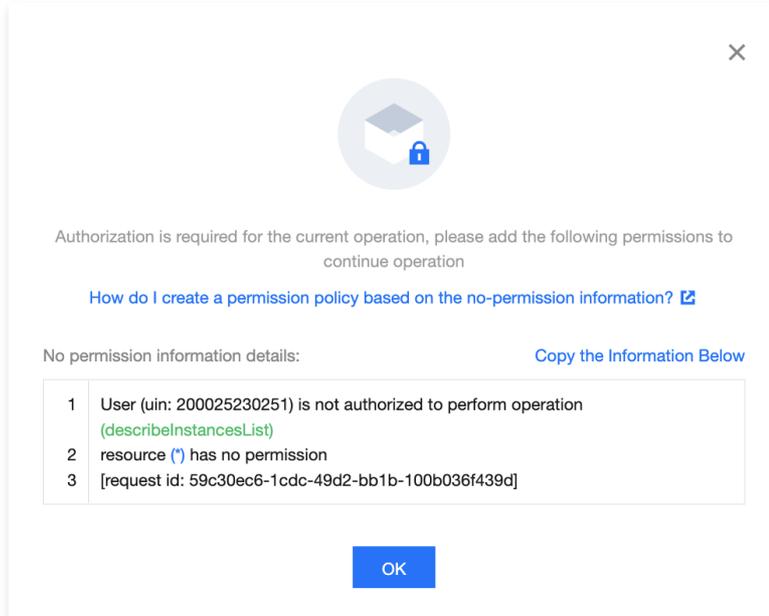
```

△ Note

For interfaces that support resource-level authentication, please refer to the **API list that supports resource-level authentication**. As the interfaces can only be listed in full for the time being, it is not ruled out that operational errors may occur when new interfaces are added. If you encounter permission errors, you can supplement the policy based on the error prompt indicating a lack of corresponding interface permissions.

2. Assign and Verify Sub-Account Permissions

- Locate the sub-user to be authorized on the user list page and click the **Authorize** button on the right; for detailed operations on creating a new sub-user, please refer to [Creating a New Sub-user](#).
- Select the custom policy name DescribeRegionAndZoneSaleInfo-test and click **Finish** to complete the process.
- After the sub-account logs in, if a pop-up window similar to the one below appears when the above interface is called, indicating that authorization is required, it means that the authentication has been successful.



Cluster COS Service Role

Last updated: 2023-12-21 16:04:39

The Cluster COS Service Role refers to the ability of users to add a CAM service role for accessing Cloud Object Storage (COS) resource permissions. The selected service type is "Tencent Cloud Product Service", supporting the role selection of "Elastic MapReduce". If no role is set, the system defaults to using the **EMR_QCSLinkedRoleInApplicationDataAccess service-related role** for accessing Cloud Object Storage (COS) resources.

The following steps outline the process of creating a custom service role and associating it with the Cluster COS Service Role.

Note:

1. The Cluster COS Service Role by default displays the currently bound role identity, and the current cluster uses this role identity to read and write to COS resources.
2. When both the EMR service-related role and the service role are authorized simultaneously, the newly created cluster defaults to using the service-related role for COS access. In the cluster instance information authorization policy, the object storage is shown as authorized, and the Cluster COS Service Role displays the **EMR_QCSLinkedRoleInApplicationDataAccess service-related role**.
3. When both the EMR service-related role and the service role are authorized, and the service-related role is not authorized through the EMR on CVM console cluster instance information Cluster COS Service Role, the default is to use the service role EMR_QCSRole and bind the QcloudAccessForEMRRoleInApplicationDataAccess identity for COS access. If the service-related role is required, manually bind the EMR_QCSLinkedRoleInApplicationDataAccess service-related role in the cluster instance information Cluster COS Service Role.

Step 1: Customize the Permission Policy

1. Log in to the [Access Management Console](#), click on **Create New Custom Policy**, and in the pop-up "Select Policy Creation Method" page, choose **Create by Policy Syntax**.
2. On the "Create by Policy Syntax" page, select **Template Type** as **Blank Template**.
3. The syntax policy is set as follows:

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

Wherein, appld is the primary account AppID, and bucketName is the name of the bucket you wish to authorize. Generate a policy named **TestPolicy**, customers can customize this name.

Step 2: Create a Custom Role

1. In the [Access Management Console](#), click on **Create New Role**. In the pop-up "Select Role Carrier" page, choose **Tencent Cloud Product Service**. Proceed to the "Create New Custom Role" page, and select **Product Service** as "Elastic MapReduce (emr)".
2. Bind the policy generated in Step 1, here it is **TestPolicy**. Customers can bind according to the policy they wish to authorize.
3. Mark the tag key and tag value of the role, then click **Next**.
4. Generate a custom role named **EMRCosRole**. Customers have the option to customize this name.

Step 3: Bind the Role to the EMR Cluster

In the [EMR Console](#), select the corresponding cluster and click on **Cluster ID/Name** to enter the instance details. In **Instance info > Basic configuration > Cluster COS role**, click on **Settings**. Set the "Custom Service Role" to the custom role generated in Step 2, here it is **EMRCosRole**. Users have the option to customize this name.

The screenshot displays the 'Instance info' page for an Elastic MapReduce cluster. The page is divided into several sections:

- Basic configuration:** Instance ID, Region info (Region: Beijing, AZ: Beijing Zone 7), Network info (Default-VPC), Security group (blue security group), Creation time (2023-11-23 10:53:05), Master public IP, Billing mode (Pay-as-you-go), Auto replenishment (toggle), Host login (Password Set), and a 'Check for updates' link.
- Authorizations:** Account COS access (Authorized) and Cluster COS role (EMR_QCCLinkedRoleInApplicationDataAccess Set).
- Software info:** Cluster type (Hadoop), Use cases (Hadoop-Default), Product version (EMR-V3.5.0), Components to deploy, MetaDB, Hive metadatabase, and Kerberos mode (Kerberos authentication enabled).
- Deployment info:** High availability cluster (Enabled), Cloud disk encryption (Not enabled), Beijing Zone 7, Subnet (Default-Subnet), and Node (Common x3, Master x2, Core x3, Task x0, Router x0).

Note:

1. Please ensure that the custom service role bound in the cluster COS service role includes specific COS operation permissions. If the bound role does not include specific COS operation permissions, access to COS will be impossible.
2. Query the custom service role or modify the preset permission policy in the custom service role's [Access Management](#).

Setting Tags

Last updated: 2023-12-21 16:05:31

Feature Overview

The **Tag** is a key-value pair (Key-Value) provided by Tencent Cloud Elastic MapReduce for identifying cluster types or node resources. You can use tags to categorize cluster or node resource types based on various dimensions (such as business, usage, person in charge, etc.). Tags also provide a convenient way to identify cluster or node resources. The key-value pairs of tags have no semantic meaning to Elastic MapReduce and are strictly parsed and matched as strings.

Usage Limits

A tag is a key-value pair (Key-Value) that allows you to categorize and manage cluster and node resources in Elastic MapReduce by setting tags. Tags provide a convenient way to view corresponding clusters and node resources. You can edit tags for cluster or node resources through the Elastic MapReduce console.

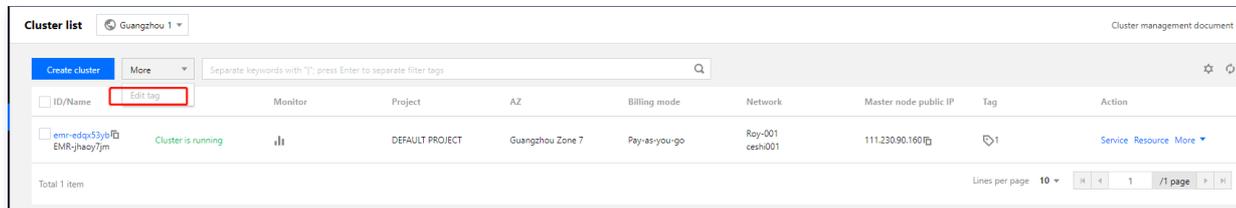
There are several limits on tags:

- **Quantity Limit:** The maximum number of tags allowed for each cluster or node is 50 (**up to 5 can be added at a time**).
- **Tag key limitations:**
 - The prefixes "qcloud", "tencent", and "project" are reserved for system tag keys and their creation is prohibited.
 - Tags can only consist of numbers, letters, and the symbols +=.@-, with a maximum tag key length of 255 characters.
- **Tag Value Limit:** Tags can only be empty strings, numbers, letters, or the symbols +=.@-, with a maximum tag value length of 127 characters.

Instructions

Editing Tags for Clusters

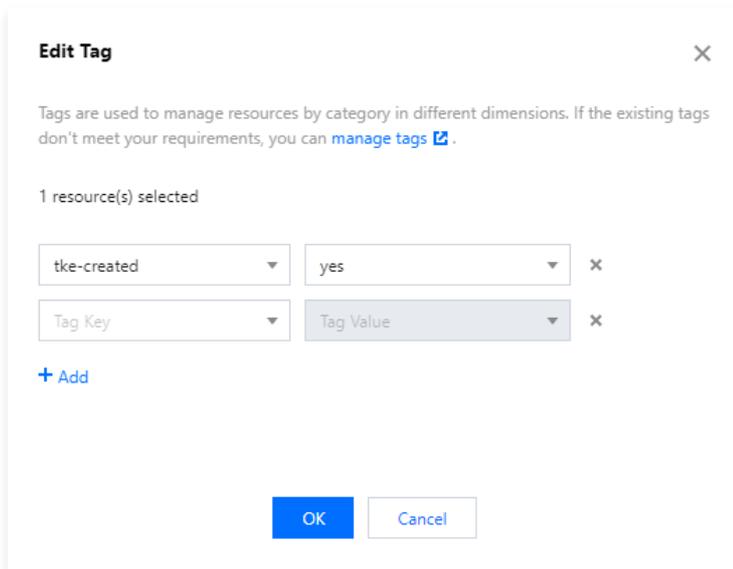
1. Log in to the [Elastic MapReduce Console](#). In the cluster list page, select the cluster for which you need to edit tags, then click on **More > Edit tag** at the top of the page, as shown in the image below:



2. In the pop-up window titled "You have selected 2 cloud resources", add, modify, or delete tags according to your actual needs.

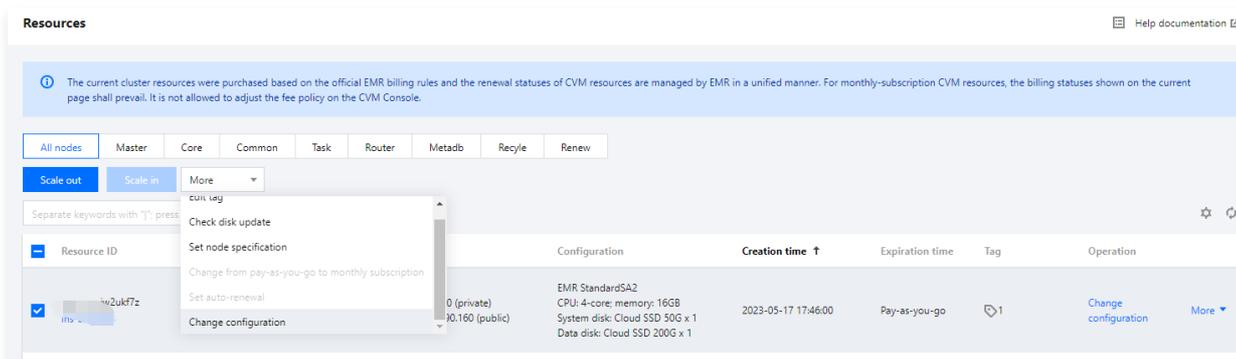
Note

A maximum of 20 clusters can be edited for tag operations in bulk.



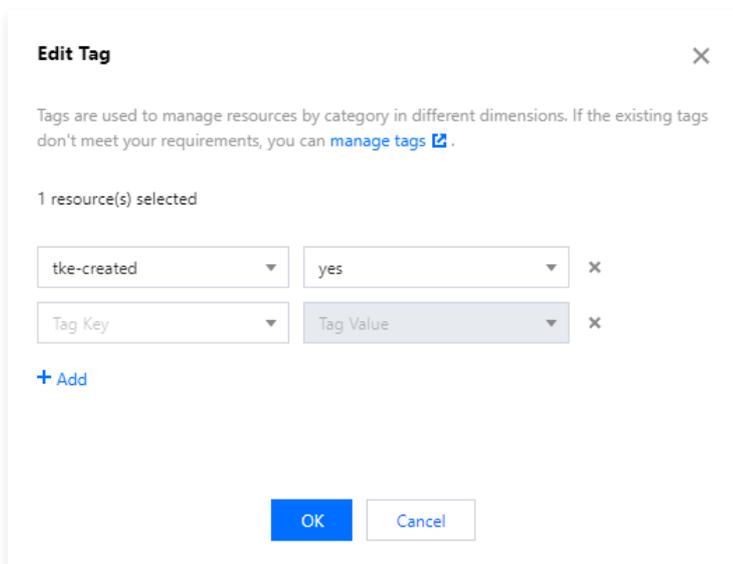
Editing Tags for Nodes

1. Log in to the [Elastic MapReduce Console](#) and click on the cluster ID/Name to access the cluster details page. In the cluster details page, select **Cluster Resources > Resources**, then choose the node resources for which you need to edit tags, and select **More > Edit tag**, as depicted in the image below:



2. In the pop-up window titled "You have selected 2 cloud resources", add, modify, or delete tags according to your actual needs.

Note
A maximum of 20 clusters can be edited for tag operations in bulk.



Software Configuration

Last updated: 2023-12-21 16:06:50

Feature Description

The software configuration allows you to customize the settings of components such as HDFS, YARN, Hive, etc., during the creation of your cluster.

Custom Software Configuration

Software such as Hadoop and Hive contain a multitude of configurations. With the software configuration feature, you can independently configure component parameters during the process of creating a new cluster. The configuration process requires you to provide the corresponding JSON file as per the requirements. This file can be customized by you, or you can export the software configuration parameters of an existing cluster, and then quickly create a new one. For details on exporting software configuration parameters, please refer to [Export Software Configuration](#).

JSON File Example and Explanation:

```
[
  {
    "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-stylesheet type='text/xsl' href='configuration.xsl'?>\n<configuration>\n  <property>\n    <name>yarn.scheduler.capacity.maximum-am-resource-percent</name>\n    <value>0.8</value>\n  </property>\n  <property>\n    <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    <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>"
    }
  }
]
```

Configuration Parameter Explanation:

- The 'serviceName' is the component name, which must be in uppercase.
- The 'classification' is the filename, which must be used in its entirety, including the suffix.
- The 'serviceVersion' is the version name of the component. This version must be consistent with the component version corresponding to the EMR product version.
- In 'properties', fill in the parameters that need to be configured independently.

- If you need to modify the configuration parameters in `capacity-scheduler.xml` or `fair-scheduler.xml`, the property key in 'properties' should be specified as 'content', and the value should be the entire content of the file.

If you need to adjust the configuration of existing clusters, you can proceed with the [Configuration Distribution](#).

Accessing External Clusters

After configuring the access address information for the external cluster's HDFS, you can read the data from the external cluster.

Configuration at Purchase

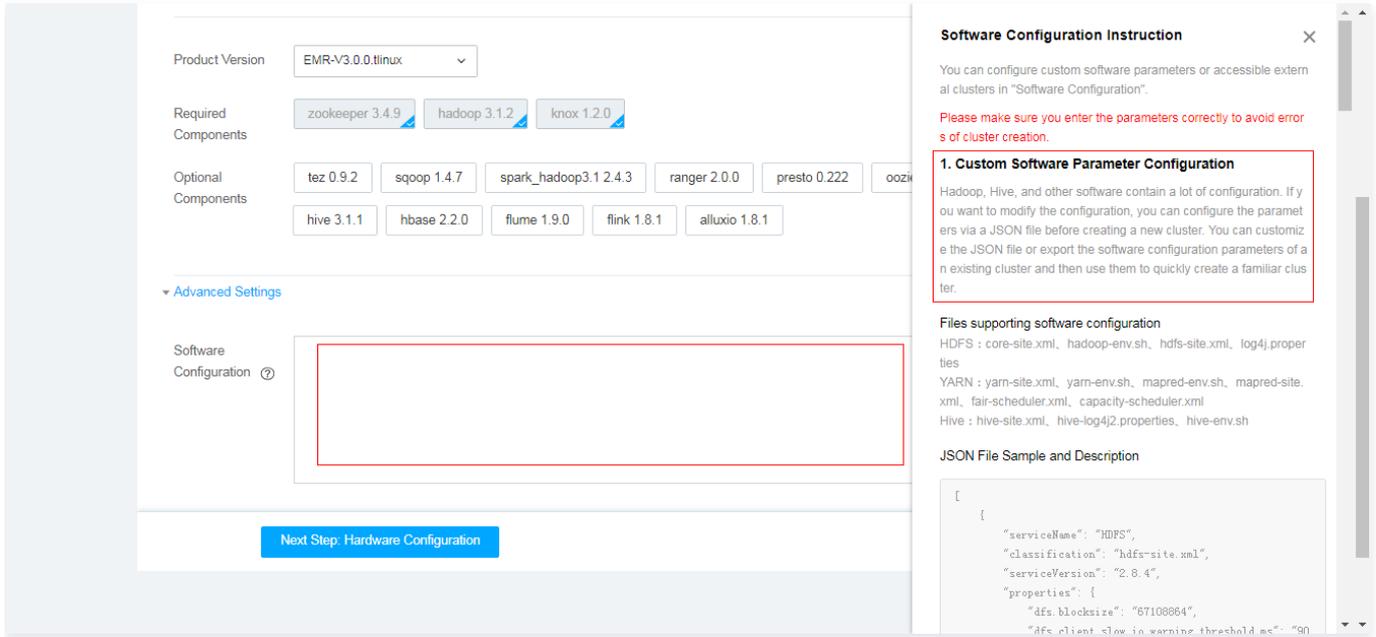
EMR supports configuring access to external clusters when creating a new cluster. This can be done by entering a compliant JSON file in the software configuration section on the [Purchase Page](#). The following example illustrates this under hypothetical conditions:

Assumed Conditions

Assume that the nameservice required to access the external cluster is HDFS8088, and its 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>
<property>
  <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>
```

If you need to access an external cluster from a newly created cluster, navigate to the [Purchase Page](#) and open the advanced settings.



JSON File and Explanation:

Using the assumed conditions as an example, the JSON file content should be filled in the box (the JSON content requirements are the same as the custom software configuration).

```
[
  {
    "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.rpc-address.HDFS8088.nn2": "172.21.16.40:4007"
    }
  }
]
```

Configuration Parameter Explanation:

- The serviceName component name must be "HDFS".
- The classification filename must be "hdfs-site.xml".
- The 'serviceVersion' is the version name of the component. This version must be consistent with the component version corresponding to the EMR product version.
- The content filled in the properties is consistent with the assumed conditions.
- newNameServiceName (optional) represents the nameservice of the newly created cluster. If left blank, it will be generated by the system; if not blank, it can only consist of a combination of strings, numbers, and hyphens.

Note
 The external cluster being accessed only supports highly available clusters.
 The external cluster being accessed only supports clusters that have not enabled kerberos.

Configuration after purchase

After the creation of the EMR cluster, it supports accessing the external cluster through the [Configuration Distribution](#) feature of EMR.

Assume the following conditions:

Assume the nameservice of this cluster is HDFS80238 (if it is a non-highly available cluster, it is generally masterlp:rpcport, for example, 172.21.0.11:4007).

The nameservice of the external cluster that needs to be accessed is HDFS8088, and its access method is:

```
<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>
<property>
  <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>
```

If this information is in the EMR cluster, it can be viewed on the [Configuration Distribution](#) management page, or by logging into the machine to view the `/usr/local/service/hadoop/etc/hadoop/hdfs-site.xml` file.

1. Enter the [Configuration Distribution](#) page, select the `hdfs-site.xml` file of the `hdfs` component.
2. Modify the configuration item `dfs.nameservices` to `HDFS80238,HDFS8088`.
3. Add configuration item and value

| Configuration items | Description |
|--|--|
| <code>dfs.ha.namenodes.HDFS8088</code> | <code>nn1, nn2</code> |
| <code>fs.namenode.http-address.HDFS8088.nn1</code> | <code>172.21.16.11:4008</code> |
| <code>dfs.namenode.https-address.HDFS8088.nn1</code> | <code>172.21.16.11:4009</code> |
| <code>dfs.namenode.rpc-address.HDFS8088.nn1</code> | <code>172.21.16.11:4007</code> |
| <code>fs.namenode.http-address.HDFS8088.nn2</code> | <code>172.21.16.40:4008</code> |
| <code>dfs.namenode.https-address.HDFS8088.nn2</code> | <code>172.21.16.40:4009</code> |
| <code>dfs.namenode.rpc-address.HDFS8088.nn2</code> | <code>172.21.16.40:4007</code> |
| <code>dfs.client.failover.proxy.provider.HDFS8088</code> | <code>org.apache.hadoop.hdfs.server.namenode.ha.ConfiguredFailoverProxyProvider</code> |
| <code>dfs.internal.nameservices</code> | <code>HDFS80238</code> |

Note

The `dfs.internal.nameservice` needs to be added, otherwise, after expanding the cluster, it may cause the datanode to report an exception and be marked as dead by the namenode.

4. Distribute the configuration using the [Configuration Distribution](#) feature.

For more detailed information and principles related to the configuration, please refer to the [Community Documentation](#).

Mounting CHDFS Instance

Last updated: 2023-12-21 16:07:06

Scenario

CHDFS is a high-performance distributed file system with standard HDFS access protocol and hierarchical namespace. EMR supports reading and writing data on CHDFS. This document primarily elucidates how to mount CHDFS to the EMR cluster.

Instructions

Scenario One: Mounting CHDFS on a New Cluster

Note

New Cluster: For clusters created on and after December 31, 2019, EMR defaults the CHDFS mount address to `/data/emr/hdfs/tmp/chdfs`.

The EMR cluster has been automatically adapted to CHDFS. Create CHDFS and set permissions appropriately to facilitate network communication between CHDFS and the EMR cluster. The configuration steps are as follows:

1. Create a CHDFS in the same region as the EMR cluster. Refer to [Creating CHDFS](#) for guidance.
2. Create a permission group as needed. Refer to [Creating a Permission Group](#) for guidance.
3. Create permission rules as needed. Refer to [Creating Permission Rules](#) for guidance.
4. Create a mount point under the same network as the EMR cluster. Refer to [Creating a Mount Point](#) for guidance.
5. Check the connectivity between CHDFS and the EMR cluster using the `hadoop fs` command line tool. Run the `hadoop fs -ls ofs://{mountpoint}/` command, where `mountpoint` is the mount address. If the file list is displayed correctly, this indicates that CHDFS has been successfully mounted.

Scenario Two: Mounting CHDFS on an Existing Cluster

Note

Existing Cluster: A cluster that was created before December 31, 2019.

For mounting CHDFS on an existing EMR cluster, refer to [Mounting CHDFS](#).

Hive Metadata Database

Last updated: 2024-01-12 11:12:27

Feature Overview

When deploying the optional component Hive in a newly created EMR cluster, the system provides two methods of Hive metadata storage, thereby achieving unified management of Hive metadata. The first method is the cluster default, where Hive metadata is stored in the MetaDB purchased independently by the cluster. The second method involves linking to an external Hive metadata database, where you can choose to link to EMR-MetaDB or a self-built MySQL database. The metadata will be stored in the linked database and will not be destroyed with the cluster.

Preparations

- **Cluster Default:** This involves the independent automatic purchase of a MetaDB cloud database instance storage unit as the metadata storage location. It stores metadata along with other component metadata and is destroyed along with the MetaDB cloud database when the cluster is destroyed. If you need to preserve metadata, you must manually save the metadata in the cloud database in advance.
- Hive metadata is stored together with the metadata of Hue, Ranger, Oozie, Presto, Druid, and Superset components.
- The cluster requires the separate purchase of a MetaDB as a metadata storage unit.
- The MetaDB is destroyed along with the cluster, meaning that the metadata is also destroyed with the cluster.
- **Associated EMR-MetaDB:** During the creation of the cluster, the system will fetch the available MetaDB from the cloud for the new cluster's Hive component to store metadata, eliminating the need for a separate purchase of MetaDB and saving costs. Moreover, the Hive metadata will not be destroyed with the current cluster.
- The available MetaDB instance ID corresponds to an existing MetaDB within the EMR cluster under the same account.
- When one or more components such as Hue, Ranger, Oozie, Druid, or Superset are selected, the system will automatically purchase a MetaDB for the storage of component metadata, excluding Hive.
- To destroy the associated EMR-MetaDB, one must proceed to the cloud database for destruction. Once destroyed, the Hive metadata database will be irretrievable.
- It is necessary to ensure that the associated EMR-MetaDB network and the newly created cluster exist within the same network environment.
- **Associated with a self-built MySQL database:** Associating your locally self-built MySQL database as Hive metadata storage also eliminates the need for a separate purchase of MetaDB, thus saving costs. It is necessary to accurately fill in the local address starting with "jdbc:mysql://", the database name, and the database login password, and ensure that the network is connected with the current cluster network.
- Please ensure that the self-built database and the EMR cluster are within the same network.
- Accurately fill in the database username and password.
- When one or more components such as Hue, Ranger, Oozie, Druid, or Superset are selected, the system will automatically purchase a MetaDB for metadata storage, excluding Hive.
- It is essential to ensure that the Hive metadata version in the custom database is greater than or equal to the Hive version in the new cluster.

Instructions

Create cluster

1. Log in to your [Tencent Cloud account](#), click on [Purchase Now](#), and in the **Optional Components** section of the Available Zone and Software Configuration page, select the Hive component.
2. For Hive metadata storage, you can choose based on your needs, with the default options being EMR-MetaDB or a self-built MySQL database.
3. Configure according to your selection and the aforementioned restrictions.

Install the HIVE component afterwards.

1. After the successful creation of the cluster, log in to the [EMR Console](#), enter the **Cluster List Page**, and click on the **Cluster ID/Name** you wish to manage.
2. Select **Add Component** in the **Cluster Services** and install 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

| | | | |
|---|--|--|---|
| <input checked="" type="checkbox"/> zookeeper-3.5.5 | <input checked="" type="checkbox"/> hadoop-2.8.5 | <input checked="" type="checkbox"/> Knox-1.2.0 | <input type="checkbox"/> zeppelin-0.8.2 |
| <input type="checkbox"/> livy-0.7.0 | <input type="checkbox"/> hbase-1.4.9 | <input checked="" type="checkbox"/> hive-2.3.5 | <input type="checkbox"/> hue-4.6.0 |
| <input checked="" type="checkbox"/> oozie-5.1.0 | <input type="checkbox"/> presto-0.228 | <input type="checkbox"/> ranger-1.2.0 | <input type="checkbox"/> spark_hadoop2.8-2.4... |
| <input type="checkbox"/> storm-1.2.3 | <input type="checkbox"/> sqoop-1.4.7 | <input type="checkbox"/> tez-0.9.2 | <input type="checkbox"/> ganglia-3.7.2 |
| <input type="checkbox"/> flume-1.9.0 | <input type="checkbox"/> impala-2.10.0 | <input type="checkbox"/> alluxio-1.8.1 | <input type="checkbox"/> flink-1.9.2 |
| <input type="checkbox"/> kyllin-2.5.2 | <input type="checkbox"/> superset-0.35.2 | <input type="checkbox"/> hudi-0.5.1 | |

Hive Metadatabase: Default Associate EMR-MetaDB Associate self-built MYSQL

Instance ID: Select data ▼

MetaDB [?]: Pay-as-you-go 1 high IO TencentDB with 4,000 MB memory and 100 GB disk

Cost:

Confirm
Cancel

3. For Hive metadata storage, you can choose based on your needs, with the default options being EMR-MetaDB or a self-built MySQL database.
4. Configure according to your selection and the aforementioned restrictions.

Security Group Configuration

Last updated: 2023-12-21 16:07:31

EMR employs Tencent Cloud's Private Network (VPC) as its foundational network, utilizing security group settings within EMR to regulate internal node intercommunication and external node access to internal node-side virtual firewalls. This document primarily elucidates the best practices for employing security groups within EMR, assisting you in selecting an appropriate security group strategy.

Security Group

A security group functions as a stateful packet-filtering virtual firewall, employed to establish network access control for single or multiple cloud servers (nodes), serving as a crucial means of network security isolation provided by Tencent Cloud. In the absence of a security group, the creation of an EMR cluster will automatically generate a new security group for you. If the number of security groups has reached its limit and no new ones can be created, it is possible to delete some of the unused security groups. If a security group already exists, it can be directly viewed and utilized via the Private Network Console.

Usage Limitations and Regulations

For information regarding the usage limitations and quotas of security groups, please refer to the relevant restrictions on security groups in the Overview of Usage Limitations.

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, UDP, etc.
- Policy: allow or reject the access request.

Principles for Selecting EMR Security Groups

By default, an existing security group is selected, specifically the EMR security group. However, users have the option to create a new EMR security group or select a non-EMR security group.

1. A newly created EMR security group will open ports 22 and 30001, as well as the necessary internal communication network segments. The new security group will be named in the format `emr-xxxxxxx_yyyyMMdd`. Please refrain from manually altering the security group's name.
2. Select an existing security group as the security group for the current instance, supporting all available security groups in the current region. It is recommended to prioritize security groups beginning with `emr-xxx`, as these groups have already enabled the necessary policies for the normal operation of EMR services. Security groups not beginning with `emr` may lack the necessary inbound and outbound rules, leading to cluster creation failure or cluster unavailability. Please exercise caution when selecting security groups not beginning with `emr`.
3. When expanding nodes, the security group by default integrates the security group policy selected during the creation of the new cluster.

Details of EMR Security Group Policies

When creating a new EMR cluster using a non-EMR security group, the following inbound and outbound rules must be included, otherwise the cluster cannot be created.

Inbound rules

| Source | Protocol Port | Policy | Note |
|----------------|---------------|--------|---------------------------|
| 10.0.0.0/8 | ALL | ACCEPT | Unblock A network segment |
| 172.16.0.0/12 | ALL | ACCEPT | Unblock B network segment |
| 192.168.0.0/16 | ALL | ACCEPT | Unblock C network segment |
| 0.0.0.0/0 | ICMP | ACCEPT | Unblock local ICMP |

Outbound rules

| Source | Protocol Port | Policy | Note |
|--------|---------------|--------|------|
| | | | |

| | | | |
|-----------|-----|--------|------------------------------|
| 0.0.0.0/0 | ALL | ACCEPT | Unblock all outbound traffic |
|-----------|-----|--------|------------------------------|

Inbound rules for accessing the webUI

When utilizing a non-EMR security group, to ensure normal access to the cluster service webUI, the inbound rules should encompass the following policies:

| Source | Protocol Port | Policy | Note |
|-----------|---------------|--------|-------------------------------|
| 0.0.0.0/0 | TCP:13000 | ACCEPT | Port 13000, the hue port |
| 0.0.0.0/0 | TCP:30001 | ACCEPT | Unblock port 30001 |
| 0.0.0.0/0 | TCP:30002 | ACCEPT | Unblock port 30002 |
| 0.0.0.0/0 | TCP:22 | ACCEPT | Unblock the remote login port |

For a more comprehensive introduction to security groups, please refer to [Security Groups](#).

Cluster Management

Instance Information

Last updated: 2023-12-21 16:09:54

Feature Overview

Instance information is a record of the fundamental details of a user's EMR cluster. Users can view the basic configuration, software information, and deployment details of the cluster on the instance information page.

- The basic configuration displays the fundamental details of the cluster, such as network information, creation time, security group, master public IP, billing mode, host login method, whether an automatic service role has been set, and if object storage (COS) has been enabled.
- The software information displays the cluster type, application scenario, product version, deployment components, MetaDB, Hive metadata database, and whether the Kerberos mode has been enabled.
- The deployment information displays the quantity details of each node type.

This document provides guidance on how to view cluster instance information via the console.

Instructions

1. Upon successful creation of the cluster, log in to the [EMR Console](#) and click on the **ID/Name** of the cluster you wish to manage from the cluster list page.
2. If the current cluster has not enabled object storage, you can click on **Authorize** to grant permission.

Instance Information

Help Documentation

Basic Configuration

| | | | | |
|---------------------|---------------------|-----------------------------|---------------------------|---------------|
| Instance ID | Region Info | Guangzhou(Guangzhou Zone 7) | High Availability Cluster | Enabled |
| Network information | Security Group | Configure Security Groups | Cloud disk encryption | Not enabled |
| Creation time | Master Public IP | Check for Updates | Billing mode | Pay-as-you-go |
| COS | Custom Service Role | Not set | Host Login | Password |

Software Info

| | |
|----------------------|---|
| Cluster Type | Hadoop |
| Use Cases | Hadoop-Default |
| Product Version | EMR-V2.2.0.tlinux |
| Components to Deploy | hdfs-2.8.5,yarn-2.8.5,zookeeper-3.5.5,knox-1.2.0,hive-2.3.5 |
| MetaDB | cdb-kcx3dajr |
| Hive Metadatabase | cdb-kcx3dajr |
| Kerberos mode | Not enabled |

Hardware Info

| | |
|-------------|---|
| Master Node | 2 |
| Core Node | 4 |
| Common Node | 3 |

Click on **Grant** to complete the authorization process.

Role Management

Service Authorization

After you agree to grant permissions to EMR, a preset role will be created and relevant permissions will be granted to EMR

Role Name: EMR_QCRole

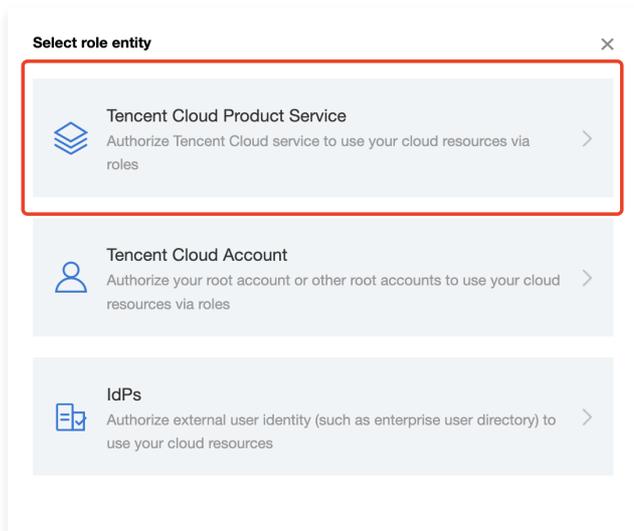
Role Type: Service Role

Description: 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 QcloudAccessForEMRRole

Grant Cancel

For more granular authorization, you can set a custom service role for accessing cloud resources during big data job execution. The custom service role type is "Tencent Cloud Product Service", with "Elastic MapReduce" as the supported role service.



If you need to change the cluster node login password or method after expansion, click on **Confirm** to modify the host login method.

Reset Default Password/SSH Key

1. Modification of the default host login method applies only to new nodes of the cluster.
2. You can go to [CVM console](#) to change host passwords/keys of existing nodes of the cluster. Before change, please see [Shutting Down Instances](#) to understand the potential shutdown risks.

Set Password Associate SSH Key Pair

Username: root

New Password *

Confirm Password *

Confirm Cancel

Note

- The login method set during the initialization of the EMR cluster node host only applies to newly created nodes after the cluster has been established.
- Please go to the CVM interface to modify the existing node host password/key in the cluster CVM console. When you modify the existing node host password/key, please read and understand the potential risks of shutting down [Shutdown Instance](#).

Node Specification Management

Last updated: 2023-12-21 16:10:05

Feature Overview

Node Specification Management is utilized for setting the hardware specifications during the expansion of CORE, TASK, and ROUTER. Each node type can have up to three specifications, and the current default specification will be automatically selected during each node type's expansion. If the current default expansion specification lacks quota, you can reset the default expansion specification. Endeavor to align the new specification's CPU, memory, disk, and other resources with the existing nodes to ensure compatibility with the current component configuration parameters.

Note:

Local disk models are not supported for deployment on Master and Common nodes. Please select a non-local disk model.

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding cluster ID/Name to access the cluster details page.
2. In the cluster details page, select **Cluster Resources > Resource Management > More** and then **Set Node Specifications**. Set the default specifications for different node types based on the billing type. The default is the pay-as-you-go specification, but you can switch to the annual or monthly package specification.

The screenshot shows the 'Resource Management' page in the EMR Console. A table lists four node types: Master, Core, Core, and Task. Each row includes columns for Node Type, Host, Resource Name, IP, Configuration, Deployment Process, Expiration Time, Creation Time, Tag, and Operation. The 'More' dropdown menu is open, showing 'Set Node Specifications' and 'Edit Tag' options.

| Node Type | Host | Resource Name | IP | Configuration | Deployment Process | Expiration Time | Creation Time | Tag | Operation |
|-----------|------|---------------|-------------|---|---|-----------------|---------------------|-----|---|
| Master | HOST | emr-1djq4a | 10.10.10.10 | EMR StandardSA2 CPU: 2-core, memory: 8GB Premium cloud disk: 100G x 1 | knox,HiveMetaStore,HiveServer2,HiveWebHcat,Impala-Catalog,Impala-StateStore,JobHistoryServer,NameNode,Zookeeper,ResourceManager | Pay-as-you-go | 2021-06-15 20:01... | | Change Configuration Start/Stop Services |
| Core | HOST | emr-1djq4a | 10.10.10.11 | 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 Configuration Start/Stop Services |
| Core | HOST | emr-1djq4a | 10.10.10.12 | 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 Configuration Start/Stop Services |
| Task | HOST | emr-1djq4a | 10.10.10.13 | 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 Configuration Start/Stop Services More |

3. The initial default specification for CORE nodes is the specification selected when creating the cluster. You can adjust the default specification by adding or removing, to facilitate normal expansion.
4. TASK nodes and ROUTER nodes do not have initial default specifications. Please add them according to the page prompts.

Set Node Specification

Node specification management is used to set the specifications for core, task, and router node scale-out. Up to 3 node specifications can be set for each node type. The default specifications for each node type are automatically selected for scale-out. If there is no quota of default specifications, you can set a new default specification. Please try to make the CPU, memory, disk, and other resources of the new specifications consistent with the existing nodes so that the component configuration parameters are compatible.

Pay-as-you-go Specification

Core

| Type | 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

| Type | 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

| Type | Model | Disk | Add Time | Operation |
|-------------------------------------|-------|------|----------|-----------|
| + New Specification | | | | |

Confirm

Inspect and Update Public IP

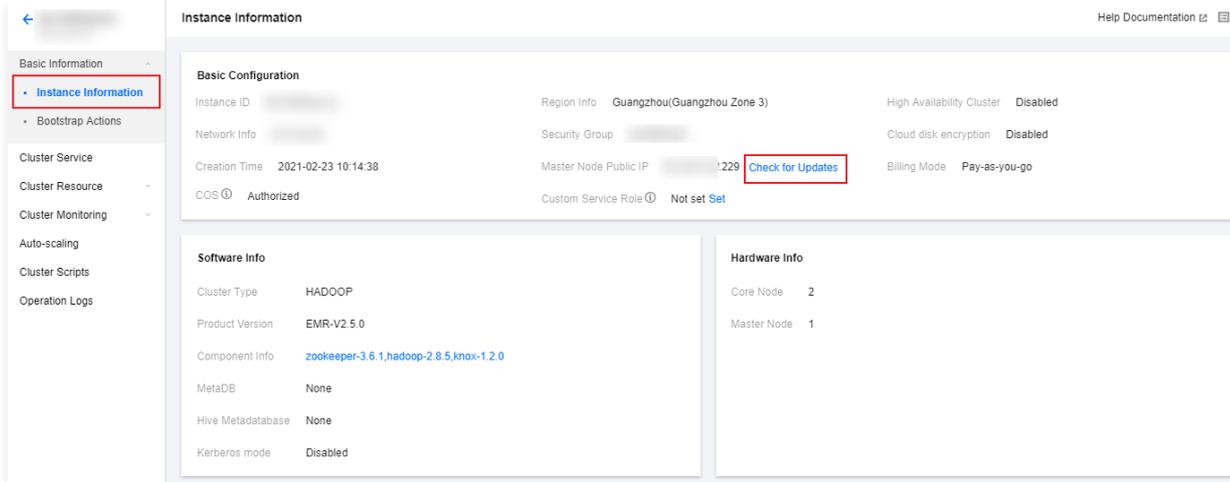
Last updated: 2023-12-21 16:10:32

Feature Overview

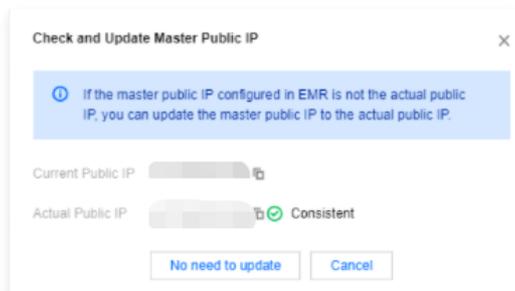
The EMR Console supports the inspection of updates to the public IP on master1, synchronously updating the cluster configuration information after updating or adding the Master public IP, and synchronously updating the service WebUI access address.

Instructions

1. Log into the [EMR Console](#), click on the corresponding **Cluster ID** in the cluster list, and after entering the instance information, you can check for updates to the current Master public IP on the right side of **Basic Configuration > Master Node Public IP**.



2. This feature allows for a comparison between the actual public IP on the master and the public IP stored in EMR. If there is a discrepancy, the latest master public IP can be updated in the relevant EMR business database through further operations. Once synchronization is complete, the current page will automatically refresh.



Logging in to Clusters

Last updated: 2023-12-21 16:10:49

Remote Access Software Login (Local System: Windows)

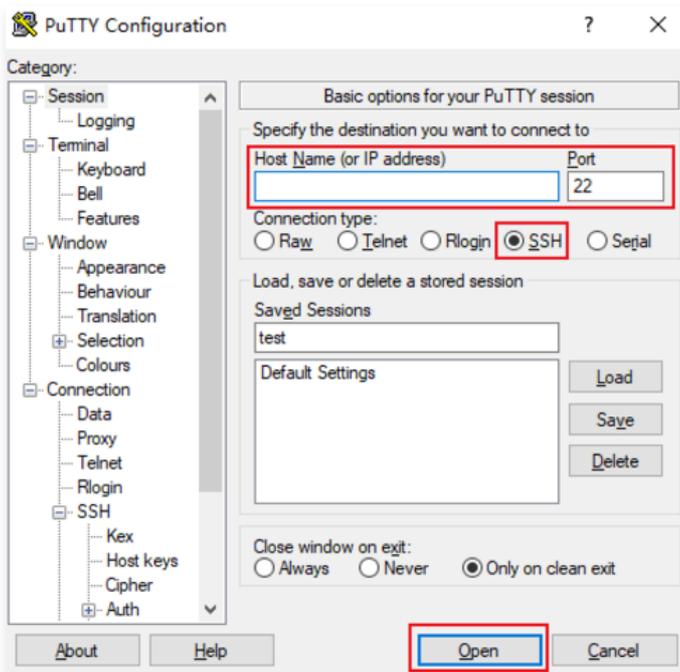
Using Xshell as an example, this section illustrates how a computer with a Windows operating system can utilize remote access software to log into the EMR cluster via password authentication.

Applicable OS

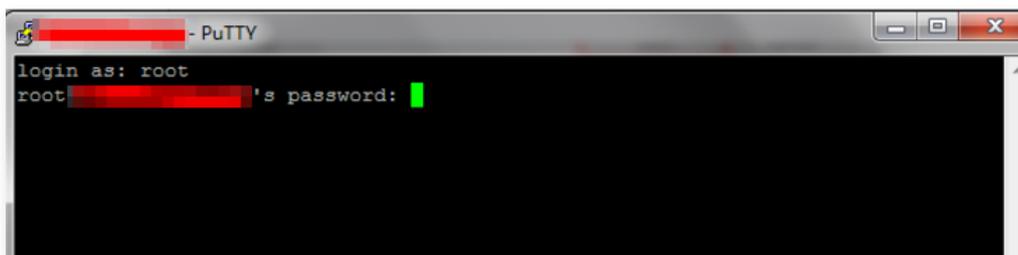
Windows

Log in using a password

1. Install the remote access software for Windows, namely PuTTY. Refer to the [acquisition link](#) for obtaining PuTTY.
2. Launch the PuTTY client, input the following details in the PuTTY Configuration window, and click on **Open** to initiate a new session. As illustrated below:
 - **Host Name:** The public IP of the EMR cluster. Log into the [EMR Console](#) to obtain the cluster's public IP from the list page and the details page.
 - **Port:** The port of the cloud server, which must be set to 22. Ensure that port 22 of the cloud server is open. For more details, refer to [Security Group](#) and [Network ACL](#).
 - **Connection type:** Select "SSH".



3. In the PuTTY session window, input the acquired administrator account and press the Enter key.
4. Enter the acquired login password and press the Enter key to complete the login. As depicted below:



Logging in using SSH (for local systems running Linux/Mac OS)

This section describes how a computer running a Linux/Mac OS system can log into an EMR cluster via SSH.

Applicable OS

Linux or Mac OS

Log in using a password

1. Mac OS users, please open the system's built-in Terminal and execute the following command. Linux users, please directly execute the following command:

```
ssh <username>@<hostname or IP address>
```

- Username: This refers to the administrator account, for instance, root.
 - Hostname or IP address: This refers to the public IP of your EMR instance or your custom domain name.
2. Enter the obtained password (at this point, input is accepted but no output is displayed), then press the Enter key to complete the login process.

Cluster Scale-Out

Last updated: 2023-12-21 16:11:05

Feature Overview

When the EMR cluster's computational and storage resources are insufficient, it is possible to scale up the Core and Task nodes via the console. If the cluster's main node (master node) is under high load or insufficient, the load can be shared by scaling up or adding new routing nodes (router nodes), which can also serve as the cluster's task submission machine. This system supports scaling up and down at any time.

Note

- The current expansion instance specification is the default when creating a new cluster. If the current default specification is out of stock or needs to be adjusted during expansion, it must be set in the [Node Specification](#).
- The selected expansion component will, by default, inherit the cluster dimension configuration, and the expanded nodes will belong to the default configuration group of that node type. If there is a need to adjust the expansion component configuration, it can be set through the **Specify Configuration** option.
- For ClickHouse clusters, the feature to specify a configuration group is not supported during expansion.

Preparations

- **Expanding Router Nodes:** Router nodes can be used as submission machines. They can normally submit computing tasks such as Yarn, Hive, Spark, etc. to the cluster. Therefore, it is recommended to choose a machine type with larger memory, preferably not lower than the Master specification.
- When the cluster is billed on a yearly or monthly basis, the billing mode for expansion nodes (task nodes, router nodes) supports hybrid deployment, i.e., it supports both pay-as-you-go and yearly/monthly billing. Core nodes do not support hybrid deployment and only support yearly/monthly billing.
- When the cluster is billed on a pay-as-you-go basis, all node expansions support the pay-as-you-go billing mode.

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding cluster **ID/Name** to access the cluster details page.
2. In the cluster details page, select **Cluster Resources > Resource Management > Expansion**. Based on business needs, choose the node type (Core, Task, Router) for expansion, billing mode, optional services for expansion, expansion quantity, and other operational configurations.

Cluster Scale-out ✕

Name: panjian123

Project: All projects

Node Type: Core Node Task Node Router Node

Billing Mode: Pay-as-you-go

Scale Out Components: HDFS-3.1.2 YARN-3.1.2

Deployment Process: Show deployment process

Current Specification: EMR StandardS5 / 2-core 8 GB / Efficient cloud disk 100G * 1
The default node specification for scale-out. You can adjust the specification in [Node Specification](#).

Scale-out Quantity: - 1 +

Tag: [+ New Tag](#)
Up to 5 tags can be bound

- **Expansion Service:** Once a service is selected, the newly added nodes will default to deploying the service client.
 - **Specified Configuration:** Locate the component that requires a specified configuration, then select the configuration dimension it should inherit.
 - If you choose to inherit the cluster dimension configuration, the expanded nodes will inherit the cluster dimension configuration, and these nodes will belong to the default configuration group of that node type.
 - If you opt to inherit the configuration group dimension configuration, the expanded nodes will inherit the selected configuration group settings, and these nodes will belong to the chosen configuration group.
 - **Deployment Process:** This refers to the service process information deployed corresponding to the expansion component selected for different node types. If there is a need to adjust the deployment process, the process can be edited.
 - **Post-Expansion Service Non-Start:** If this option is selected during expansion, the nodes expanded in this instance will not initiate the service. To start the service, please use the [Start/Stop Service](#) to initiate the service on the corresponding nodes.
 - **Tags:** Establish identifiers for expanded node resources.
 - **Current Specification:** This is the default specification.
 - If the default expansion specification has not been set, it can be configured in the [Node Specification Management](#) section.
 - During expansion, this is the default specification for nodes. If adjustments are needed, please proceed to [Node Specification Adjustment](#).
3. After selecting the required components and quantity for node expansion, click on **Confirm**. Upon successful payment, the cluster will commence the expansion operation, which typically takes between 10 to 20 minutes.
 4. When the ClickHouse cluster is expanded, the number of nodes for High Availability (HA) instances is even, while there is no limit for non-High Availability (HA) instances. It is necessary to select the associated cluster, with options to choose from existing clusters or create a new one.

△ Note

Upon successful expansion of the ClickHouse cluster, the newly added nodes in the virtual cluster will not contain any data. The system will not perform data migration automatically; it must be done manually to achieve data balance and significantly enhance resource utilization. Please ensure to carry out data migration promptly after the successful expansion of the cluster.

Cluster Scale-In

Last updated: 2023-12-21 16:12:14

Feature Overview

Cluster scaling-down refers to the process of reducing Task nodes through the console when there is an excess of computing resources in the EMR cluster. Moreover, when the cluster has routing nodes (router nodes) that do not share the load of the Master node, or serve as the task submission machine for the cluster, the cluster scaling-down function can be achieved by reducing nodes.

Preparations

Core nodes store data, and it is not recommended to terminate them after expansion to avoid potential data security risks. Once a node is terminated, the data on that node will not be retained. If you choose to scale down (terminate) nodes, it implies that you have confirmed that the data on the selected nodes can be destroyed.

- For annual and monthly subscription clusters: After termination, the cluster will be retained in the recycle bin for 7 days, during which you can choose to restore the cluster. After 7 days, the cluster will be permanently destroyed and cannot be restored. Please proceed with caution.
- For pay-as-you-go clusters: After termination, the cluster will not be retained in the recycle bin and will be permanently destroyed, unable to be restored. Please proceed with caution.

Note

Ensure that you have backed up your data before terminating the cluster, as data cannot be retrieved once the cluster is destroyed.

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding cluster **ID/Name** to access the cluster details page.
2. Within the cluster details page, select **Cluster Resources > Resource Management**. In the node list, select the cloud resources that support scaling down (task nodes or router nodes). When other types of nodes are selected concurrently, the **Scale Down** option cannot be clicked.

Scale In Nodes
✕

1

>

2

Check the Info of the Cluster to Be Terminated

Confirm Termination

1. If there are EIPs (including the IPs on the secondary ENI), the server 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.

2. Please make sure that you back up the data before terminating the instance. After the instance is terminated, the data cannot be restored.

1 node selected. [View Details](#) ▼

| No. | Resource ID | Node Type | Status |
|-----|-------------|-----------|------------|
| 1 | ██████████ | Task | Terminable |

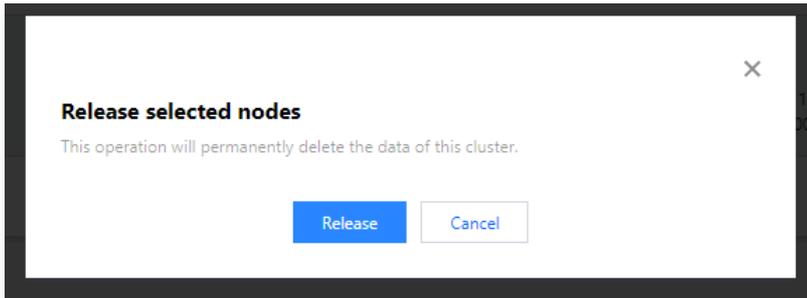
I have read and agree to [Termination Documentation](#) [↗](#)

Next

3. In the pop-up window, you need to confirm the termination (scaling down) of instances. If the cluster's metaDB is a shared meta-database, it will be retained and you will need to manually operate it in the CDB console.

The screenshot shows the 'Resources' page in the Tencent Cloud EMR console. A blue banner at the top contains a warning: "The current cluster resources were purchased based on the official EMR billing rules and the renewal statuses of CVM resources are managed by EMR in a unified manner. For monthly-subscription CVM resources, the billing statuses shown on the current page shall prevail. It is not allowed to adjust the fee policy on the CVM Console." Below the banner, there are tabs for 'All nodes', 'Master', 'Core', 'Common', 'Task', 'Router', 'Metadb', 'Recycle', and 'Renew'. A 'More' dropdown menu is open, showing options like 'Recycle' and 'Release'. A search bar is present with the text "Separate keywords with '|', press Enter to separate filter tags". Below the search bar is a table with columns: Resource ID, Node type, Resource type, IP, Configuration, Creation time, Expiration time, and Operation. The table contains one row for a 'Task' node. The 'Expiration time' column shows "Monthly subscription 6 day(s) later, it will expire". At the bottom right, there is a pagination control showing "Total 1 item" and "Lines per page 20".

- Under "Operations", select **More > Immediate Release**, and in the "Confirm Immediate Release of Selected Nodes" pop-up, click on **Confirm Release**.



Auto Scaling Overview

Last updated: 2023-12-21 16:12:50

Feature Overview

As cluster load fluctuates with business changes, Elastic MapReduce allows for the configuration of scaling rules to automatically supplement or reduce Task node computing resources, swiftly responding to changes in computing demands while saving costs. Auto-scaling supports two scaling strategies: load-based scaling and time-based scaling. Load-based scaling is only applicable to cluster types that include YARN components.

Supports and Limits

1. Auto-scaling is turned off by default. The scaling types supported are **Custom Scaling** and **Managed Scaling**, and only one can be selected. Currently, Managed Scaling is open to a whitelist. If you wish to experience it, please [submit a ticket](#) to enable it.
2. Custom Scaling supports both load-based and time-based strategies. Depending on business needs, the corresponding strategy can be chosen to set scaling rules. It also supports a mix of time-based and load-based elastic rule settings. The rule for triggering is **"First triggered, first executed; if triggered simultaneously, execution is based on rule priority order."**
3. Managed Scaling only supports the HOST resource type, while Custom Scaling supports both HOST and POD resource types, but only one can be selected and they cannot be used simultaneously. If the resource type is switched, the resource specifications and instance deployment methods set by the original resource type will be retained, but they will be in an inactive state and will not be triggered. The currently expanded nodes will also be retained unless the scaling-down rule is triggered, they will not be scaled down. Currently, the POD resource is open to a whitelist. If you wish to experience it, please [submit a ticket](#) to enable it.
4. The instance deployment strategy supports pay-as-you-go and spot instance priority, while POD resources only support pay-as-you-go deployment.

Custom Scaling Configuration

Last updated: 2024-01-12 11:15:12

Basic settings

The basic settings restrict the range of node numbers for the custom scaling function, configure the type of elastic resources, and determine whether elastic scaling supports graceful reduction. They display the current number of elastic nodes in the cluster and provide the option to release elastic instances with a single click.

| Parameter settings | Description |
|------------------------------------|--|
| Minimum number of nodes | When the automatic reduction strategy is triggered, the cluster retains the minimum number of task nodes for elastic scaling. |
| Maximum number of nodes | Upon the activation of the automatic expansion strategy, the cluster retains the maximum number of task nodes for elastic scaling. The cumulative expansion number of a single or multiple specifications cannot exceed the maximum number of nodes. |
| Release all | This refers to the one-click removal of all nodes expanded by automatic scaling, while nodes not affected by automatic scaling remain unaffected. |
| Release of bidding instances | This refers to the one-click removal of bidding instance nodes expanded by automatic scaling, while non-bidding instance resource nodes remain unaffected. |
| Release of pay-as-you-go instances | This refers to the one-click removal of pay-as-you-go instance nodes expanded by automatic scaling, while pay-as-you-go nodes not expanded by automatic scaling remain unaffected. |
| Graceful Scaling Global Switch | By default, this feature is disabled. Once the Graceful Scaling Global Switch is enabled, all scaling-down rules will adopt the graceful scaling strategy. However, individual scaling-down rules can be configured to disable the graceful scaling strategy. Please note: When both the Global Graceful Scaling and individual scaling-down rules are enabled, the Graceful Scaling takes effect. |
| ResourceType | The HOST resource type supports both pay-as-you-go and spot instance billing, whereas the POD resource type only supports pay-as-you-go billing and can exclusively be used for deploying the NodeManager role in Yarn. |

Note:

When the resource type is switched, the corresponding scaling specifications and node selection strategies are simultaneously activated.

Edit basic settings:

Edit basic settings
✕

Min node count (0 - 9999)

Max node count (1 - 9999)

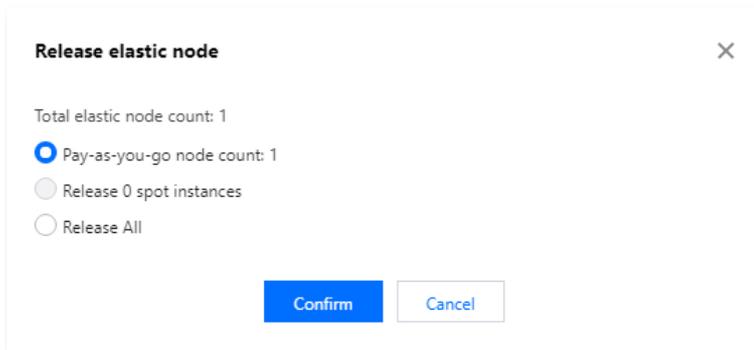
Resource type HOST

Switch to POD if resources are in shortage

Global graceful scale-in ?

Confirm
Cancel

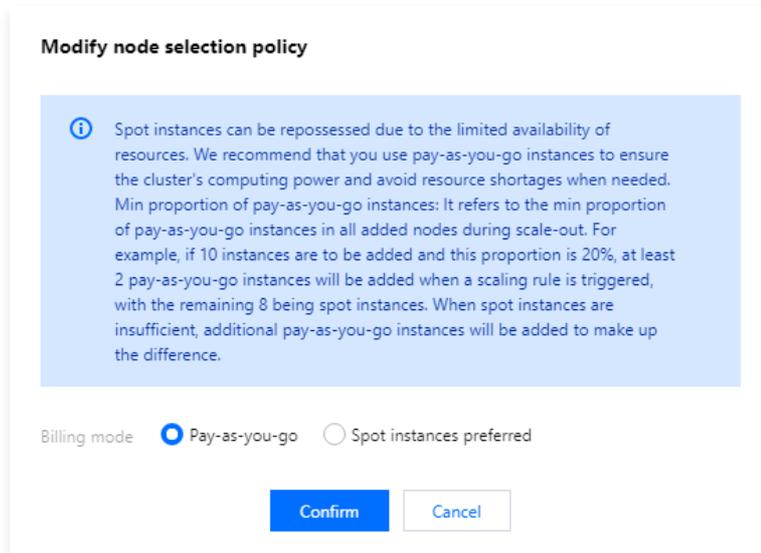
Release elastic node:



Scaling specification management

Scaling specifications refer to the customization of expansion node specifications and node payment strategies through scaling. To maintain linear changes in cluster load, it is recommended to keep the CPU and memory of the scaling specifications as consistent as possible.

- Node Selection Strategy: Supports two strategies, "Pay-as-you-go" and "Spot Instance Priority".



- Pay-as-you-go: When the expansion rule is triggered, all additional nodes are billed on a pay-as-you-go basis to supplement computing power.
 - Spot Instance Priority: When the expansion rule is triggered, spot instances are preferentially added to supplement computing power. When spot instance resources are insufficient, pay-as-you-go resources make up the computing power.
- Minimum Pay-as-you-go Proportion: Ensures the minimum proportion of pay-as-you-go nodes in a single expansion.

ⓘ For instance:

In a single expansion of 10 nodes, if the minimum proportion of pay-as-you-go nodes is 20%, then at least 2 nodes will be supplemented by pay-as-you-go nodes when the expansion rule is triggered. The remaining 8 nodes will be supplemented by spot instances. If the spot instance resources are insufficient for 8 nodes, the resources will be supplemented by pay-as-you-go nodes.

- In the scaling specifications, nodes support addition, deletion, modification, and query. The priority of scaling specifications can be adjusted as needed; the rule priority sequence is in descending order (1>2>3>4>5).

⚠ Note:

When the preset in the basic settings is "Resource Type: POD", the node payment strategy only supports pay-as-you-go billing.

Scaling Specification Management Add Specification

Node selection policy: Pay-as-you-go [Modify](#)

| Model | AZ | Hardware Configuration | Priority | Operation |
|----------------|------------------|--|----------|------------------------|
| EMR Standard55 | 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 strategies for configuring the triggering conditions for scaling actions and the number of changing nodes, supporting both load-based scaling and time-based scaling strategies. Depending on business needs, the corresponding strategy can be chosen to set scaling rules. It also supports a mix of time-based and load-based elastic rule settings. The rule triggering follows the principle of "first triggered, first executed, and simultaneous triggering is executed according to rule priority order".

Setting Load-Based Scaling

When it is impossible to accurately estimate the peaks and troughs of cluster computation, load-based scaling can be used for strategy configuration to ensure that important tasks are completed on time. The load is primarily based on the preset YARN metric statistical rules, and task nodes are automatically adjusted when preset conditions are triggered.

Note:

For detailed information on cluster queue load metrics, please refer to [Queue Load Metrics Correspondence](#).

Click on **Add Rule**, select **"By Load"** as the strategy type on the "New Rule" page, and proceed with rule settings as follows:

Create rule ✕

1. If there are multiple preset resource specifications, we recommend you keep CPUs and memory sizes for different specifications consistent and match them with the resource occupation configuration of the service, so as to avoid the service process start failure due to the inconsistency between node hardware configuration and service configuration.

2. When the scale-in rule is triggered, nodes will be removed in reverse chronological order of their creation time.

Rule type: Scale out Scale in

Policy type: By load By time

Rule name:
1-64 characters; supports Chinese characters, letters, digits, -, and _

Validity: Unlimited

Statistical rule: Rule root AvailableVCores#root >

Add metric

Statistical period: Statistical period 300Second

Repeat count: 1

Scale-out mode: Node

More ▼

Confirm
Cancel

| Configurati on | Description |
|----------------|-------------|
| | |

| | |
|------------------------------|---|
| Rule | Scale Up/Down |
| Policy types | By Load |
| Rule name | The name of the scaling rule, within a cluster, the scaling rule name must be unique (including scale-up and scale-down rules). |
| Validity Period | The load scaling rule is triggered only within the validity period; the default time range is unrestricted, supporting the configuration of scaling rules based on load during custom time periods. |
| Statistical rule | <p>Set single or multiple trigger thresholds based on the selected cluster load metrics; a maximum of five statistical rules can be set; rule statistics based on sub-queues are supported.</p> <ul style="list-style-type: none"> • Rule: Specify the queue and load metrics, set the conditions for triggering thresholds, referring here to the load metrics of YARN. • Statistical Period: The selected load metric, within a statistical period and according to the chosen aggregation dimension (average, maximum, minimum), reaching the trigger threshold constitutes one trigger; currently, three statistical periods are supported: 300 seconds, 600 seconds, and 900 seconds. • Repetition Count: The number of times the load metric reaches the threshold after aggregation, triggering the action of elastic scaling of the cluster after reaching this count. |
| Expansion/Contraction Method | Supports selection of nodes, memory, and core count as three methods; these three methods only support non-zero integer input; when the method is set to core count or memory, expansion ensures maximum computing power for node quantity conversion; contraction ensures normal business operations with minimum unit count for node quantity contraction, contracting in reverse chronological order and ensuring at least one unit is contracted. |
| Scale-out service | The expansion component inherits the cluster dimension configuration by default, and the expanded nodes will belong to the default configuration group of that node type. If adjustment of the expansion component configuration is needed, it can be achieved through specified configuration settings. |
| Node Label | By default, resources expanded will be placed into the Default Label. After setting, expanded resources will be placed into the specified Label. |
| Resource supplement retry | During peak order times, automatic scaling may not reach the elastic target number due to resource contention. When you enable the resource supplement retry strategy, if the configured scaling specification resources are sufficient, the system will automatically retry to apply for resources until it meets or approaches the target number. If resource shortages frequently cause automatic scaling to fall short of expectations, you may consider enabling this configuration. However, be aware that if a retry is triggered, it may extend the automatic scaling time. Please pay attention to the impact on your business after adjusting the strategy. |
| Cooldown period | The interval time (cooldown period ranging from 0 – 43200 seconds) to initiate the next automatic scaling action after the current rule has been successfully executed. |
| Graceful scale-in | Upon enabling the graceful scale-in mode, if a scale-in action is triggered while a node is executing a task, the node will not be immediately released. Instead, it will wait for the task to complete within a custom time frame before scaling in. If the task is not completed by the end of the custom time, scale-in will proceed regardless. |

Scheduled scaling configuration

The cluster's computational load exhibits significant peaks and troughs within certain cycles. To ensure that important jobs are completed on time, time-based scaling can be used for policy configuration. Time-based scaling policies can be set to add or reduce task nodes at fixed times daily, weekly, or monthly.

Click on **Add Rule**, select "**Time-based**" as the policy type on the "New Rule" page, and set the rule as follows:

Create rule
✕

ⓘ 1. If there are multiple preset resource specifications, we recommend you keep CPUs and memory sizes for different specifications consistent and match them with the resource occupation configuration of the service, so as to avoid the service process start failure due to the inconsistency between node hardware configuration and service configuration.

2. When the scale-in rule is triggered, nodes will be removed in reverse chronological order of their creation time.

Rule type Scale out Scale in

Policy type By load By time

Rule name ⓘ
1-64 characters; supports Chinese characters, letters, digits, -, and _

Execution type Once Recurring

Execution time 2024-01-08 15:56

Scale-out mode ⓘ Node

More ▾

Confirm
Cancel

| Configuration | Description |
|------------------------------|---|
| Rule | Scale Up/Down |
| Policy types | Time-based |
| Rule name | The name of the scaling rule, within a cluster, the scaling rule name must be unique (including scale-up and scale-down rules). |
| Execution type | <ul style="list-style-type: none"> Execute once: This refers to triggering a scaling action at a specific time, accurate to the minute. Recurring execution: This refers to setting a scaling action to trigger at each time interval or at a specific time, with support for "daily", "weekly", and "monthly" schedules. Execution Time: Specifies the exact time each day when the scaling action is performed. Rule Validity: The maximum validity period for triggering a single rule in recurring execution. |
| Expansion/Contraction Method | Supports selection of nodes, memory, and cores in three ways; these three methods only support non-zero integer input. When the method of cores and memory is selected, expansion ensures the maximum computing power for the conversion of expansion node quantity. Contraction ensures normal business operations according to the minimum number of units for the conversion of contraction node quantity, contracting in reverse chronological order while ensuring at least one unit is contracted. |
| Scale-out service | The expansion component inherits the cluster dimension configuration by default, and the expanded nodes will belong to the default configuration group of that node type. If there is a need to adjust the expansion component configuration, it can be done through specified configuration settings. |
| Node Label | By default, resources expanded will be placed in the Default Label. If set, the expanded resources will be placed in the specified Label. |
| Resource supplement retry | During peak order times, automatic scaling may not reach the elastic target number due to resource contention. When you enable the resource supplement retry strategy, if the configured scaling specification resources are sufficient, the system will automatically retry to apply for resources until it meets or approaches the target number. If resource shortages frequently cause automatic scaling to fall short of expectations, you may consider enabling this configuration. However, be aware that if a retry is triggered, it may extend the automatic scaling time. Please pay attention to the impact on your business after adjusting the strategy. |
| Retry time | Elastic scaling may not be able to execute at the specified time due to various reasons. By setting a retry |

| | |
|-----------------------|---|
| after expiration | expiration time, the system will attempt to execute at intervals within this time range, until it executes scaling when the conditions are met. |
| Cooldown period | The interval time (cooling time range 0 – 43200 seconds) to initiate the next automatic scaling action after the current rule has been successfully executed. |
| Scheduled termination | Specify the usage duration of expanded resources, and when the scaling-down rule is triggered, the nodes of the current batch are not affected by the scaling-down rule. By default, "No limit" is selected. It supports customizing the termination duration, the input value should be an integer, and the input value range is (1–24) hours. Use Case Description: This is used when there is a need to supplement computing power during a fixed period and maintain the computing power within a day, and when other scaling-down rules do not affect this batch of resources. |
| Graceful scale-in | After enabling the graceful scaling-down mode, if a scaling-down action is triggered while a node is executing a task, the node will not be released immediately, but will wait for the task to be completed within a custom time before scaling down; if the task is not completed when the custom time ends, scaling down will also be carried out. |

Managed Scaling Configuration

Last updated: 2023-12-21 16:13:09

Upon activation of the managed scaling feature, the system persistently monitors the load of the Yarn cluster, calculating fluctuations in peak load over the recent ten-minute period, thereby automatically supplementing or reducing task nodes. Managed scaling is exclusively applicable to cluster types incorporating Yarn components.

Basic settings

The basic settings delineate the range of node numbers for the expansion and contraction of the managed scaling feature, as well as the minimum node count for pay-as-you-go billing.

- **Minimum Node Count:** When the managed contraction strategy is triggered, the cluster retains at least a certain number of task nodes for elastic scaling.
- **Maximum Node Count:** When the managed expansion strategy is triggered, the cluster retains at most a certain number of task nodes for elastic scaling. The cumulative expansion count of a single or multiple specifications cannot exceed the maximum node count.
- **Minimum Node Count for Pay-as-you-go Billing:** This refers to the least number of pay-as-you-go nodes to be expanded after the expansion is triggered. It is used to set the proportion of pay-as-you-go nodes and bidding instances, defaulting to the maximum node count.

For instance

With the minimum node count set to 0, the maximum node count set to 100, and the minimum pay-as-you-go node count set to 10, upon triggering expansion, at least 10 pay-as-you-go nodes will be expanded. The remainder will be supplemented by bidding instances. If there are insufficient bidding instances, they will be supplemented by pay-as-you-go billing.

Scaling specification management

Scaling specifications refer to the node specifications that can be expanded through managed scaling. Each cluster can configure up to five scaling specifications. When the expansion rule is triggered, expansion will be carried out according to the priority of the specifications. If the quantity of high-priority specifications is insufficient, the next priority resource specifications will supplement the high-priority specifications for expansion to supplement computing resources. To maintain the linear change of cluster load, it is recommended to keep the CPU and memory of the scaling specifications consistent. The managed scaling function only supports the host resource type.

- Nodes within the scaling specifications support addition, deletion, modification, and inquiry, allowing for the adjustment of scaling specification priority as needed.
- The expansion order for the five specifications is as follows (the execution order for pay-as-you-go and bidding instances is the same):
 - When resources are abundant: $1 > 2 > 3 > 4 > 5$

For instance

Given five preset specifications and sufficient resources, when the expansion rule is triggered and requires the expansion of 10 nodes, the expansion will proceed with 10 nodes according to the first specification, while the remaining preset specifications will not be selected.

- When resources are insufficient: $1 + 2 > 1 + 2 + 3 > 1 + 2 + 3 + 4 > 1 + 2 + 3 + 4 + 5$

For instance

Given that preset specification 1 has 8 nodes, specification 2 has 4 nodes, and specification 3 has 3 nodes, when the expansion rule is triggered and requires the expansion of 13 nodes, the expansion will proceed with 8 nodes according to specification 1, 4 nodes according to specification 2, and 1 node according to specification 3.

- When a resource specification is out of stock, assuming specification 2 is unavailable: $1 + 3 > 1 + 3 + 4 > 1 + 3 + 4 + 5$

For instance

- Given that preset specification 1 has 8 nodes, specification 2 is out of stock with no nodes, and specification 3 has 3 nodes, when the expansion rule is triggered and requires the expansion of 10 nodes, the expansion will proceed with 8

nodes according to specification 1, specification 2 will not be selected, and 2 nodes will be expanded according to specification 3.

- Given that preset specification 1 has 8 nodes and the remaining preset specifications are out of stock, when the expansion rule is triggered and requires the expansion of 10 nodes, the expansion rule will be activated, and 8 nodes will be expanded according to specification 1, resulting in a partially successful expansion.

Managed Scaling Monitoring Metrics

Managed scaling will monitor multiple metrics and calculate the recommended node count for each, subsequently making expansion and contraction decisions based on the provided node quantities.

| Managed Scaling Monitoring Metrics | MetricsDescription |
|------------------------------------|--|
| AvailableMemPercentage | Percentage of Remaining Memory |
| AvailableVCoresPercentage | Percentage of Available YARN Virtual Cores |

Statistical Rules: The cluster load metrics within a statistical cycle, according to the selected aggregation dimension, are set to handle the peak load within the last 10 minutes.

Statistical Cycle: The duration of the metric statistics, one minute.

The default principle of managed scaling is: rapid expansion, cautious contraction, with contraction being for graceful reduction.

Repairing Disks

Last updated: 2024-01-12 11:16:12

Feature Overview

The EMR console supports automatic detection of local disk replacement events. After the disk is replaced, you can independently initialize the new disk operation via the console.

Note

- Upon receiving a CVM disk failure notification and subsequently repairing the physical disk or replacing the disk according to the CVM notification content, the EMR console can then initiate the 'Disk Repair' operation.
- Following a disk replacement, the data on the disk will be lost. Please ensure that the data on the disk has been backed up.

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding **Cluster ID/Name** to enter the cluster details page.
2. In the cluster details page, select **Cluster Resources > Resource Management**, and perform the **Disk Repair** operation on the node for which the disk has been replaced.
3. During the operation, services on the current node will be restarted, rendering the services and node unavailable during the restart process. It is recommended to perform the repair operation during off-peak business hours.

Kudu Service Recovery

Note

- When there are multiple local disks, and one or more disks have been repaired and the EMR disk repair function is used, the node where the KuduServer service is deployed;
- Due to the capabilities of Kudu's `fs_data_dirs`, if one or more disks have been formatted, to ensure the normal startup of KuduServer, it only supports that all data directories configured on the KuduServer node are empty. Customers need to assist in confirming these data directories, and apart from Kudu storage data, they have not been misused by other customer's own businesses.

Scenario:

Specifically, in the EMR console cluster service, the health status of the KuduServer in the node where the disk has been replaced is "Unavailable".

Confirm Data Consistency and Recovery:

- 1.1 Confirm that the data in the directory (specific viewing method is as follows) has no other use besides Kudu; if there are other uses, please first migrate the relevant data to other directories not configured by `fs_data_dirs`, and then perform the following operations.

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

```
--fs_data_dirs=/data/emr/kudu/tserver,/data1/emr/kudu/tserver
--webserver_doc_root=/usr/local/service/kudu/www
--fs_wal_dir=/data/emr/kudu/tserver/wal
--raft_get_node_instance_timeout_ms=300000
--log_dir=/data/emr/kudu/log
--fs_metadata_dir=/data/emr/kudu/tserver/meta
```

- 1.2 Log into the node with abnormal local disk to view the log: `/data/emr/kudu/log/kudu-tserver.INFO :`

```
10110 14:16:50.918087 96777 minidump.cc:244] Setting minidump size limit to 20M
10110 14:16:50.918205 96777 env_posix.cc:2179] Not raising this process' open files per process limit of 1000000; it is already as high as it can go
10110 14:16:50.918251 96777 file_cache.cc:492] Constructed file cache file cache with capacity 400000
10110 14:16:50.918663 96777 hybrid_clock.cc:249] auto-selected time source: builtin
10110 14:16:50.920179 96777 hybrid_clock.cc:583] waiting up to --ntp_initial_sync_wait_secs=60 seconds for the clock to synchronize
10110 14:16:50.927374 96777 hybrid_clock.cc:601] HybridClock initialized: now 1673331410927647 us; error 18281 us; skew 500 ppm
10110 14:16:50.927525 96777 webserver.cc:395] Webserver started at http://0.0.0.0:8050/ using document root /usr/local/service/kudu/www and password file <none>
10110 14:16:50.927855 96777 server_base.cc:612] This appears to be a new deployment of Kudu; creating new FS layout
10110 14:16:50.931181 96777 tablet_server_main.cc:401 Already present: RunTabletServer() failed: FS layout already exists; not overwriting existing layout: unable to create file system roots: FSManager roots alr
ready exist: /data/emr/kudu/tserver,/data1/emr/kudu/tserver
```

Execute the following command as root user to clean up the related inconsistent data:

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

This command assumes that `fs_data_dirs` is configured as `/data/emr/kudu/tserver/` , `/data1/emr/kudu/tserver/` . The specifics can be viewed according to `/usr/local/service/kudu/conf/tserver.gflags` .

1.3 Observe the status of the KuduServer service.

Note

If you encounter any issues during the operation, please [submit a ticket](#) promptly for us to verify and address.

Disk Update Check

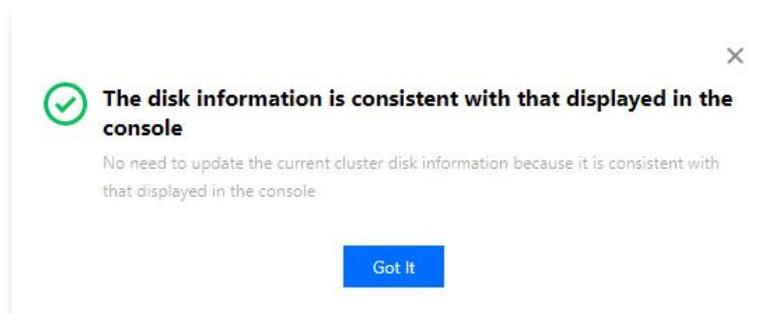
Last updated: 2023-12-21 16:13:39

Feature Overview

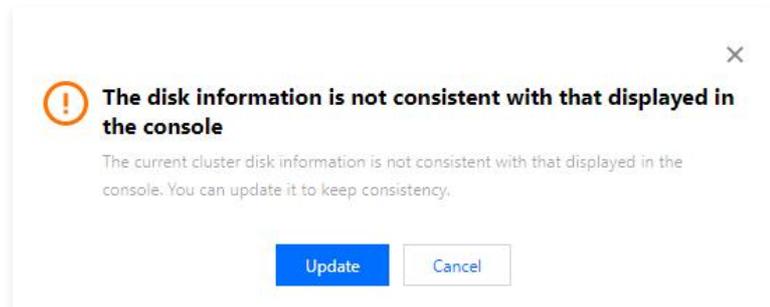
Verify the consistency of the cluster disk information with the actual display information on the console, and support updates to maintain consistency.

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding cluster **ID/Name** to access the cluster details page.
2. In the cluster details page, select **Cluster Resources** to enter the resource management page. Choose **More Actions > Disk Check Update**.



When the disk information in the cluster is inconsistent with the console display, clicking update will refresh the disk usage, total disk capacity, and disk space usage rate in the cluster.



Scaling up Cloud Disks

Last updated: 2024-01-12 11:17:35

Feature Overview

As business growth necessitates the expansion of data storage space within the cluster, it becomes essential to carry out scaling operations. This document provides a comprehensive guide on how to expand cloud disk capacity via the EMR console.

Note

- Only cloud data disk expansion is supported; system disk and local disk expansion are not supported.
- Batch expansion of cloud data disks across multiple nodes is not supported.
- It is recommended to create a snapshot of the cloud disk prior to executing the expansion operation, to prevent file system damage due to erroneous operations.
- To prevent data loss, disk capacity can only be increased and not decreased.

Expanding cloud disks

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding **Cluster ID/Name** to enter the cluster details page.
2. Navigate to the Resource Management page, select the node that needs to be expanded, and click on **Cloud Data Disk Expansion** in the list of operations. This will take you to the cloud data disk expansion configuration.

The screenshot shows the EMR console 'Resources' page. A table lists resources with columns: Resource ID, Node type, Resource type, IP, Configuration, Creation time, Expiration time, and Operation. The third row is selected, and a context menu is open over the 'Operation' column, with 'Scale up cloud data disk' highlighted.

| Resource ID | Node type | Resource type | IP | Configuration | Creation time | Expiration time | Operation |
|--|-----------|---------------|-----------------------|---|---------------------|-----------------|--|
| emr-...v38hcxh imp-...eg | Master | HOST | 10... | EMR StandardSA2 CPU: 8-core; memory: 32GB System disk: Premium cloud disk 70G x 1 Data disk: Premium cloud disk 200G x 1 | 2024-01-04 11:30:23 | Pay-as-you-go | Change configuration More |
| emr-...v2y8be3 ip-...manscu | Core | HOST | 192.168...8 (private) | EMR StandardSA2 CPU: 4-core; memory: 8GB System disk: Premium cloud disk 70G x 1 Data disk: Premium cloud disk 200G x 1 | 2024-01-04 11:30:23 | Pay-as-you-go | Change configuration More |
| emr-...vbk2y91j ip-...1b3y | Core | HOST | 192.16...90 (private) | EMR StandardSA2 CPU: 4-core; memory: 8GB System disk: Premium cloud disk 70G x 1 Data disk: Premium cloud disk 200G x 1 | 2024-01-04 11:30:23 | Pay-as-you-go | Change configuration More Start/Stop services Renew Set auto-renewal Scale up cloud data disk |
| emr-...voviqay95 en-...z3q2r1640275-0 | Task | POD | 192.1...113 (private) | CPU: 4-core; memory: 8GB System disk: 0G x 1 Data disk: Premium cloud disk 100G x 1 | 2024-01-04 15:56:53 | -- | Change configuration More |
| emr-...da95uzen | Task | POD | 19...66.24 (private) | CPU: 4-core; memory: 8GB System disk: 0G x 1 Data disk: Premium cloud disk 100G x 1 | 2024-01-04 15:56:53 | -- | Change configuration More |

3. If multiple cloud data disks are mounted on the current node, it supports batch expansion of multiple disks on a single node, and the target capacity of multiple disks is consistent after expansion.

Scale up cloud data disk ✕

✔ **Select cloud data disk** > 2 **Change size**

Notes:
 We strongly recommend that you snapshot the cloud disks before scale-up to avoid misoperation-related damage to the file system.
 Only disk scale-up is allowed to prevent data loss.

Cloud data disks selected: 1 [Less](#) ▾

| Cloud disk ID | Billing mode | Expiration time | Current size |
|---------------|---------------|-----------------|--------------|
| disk-pe19afqk | Pay-as-you-go | - | 200GB |

Common target size

- 400 + GB

200 GB 500 GB 1000 GB 32000 GB

Cost

Back
Confirm
Disable

4. After the disk is expanded, it will be automatically initialized, eliminating the need for manual updates to the disk information.

Changing Configurations

Last updated: 2023-12-21 16:13:58

Feature Overview

In practical application, there may be instances where nodes within the cluster, particularly the MASTER nodes, are deficient in CPU or memory. In such cases, the CPU or memory size can be increased by modifying the configuration. This document provides an overview of the relevant operations for changing instance configurations through the [EMR Console](#).

Note

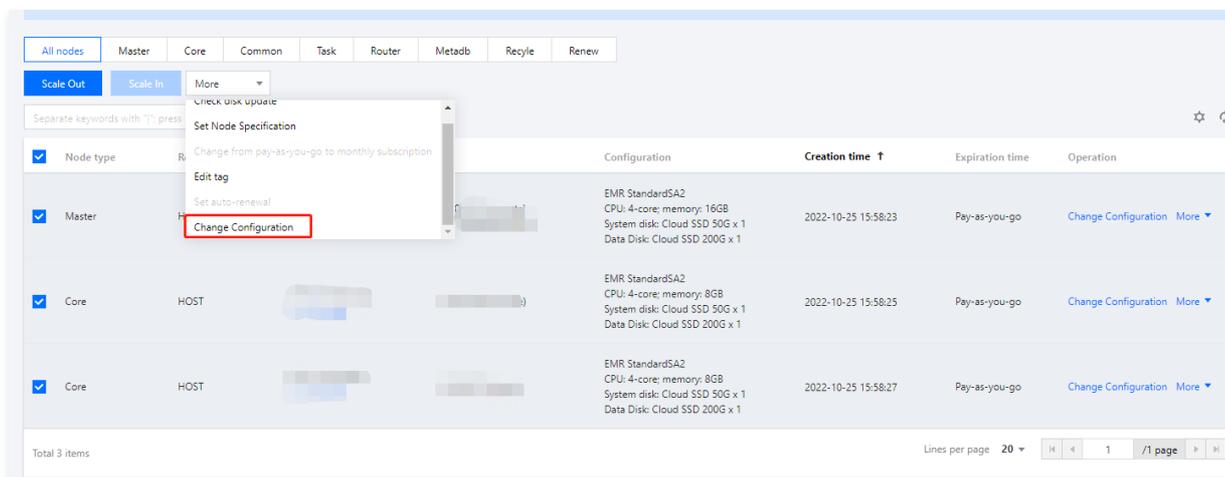
- During the process of node configuration changes, the system will shut down. This shutdown may potentially disrupt the normal operation of the cluster, and could even result in business interruption. Please proceed with caution after careful evaluation.
- During node reconfiguration, changes to the size of the data disk and system disk are not supported.

Preparations

1. After reconfiguration of pay-as-you-go nodes, the billing tier will reset to the first tier. For annual and monthly subscription clusters, the corresponding fee difference needs to be compensated.
2. Local disk models, POD resources, and spot instance billing models do not support reconfiguration.
3. Costs in a batch configuration change will be deducted item by item. Please ensure your account balance is sufficient.
4. The refund amount will be returned to your Tencent Cloud account in accordance with the ratio of cash and gift money used at the time of purchase; discounts and vouchers used during purchase are non-refundable.

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding cluster ID/Name to access the cluster details page.
2. In the cluster details page, select **Cluster Resources** to enter the resource management page. Based on business needs, select the nodes that need to be changed and perform the **Change Configuration Operation**; batch changes are supported, but only nodes of the same billing type can be changed to a uniform configuration during a batch change.



3. On the configuration adjustment page, confirm the relevant change information. It is crucial to read the important prompts carefully and check the box to agree to the change information.
4. On the target configuration page, select configurations such as machine type, instance type, and machine type list. After confirming the cost is correct, click **Confirm Changes** to adjust the configuration.
5. The cost details page will separately display the fees required to change different nodes to a uniform configuration.

Change Configuration ✕

1 **Select Target Configuration** >
 2 Cost details >
 3 Shutdown notes

i • The configuration change does not apply to local models and pod resources.

• The configuration change does not apply to data and system disks.

• If you want to change resources of the corresponding component after the configuration change, redeliver the configurations in "Configuration management" and restart the service.

Nodes selected: 1 [See Less](#) ▾

| Resource ID | Node type | Current configuration | Billing Period | Operation |
|-------------|-----------|--|----------------|--------------------------|
| [blurred] | master | EMR StandardSA2 CPU: 4-core; memory: 16GB | - | Adjustable configuration |

Select a target configuration

Model: Standard MEM Optimized Compute

Type: Standard SA2 Standard S5

Model List **i**: Specification: SA2.2XLARGE16; CPU: 8-core; m ▾

Total cost to be frozen

Next
Cancel

6. (Optional) After changing the configuration, if you need to adjust the resources of the corresponding components, you will need to reissue the configuration in configuration management and restart the service.

Change Configuration ✕

✓ **Select Target Configuration** >
 2 **Cost details** >
 3 Shutdown notes

i • Costs in a batch configuration change will be deducted item by item. Please ensure your account balance is sufficient.

• Discounts or vouchers (if any) used in the purchase will not be returned.

• If the configuration is changed, the pay-as-you-go nodes will be charged again from the first tier. Please proceed with caution. [View details](#)

| Resource ID | Node type | Current configuration | Target configuration | Billing Period | Cost |
|-------------|-----------|--|--|----------------|---------------------|
| [blurred] | master | EMR StandardSA2 CPU: 4-core; memory: 16GB | Standard SA2 CPU: 8-core, memory: 16 GB | - | i - USD/hour |

I have read and understood the notes and agree to the operation

Back
Next
Cancel

Note

- Yarn resources are automatically adjusted according to machine specifications by default. After reconfiguration, the resource size changes with the reconfiguration specifications, eliminating the need for manual configuration adjustments.

- If you have manually adjusted the Yarn resource configuration, after reconfiguration, you need to go to **Configuration Management** to modify the parameter values of `yarn.nodemanager.resource.cpu-vcores` and `yarn.nodemanager.resource.memory-mb`. Click **Save Configuration**, and after issuing the configuration, restart the NodeManager service to complete the Yarn resource configuration update.

Exporting Software Configuration

Last updated: 2023-12-21 16:14:20

Feature Overview

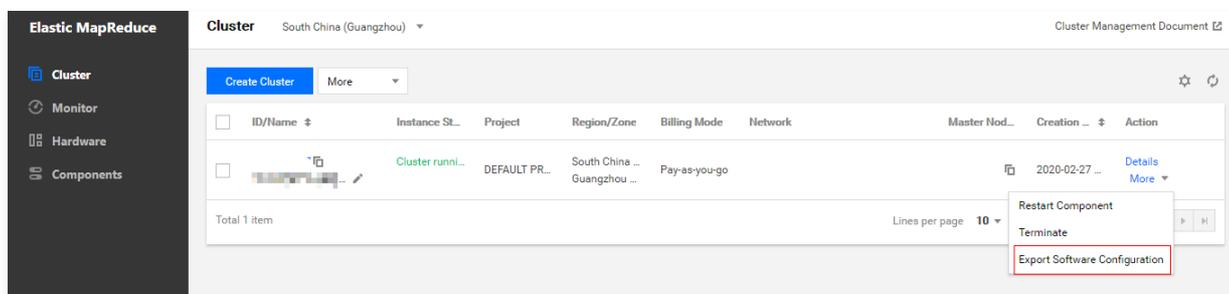
Through the [EMR Console](#), one can export the software configuration parameters of existing clusters. Subsequently, these parameters can be utilized for [software configuration](#) when creating new clusters, thereby swiftly establishing a familiar cluster.

Note

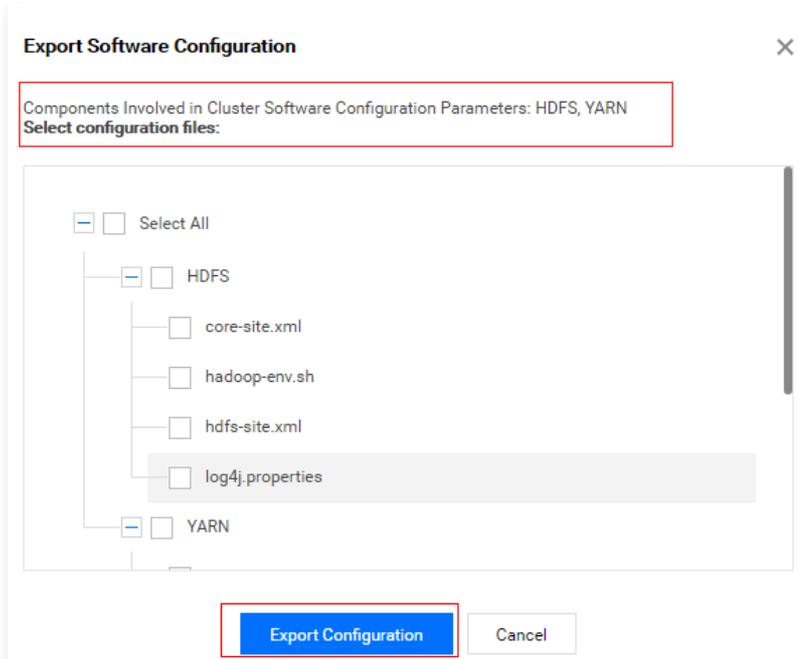
- The exported configuration file pertains to the cluster dimension and does not support the export of configurations independently modified at the group or node dimension.
- If a new cluster intends to reuse the configuration of an old cluster, it is recommended to export the modified configuration files of the old cluster, eliminating the need to export all configuration files.

Instructions

1. Log in to the [EMR Console](#) to access the cluster list page.
2. In the management column of the cluster to be exported, select **More > Export Software Configuration**.



3. Select the files you wish to export, then choose **Export Configuration** to download the software configuration files.



Cluster Scripts

Last updated: 2023-12-21 16:14:37

Feature Description

The cluster script feature facilitates the execution of specified scripts on multiple nodes simultaneously, thereby enhancing the efficiency of batch operations. At any given time, only one cluster script can be run within a cluster. If a cluster script is currently in operation, the execution of a new cluster script cannot be initiated. This feature is particularly useful for tasks such as the installation of third-party software and modification of the cluster operating environment.

Note

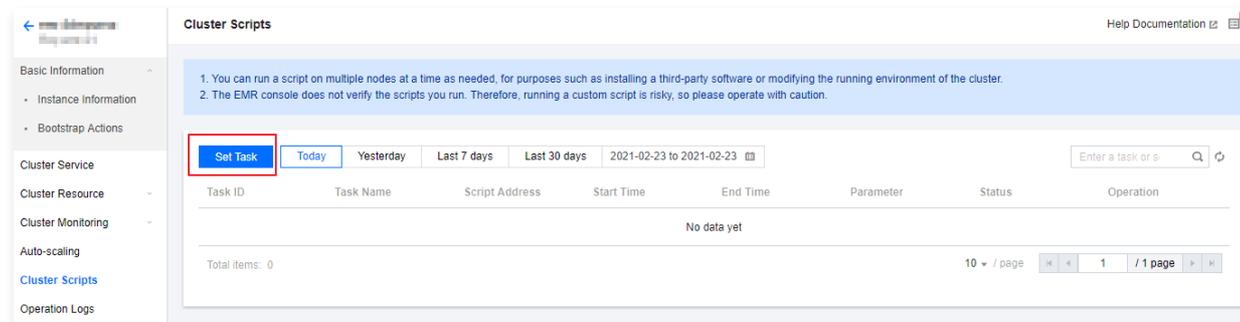
- The execution of script files currently only supports the selection of cross-regional files from Cloud Object Storage (COS).
- The selection of execution nodes is limited to those within the current cluster and does not extend to nodes across different clusters.
- Custom parameters can be configured based on specific business requirements.
- The Elastic MapReduce console does not perform any validation on the execution scripts. The operation of executing custom scripts is considered high-risk, hence it should be undertaken with utmost caution.

Preparations

- Only clusters in the running state can utilize the cluster script function.
- The execution script must be in the standard storage format of Object Storage Service (COS) and only shell script files are supported.
- Prior authorization for EMR to access COS is required to utilize the cluster script function.

Instructions

1. Log in to the [EMR Console](#). In the cluster list, click on the corresponding cluster ID. After entering the instance information, select **Cluster Scripts > Set Task**.



2. When setting up a task, you need to specify the task name, select the script file to be executed, choose the execution node, and customize parameters. After the task configuration is complete, click on **Run** to generate the task in the task list.

Set Task

Task Name *

The length should be 6-36 characters. Chinese, letters, numbers, -, and _ are supported.

Script *

Select a shell script file in the COS STANDARD storage class.

Nodes *

| <input type="checkbox"/> | Node Type | Node IP | Deploy Task | Deployment Process |
|--------------------------|-----------|---------|---------------------------------------|--|
| <input type="checkbox"/> | core | | HDFS, FILEBEAT, YARN | DataNode, Filebeat, NodeManager |
| <input type="checkbox"/> | master | | YARN, HDFS, ZOOKEEPER, FILEBEAT, KNOX | Ldap, NameNode, QuorumPeerMain, ResourceManager, Filebeat, Gateway, JobHistoryServer |

Selected nodes (0):

Custom Parameter

3. Tasks generated in the task list, depending on the execution status, may be entirely successful, entirely failed, or partially failed. You can click on **Details** for further information.

emr-6dmpunsz Roy-emr-01
Cluster Scripts Help Documentation

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.

Set Task
Today
Yesterday
Last 7 days
Last 30 days
2021-02-23 to 2021-02-23

| Task ID | Task Name | Script Address | Start Time | End Time | Parameter | Status | Operation |
|---------|-----------|--|---------------------|---------------------|-----------|---------|--|
| 191 | Roy-001 | <code>#!/bin/bash #set -e #set -x #set -o pipefail #set -u #set -o errexit #set -o nounset #set -o xtrace #set -o errtrace #set -o errwarn #set -o errfail #set -o errtrace #set -o errwarn #set -o errfail</code> | 2021-02-23 10:19:52 | 2021-02-23 10:19:53 | | Running | Details Delete |

Total items: 1 10 / page 1 / 1 page

See below:

Details Roy-001

This script is run on 1 node(s). Successful: 0; failed: 0; running: 1.

| Node IP | Node ... | Start Time | End Time | Execu... | Running Result |
|---------|----------|---------------------|---------------------|----------|--------------------|
| | core | 2021-02-23 10:19... | 2021-02-23 10:19... | Running | stdout: stderr: |

4. Cluster scripts may run successfully on some nodes and fail on others. You can replicate the failed nodes en masse for another attempt.

Cluster Restoration

Last updated: 2023-12-21 16:14:51

Feature Description

For clusters billed on an annual or monthly basis that are in isolation and have not been reclaimed, it is possible to restore these clusters through the renewal operation in the Elastic MapReduce console cluster list.

Note

When a pay-as-you-go cluster is in isolation due to insufficient account balance, it can be restored by renewing the cluster in the Renewal Center.

Preparations

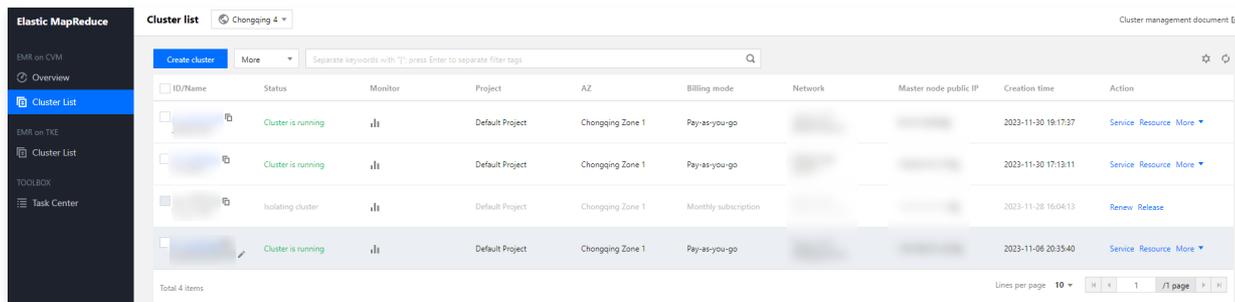
1. The account balance is sufficient to restore all nodes of the current cluster.
2. If the cluster includes a MetaDB, it is necessary to first proceed to the Cloud Database Console to restore the MetaDB.
3. In a mixed billing cluster, only the nodes billed on an annual or monthly basis can be restored; the pay-as-you-go nodes will not be restored.

Instructions

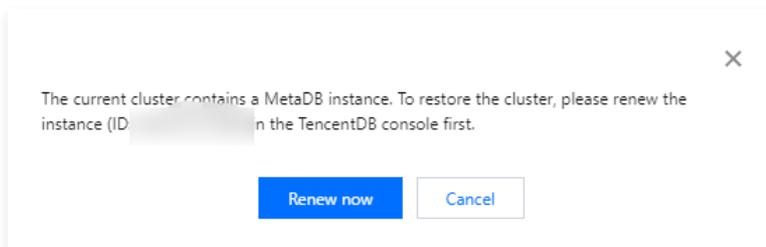
Note

- Once the isolated status renewal and cluster restoration is completed, the cluster cannot be destroyed again within six hours.
- All monitoring data from the period of isolation until the restoration of the cluster will be missing.

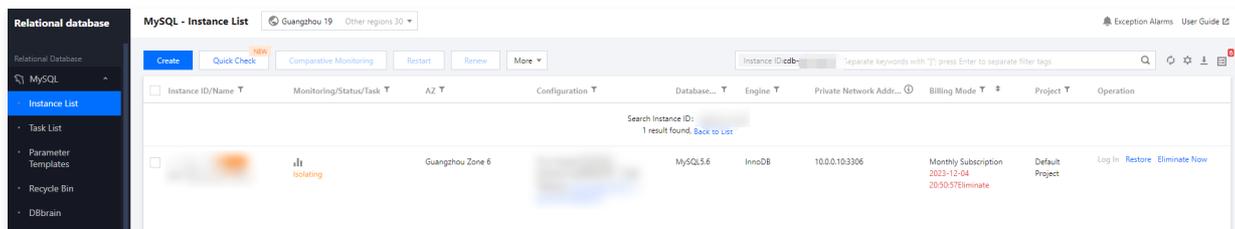
1. Log in to the [Elastic MapReduce Console](#), select the cluster that needs to be restored from the cluster list, and choose **Manage > Renew** to proceed with the renewal.



2. If the cluster contains a MetaDB, you must first go to the Cloud Database Console to restore the MetaDB. Click on **Renew now**.



3. In the [TencentDB Cloud Database Console](#), click on **Restore** for the selected isolated database instance to proceed with the renewal and recovery.



4. If the cluster does not contain a MetaDB, or if the MetaDB has already been renewed and restored, click on **Renew** for the corresponding cluster in the Elastic MapReduce Console to renew and restore all monthly nodes in the cluster.

Renewemr-6oxg5lyp



1. When a monthly subscription cluster is isolated and not repossessed, you can recover the cluster by renewing the monthly subscription nodes in it.
2. The fees shown here are for the renewal of the EMR monthly subscription nodes, excluding the pay-as-you-go nodes. You can check the bill for details.

| Renewed item | Node type | IP | Instance specification | Expiration time |
|-----------------|-----------|-------------|---|---------------------|
| emr-vm-hs7p9zr1 | Master | 172.30.1.95 | EMR StandardSA2 CPU: 4-core; memory: 16GB Premium cloud disk: 100.00 GB | 2023-11-28 20:20:26 |
| emr-vm-9s6w3vv5 | Core | 172.30.3.2 | EMR StandardSA2 CPU: 4-core; memory: 8GB Premium cloud disk: 100.00 GB | 2023-11-28 20:20:15 |
| emr-vm-62o4qxyt | Core | 172.30.0.38 | EMR StandardSA2 CPU: 4-core; memory: 8GB Premium cloud disk: 100.00 GB | 2023-11-28 20:20:21 |

Total items: 3

10 / page ⏪ ⏩ 1 / 1 page ⏪ ⏩

Renewal period 1 month 2 3 4 5 6 7 8 9 10 11 1 year 2 years 3 years 5 years

Total cost [blurred]

Confirm Cancel

Cluster Termination

Last updated: 2023-12-21 16:15:09

Feature Overview

When the EMR cluster is no longer required, it can be dismantled via the [EMR Console](#).

Preparations

- For annual and monthly subscription clusters: After termination, the cluster will be retained in the recycle bin for 7 days, during which you can choose to restore the cluster. After 7 days, the cluster will be permanently destroyed and cannot be restored. Please proceed with caution.
- For pay-as-you-go clusters: After termination, the cluster will not be retained in the recycle bin and will be permanently destroyed, rendering it irrecoverable. Please proceed with caution.
- Ensure that you have backed up your data before terminating the cluster, as data cannot be retrieved once the cluster is destroyed.
- If there are elastic IPs (including those on auxiliary network cards), they will continue to be retained after the machine is returned, and idle IPs will continue to incur charges. If retention is not required, please release them on the corresponding resource manager page.

Instructions

Note

When the MetaDB in the cluster being destroyed is associated with an external cluster as a Hive metadata database, the MetaDB in the current cluster will be retained during cluster destruction. If you need to destroy the database, please proceed to the cloud database for destruction. Once destroyed, the Hive metadata database cannot be recovered. Please proceed with caution.

Log in to the [EMR Console](#), select **Manage > More > Terminate** to access the cluster termination page. On this page, confirm the details of the cluster you wish to terminate. Once confirmed, check the box for **I have read and agree** and click **Next**.

The screenshot shows the 'Terminate Cluster' dialog box in the EMR console. The dialog is titled 'Terminate Cluster' and has two tabs: '1 Check Info of the Item to Terminate' and '2 Confirm Termination'. The 'Confirm Termination' tab is active. The dialog displays the following information:

- This resources to be terminated are as follows:**
 - 8 instances, Collapse

| No. | Instance name | Instance ID | Current upper li... | Instance Type |
|-----|---------------|-------------|---------------------|----------------|
| 1 | emr-... | ... | 0 | EMR StandardS2 |
| 2 | emr-ma... | ... | 0 | EMR StandardS2 |
| 3 | emr-cor... | ... | 0 | EMR StandardS2 |
 - 16 cloud disks, Collapse

| No. | Cloud disk name | Cloud disk ID | Attributes | Configuration |
|-----|-----------------|---------------|-------------|---------------------------|
| 1 | ... | ... | Data disk | Efficient cloud disk 200G |
| 2 | ... | ... | System disk | Efficient cloud disk 50G |
- This resources to be retained are as follows:**
 - 0 cloud disks
 - 0 public EIPs

At the bottom of the dialog, there are two buttons: 'Back' and 'Start Termination'.

On the "Confirm Termination" tab, after ensuring all details are correct, click **Start Termination** to terminate the cluster.

EMR - Instance List All projects South China (Guangzhou) Cluster Ma

+ Create Cluster Assign to

ID/Name Instance St

Cluster is running

Cluster is running

Cluster is running

Total 3 items

Terminate Cluster ×

Check Info of the Item to Terminate

 Confirm Termination

This resources to be terminated are as follows:

8 instances, [Collapse](#)

| No. | Instance name | Instance ID | Current upper li... | Instance Type |
|-----|---------------|-------------|---------------------|----------------|
| 1 | emr | | 5 | EMR StandardS2 |
| 2 | emr | | 0 | EMR StandardS2 |
| 3 | emr | | 0 | EMR StandardS2 |

16 cloud disks, [Collapse](#)

| No. | Cloud disk name | Cloud disk ID | Attributes | Configuration |
|-----|-----------------|---------------|-------------|---------------------------|
| 1 | | | Data disk | Efficient cloud disk 200G |
| 2 | | | System disk | Efficient cloud disk 50G |

This resources to be retained are as follows:

0 cloud disks
0 public EIPs

Time Cre... | Action

2019-07-23 1... [Details](#) [More](#)

2019-07-22 1... [Details](#) [More](#)

2019-07-18 1... [Details](#) [More](#)

1 / 1 page

Operation Log

Last updated: 2023-12-21 16:15:24

Feature Overview

Operation logs are records of actions executed by users on the Elastic MapReduce console for the cluster, facilitating user review. Actions such as creating a cluster or scaling a cluster, for instance, are all documented within the operation logs.

Instructions

Upon successful creation of the cluster, log in to the [EMR Console](#), enter the cluster list, and **click on the upper left corner of the cluster list to switch regions**. In the cluster list, click on the corresponding cluster **ID/Name** to enter the cluster details page, then click on the **Operation Logs** in the left menu bar.

| Operation Object | Operation | Operation Details | Security Level | Operation Time | Operator |
|------------------|------------------------|--|----------------|---------------------|--------------|
| Component | Deliver Configuration | Cluster Level | | 2021-07-23 15:08:52 | 100013118462 |
| Component | Add Component | Add Component: hive-2.3.7 tez-0.9.2 impala-2.10.0 | | 2021-07-21 10:21:47 | 100013118462 |
| Cluster | Scale Out Cluster | --- | | 2021-06-15 20:12:58 | 100013118462 |
| Cluster | Add Node Specification | Node Type: TASK Node Specification: SA2Model 2Core8G Efficient Cloud Disk 100G * 1 | | 2021-06-15 20:12:47 | 100013118462 |
| Cluster | Create Cluster | --- | | 2021-06-15 20:00:53 | 100010380631 |

Total 5 items

Lines per page 10

1 / 1 page

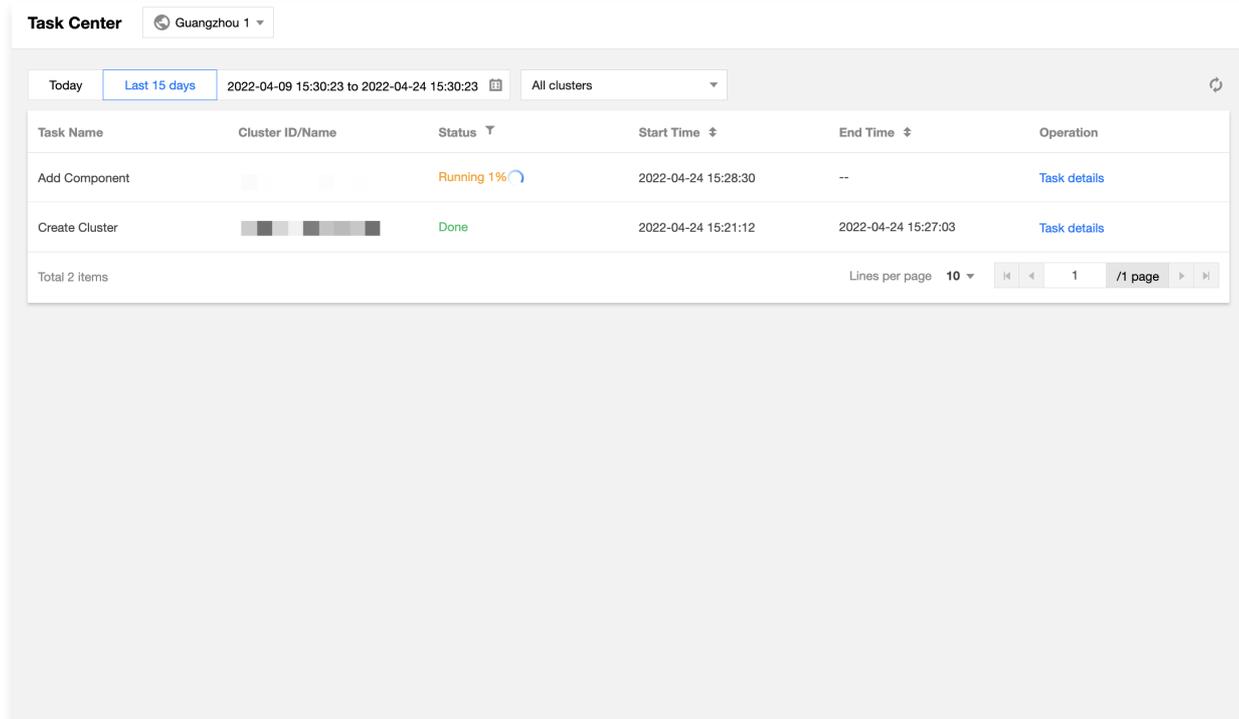
Task Center

Last updated: 2023-12-21 16:15:39

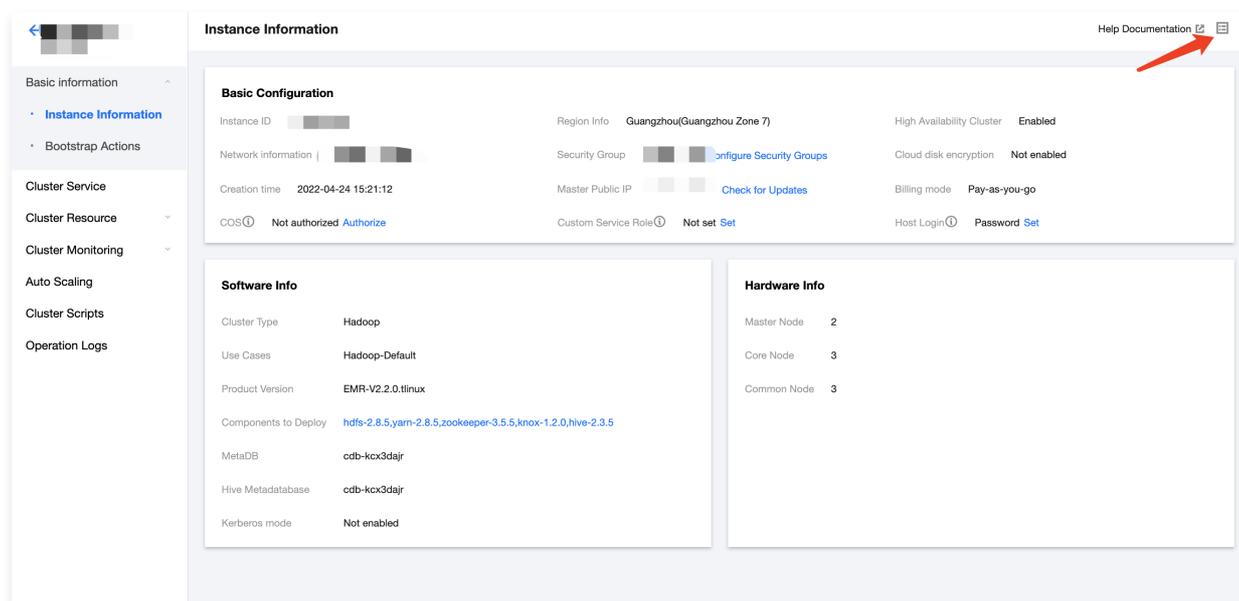
Tencent Cloud's Elastic MapReduce (EMR) console offers a **Task Center** feature, allowing for the comprehensive viewing of the execution status of all cluster tasks. It also supports functionalities such as retrying certain tasks, cancellation, and displaying error information.

Task List

- Log into the [EMR Console](#), click on the **Task Center** in the left menu bar to view the status, start time, end time, and operation information of all tasks for all clusters or a single cluster in the region.



- Access the [EMR Console](#), click on the **Cluster List** in the left menu bar, enter a single cluster and click on **Tasks** in the upper right corner to view the status, start time, end time, and operation information of all tasks in a single cluster.



Task Details

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

| Task Parameters | | | | | |
|-----------------|---|-----------|-------|------------|---|
| ClusterName | ceshi123 | CoreSize | 2 | MasterSize | 1 |
| SoftInfo | zookeeper-3.5.5,hadoop-2.8.5,knox-1.2.0 | SupportHA | 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 | - |

Run details

In the operation steps of the **Task Details** section, you can click on **Run Details** to view the completion status of the corresponding nodes.

Task details ✕

Task Parameters

CoreSize 1 DeployProcess DataNode,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 ▼ | Start Time | End Time | ErrorDetails |
|-------------|---------|------------|----------|--------------|
| No data yet | | | | |

Confirm
Cancel

Managing Service User Management

Last updated: 2023-12-21 16:16:21

Scenario

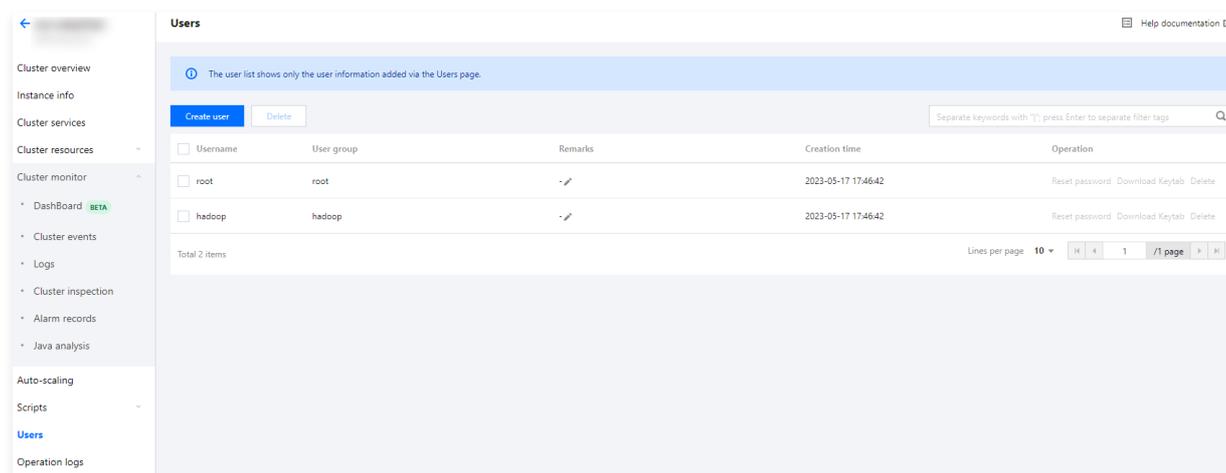
This feature facilitates the management of users within the EMR cluster, enabling swift addition, modification, and deletion of cluster users. Newly added users can be utilized for Hadoop cluster task submissions. This document will guide you through the relevant operations of user management.

Note

- Currently, the user management feature is supported by EMR-V2.6.0, EMR-V3.2.1, and subsequent newer versions.
- Deleting users or resetting passwords may potentially lead to the failure of ongoing tasks, thus necessitating careful execution.

Instructions

1. Log in to the [Elastic MapReduce Console](#), and in the cluster list, click on the cluster **ID/Name** to access the cluster details page.
2. On the cluster details page, click on **User Management**. This page supports functions such as batch user addition, batch user deletion, password reset, and Keytab download.



3. Click on **Create New User** to begin the process of user creation. The username, user group, and password are mandatory fields, while the notes field is optional.

Create user
✕

i 1. Supports creating multiple users at a time.
2. A username or user group name must be 1–30 characters long and contain letters, digits, hyphens, and underscores.

User Info

| Username* | User group* ⁱ | Remarks | Operation |
|--|------------------------------------|--------------------|-----------|
| <input type="text" value="Enter"/> | <input type="text" value="Enter"/> | Max 128 characters | Delete |
| + Add | | | |

Password *
8–30 characters; supports letters, digits, -, _, !, @, #, \$, and %

Confirm password *

Confirm
Cancel

4. Creating a new user supports automatic synchronization with Ranger, with a default sync frequency of one minute. If you need to adjust the sync frequency, you can go to the Ranger configuration management page, modify the `ranger.usersync.sleepinmillisbetween sync cycle` parameter, and distribute the configuration.

Note

For clusters created before July 1, 2023, it is necessary to manually trigger the distribution of the `ranger-ugsync-site.xml` configuration and restart the `EnableUnixAuth` service for user synchronization to take effect. You can go to the Ranger configuration management page, select the `ranger-ugsync-site.xml` configuration file, execute the configuration distribution operation, and then restart the service.

5. Reset Password

On the User Management list page, click on **Reset Password** in the operations on the right side of the user whose password needs to be changed. After entering and confirming the new password, click on **Confirm** to complete the reset.

Reset password
✕

w Resetting the password may cause running tasks to fail. Please confirm this operation.

New password
8–30 characters; supports letters, digits, -, _, !, @, #, \$, and %

Confirm new password

Confirm
Cancel

6. Download Keytab

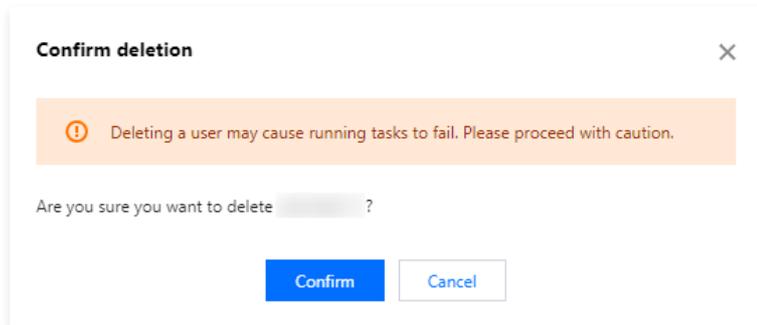
On the User Management list page, select the user for whom you need to download the Keytab, then click on **Download Keytab**. The Keytab can be used to log into the cluster.

Note

Kerberos clusters support the download of Keytab.

7. Delete User

On the User Management list page, click on **Delete** in the operations on the right side of the user you need to remove. Click on **Confirm Deletion** to complete the removal.



Restarting Service

Last updated: 2023-12-21 16:16:42

Feature Overview

Upon modification of component configurations, it is necessary to restart the corresponding services to activate the changes. To ensure minimal or no disruption to business operations during the service restart process, a rolling restart can be implemented. For instances with primary and secondary statuses, the secondary instance will be restarted first, followed by the primary instance. The duration of a rolling restart is longer than a standard restart.

- The console supports service restarts, with the rolling restart method selected by default. If the rolling restart is disabled, all nodes may restart simultaneously, potentially leading to service unavailability. Please choose wisely.
- The failure handling strategy supports two methods: blocking and waiting for resolution upon failure, and continuing processing despite individual node failure.

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding cluster ID/name to access the cluster details page.
2. If you need to restart the entire component service, you can go to the cluster service page, select the component card for the service you wish to restart, and choose **Operation > Restart Service**, or you can enter the component details page and select **More Operations > Restart All Services** from the top right corner of the page. If you need to restart a specific role instance, you can go to the cluster service page, select the component card for the service you wish to restart, and choose **Operation > Role Management**, then select the service roles you wish to restart and click **Restart Service**.
 - When restarting a service from the component card or the 'More Operations' section on the component details page, you need to select the service role to restart, the restart method, whether to use a rolling restart, and the failure handling strategy. When you select 'All' for the service role, the entire component will be restarted.

Restart Service [X]

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 a service from the Role Management page, you only need to select the restart method and the failure handling strategy.

Restart Service ✕

Service name * HDFS

Service Role * NameNode ▾

Restart Range All eligible nodes (total 2)

Restart Method Quick restart mode ▾

Rolling Restart Each restart: nodes, interval: 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 each service component are as follows:

| Component | Service | Restart Mode | Description | Remarks |
|-----------|------------------|----------------------|--|-------------------------------|
| HDFS | NameNode | Quick Restart Mode | Restart via <code>hadoop-daemon.sh stop start namenode</code> | - |
| | NameNode | Secure Restart Mode | In an HA cluster, first perform the <code>saveNameSpace</code> operation on the standby NameNode, then restart via <code>hadoop-daemon.sh stop start namenode</code> . The non-HA cluster is consistent with the Quick Restart Mode. | Supports only Rolling Restart |
| | DataNode | Default restart mode | Restart via <code>hadoop-daemon.sh stop start datanode</code> | - |
| | JournalNode | Default restart mode | Restart via <code>hadoop-daemon.sh stop start journalnode</code> | - |
| | zkfc | Default restart mode | Restart via <code>hadoop-daemon.sh stop start zkfc</code> | - |
| YARN | ResourceManager | Default restart mode | Restart via <code>sbin/yarn-daemon.sh stop start resourcemanager</code> | Supports only Rolling Restart |
| | NodeManager | Default restart mode | Restart via <code>sbin/yarn-daemon.sh stop start nodemanager</code> | - |
| | JobHistoryServer | Default restart mode | Restart via <code>sbin/yarn-daemon.sh stop start historyserver</code> | - |

| | | | | |
|-----------|-----------------------|----------------------|--|-------------------------------|
| | TimeLineServer | Default restart mode | Restart via <code>sbin/yarn-daemon.sh stop start timelineserver</code> | - |
| HBASE | HbaseThrift | Default restart mode | Restart via <code>hbase-daemon.sh stop start thrift</code> | - |
| | HMaster | Default restart mode | Restart via <code>hbase-daemon.sh stop start master</code> | - |
| | HRegionServer | Quick Restart Mode | Restart via <code>hbase-daemon.sh stop start regionserver</code> | - |
| | HRegionServer | Secure Restart Mode | Restart via <code>graceful_stop.sh --restart --reload</code> | - |
| HIVE | HiveMetaStore | Default restart mode | Restart via <code>hcat_server.sh stop start</code> | - |
| | HiveServer2 | Default restart mode | Restart via <code>hive-daemon.sh stop-h2 start-h2</code> | - |
| | HiveWebHcat | Default restart mode | Restart via <code>webhcat_server.sh stop start</code> | - |
| PRESTO | PrestoCoordinator | Default restart mode | Restart via <code>bin/launcher stop start</code> | Supports only Rolling Restart |
| | PrestoWorker | Default restart mode | Restart via <code>bin/launcher stop start</code> | - |
| ZOOKEEPER | QuorumPeerMain | Default restart mode | Restart via <code>bin/zkServer.sh stop start</code> | - |
| SPARK | SparkJobHistoryServer | Default restart mode | Restart via <code>sbin/stop-history-server.sh sbin/start-history-server.sh</code> | - |
| HUE | Hue | Default restart mode | Restart via <code>build/env/bin/start.sh</code> and <code>build/env/bin/stop.sh</code> | - |
| OOZIE | Oozie | Default restart mode | Restart via <code>oozied.sh stop start</code> | - |
| STORM | Nimbus | Default restart mode | Restart via <code>bin/storm-daemon.sh nimbus stop start</code> | - |
| | Supervisor | Default restart mode | Restart via <code>bin/storm-daemon.sh supervisor stop start</code> | - |
| | Logviewer | Default restart mode | Restart via <code>bin/storm-daemon.sh nimbus stop start</code> | - |

| | | | | |
|---------|---------------|----------------------|--|---|
| | Ui | Default restart mode | Restart via bin/storm-daemon.sh nimbus stop start | - |
| RANGER | Ranger | Default restart mode | Restart via sbin/ranger-daemon.sh stop start | - |
| ALLUXIO | AlluxioMaster | Default restart mode | Restart via bin/alluxio-stop.sh master and bin/alluxio-start.sh master | - |
| | AlluxioWorker | Default restart mode | Restart via bin/alluxio-stop.sh worker and bin/alluxio-start.sh worker | - |
| GANGLIA | Httpd | Default restart mode | Restart via /sbin/service http stop start | - |
| | Gmetad | Default restart mode | Restart via /sbin/service gmetad stop start | - |
| | Gmond | Default restart mode | Restart via /sbin/service gmon stop start | - |

Starting/Stopping Services

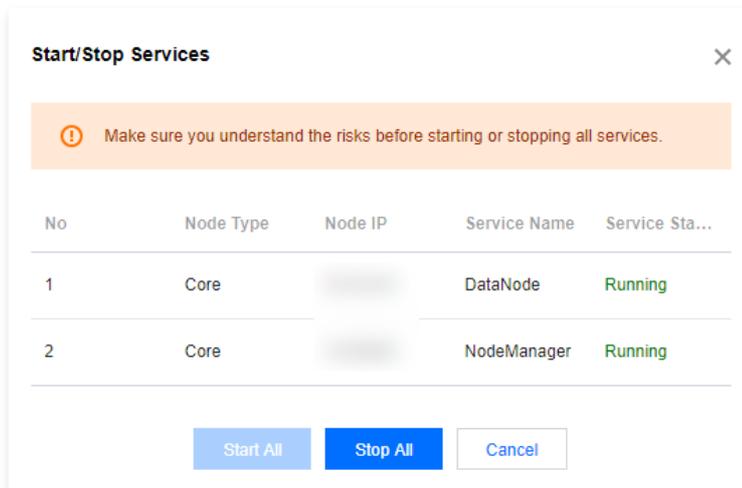
Last updated: 2023-12-21 16:16:58

Feature Overview

The start-stop service allows for the activation and deactivation of all services on a single node. This document provides a comprehensive guide on how to operate the start-stop service via the Elastic MapReduce console.

Instructions

1. Log into the [EMR Console](#), and in the cluster list, click on the corresponding cluster **ID/Name** to access the cluster details page.
2. Within the cluster details page, select **Cluster Resources**, then click on **Resource Management** to navigate to the resource management page. Choose the node that requires activation or deactivation, then click on **More > Start/Stop Service**.
3. On the Start/Stop Service page, confirm the relevant service information.



4. Upon verifying that the information is accurate, click on **Start All Services** or **Stop All Services** to activate or deactivate all services on the node.

Role Management

Last updated: 2023-12-21 16:17:15

Feature Overview

Role Management offers operational functions at the role service level, including rebooting, pausing, and maintenance operations, which are implemented at the specific node dimension. Role status supports monitoring to facilitate real-time understanding of the role process status.

Definitions

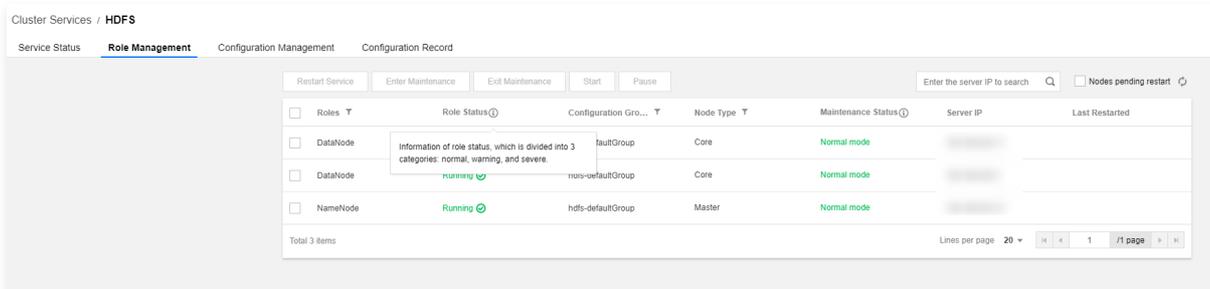
- **Reboot:** Sequentially performs a rolling restart on the selected service roles according to the nodes.
- **Pause:** This will halt the nodes of the selected service roles. The **Start** feature can be used for recovery.
- **Maintenance:** This halts process guarding for the nodes of the selected service roles. When the process is in an abnormal state due to various reasons, no alarms will be triggered or automatic recovery will occur, making it suitable for node debugging. The **Exit Maintenance** feature can be used for recovery.

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding cluster ID/Name to access the cluster details page.
2. In the cluster details page, select **Cluster Services** and click on the top right corner of the corresponding component card **Operation > Role Management**, taking HDFS as an example.
3. The **Role Management** list displays information such as the current health status, operational status, configuration group, node type, maintenance status, host IP, and most recent restart time of the service role. After selecting a role, you can perform **Restart**, **Maintenance**, **Start**, and **Pause** operations.

Note

The Health Status column displays the current operational status of the role, the Operation Status column shows the user's actions, and the Maintenance Status column indicates whether the current role is in maintenance mode.



Health status

| Health status | Description |
|---------------|---|
| Good | Port detection responds within 5 seconds. |
| Risky | Port check is responded within 5s-10s. |
| Unavailable | Port detection has not responded for more than 10 seconds. |
| Not checked | Processes of roles that have entered maintenance mode or have ceased operations are not subjected to detection. |
| Unknown | Detector exception or device downtime |

Client Management

Last updated: 2023-12-21 16:17:37

Feature Overview

Client Management offers information on the deployment of components, enabling users to view the configuration status, configuration group, node type, and node IP information on the Client Management page.

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding **Cluster ID/Name** to access the cluster details page.
2. Click on **Client Management** to navigate to the Client Management page.

| Client | Configuration status | 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

Configuration Management

Configuration Management

Last updated: 2023-12-21 16:18:03

Feature Overview

Configuration management supports the modification of key configuration parameters for commonly used open-source components such as HDFS, YARN, HIVE, SPARK, etc. It allows for the adjustment of service configurations based on actual needs at the cluster level, node level, and configuration group level. This document provides an overview of how to configure various service parameters through the console.

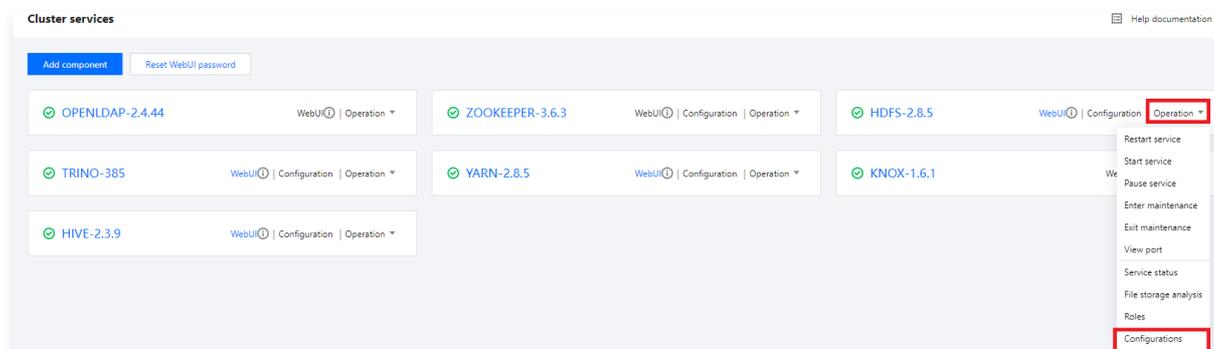
Note

- When deleting custom configuration files in console configuration management, for security reasons, the client will not synchronize the deletion action.
- If you have a scenario where you need to delete custom configuration files on the client, you can utilize the capabilities of cluster scripts to perform operations in bulk on existing clusters.
- After modifying the configuration, if the **Role Management > Configuration Status** changes to "Configuration Expired", you will need to restart this service for the configuration to take effect.

Instructions

Edit Configuration Item

1. Log in to the [EMR Console](#), select the corresponding cluster from the cluster list, and click on **Cluster Services** to enter the cluster service list.
2. In the cluster service list, select **Configuration** or **Operation > Configuration Management** in the panel of the service whose configuration is to be modified.



3. Upon entering the Configuration Management page, select the configuration dimension range as needed, with the cluster dimension selected by default.
 - If you need to configure parameters for all nodes of the selected service within the cluster, select the cluster dimension.
 - If you need to configure parameters for multiple nodes of the selected service, you can create a configuration group and select the corresponding configuration group.
 - If you need to configure parameters for a specific node of the selected service, select the node dimension.

Cluster services / HDFS

Service status File storage analysis Roles Client management **Configurations**

For configuration adjustment in the console, 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 to dimension-based comparisons](#)

Configuration filter

Parameter name or value

Dimension

Cluster level

Range [Clear all](#)

NameNode
 DataNode
 Client

Category [Clear all](#)

About Address Or Port
 Advanced
 Basic
 CHdfs
 Compression
 About Cos
 About Data Store
 About Disk
 group
 About HA
 Host

core-site.xml hadoop-env.sh hdfs-site.xml https-env.sh https-site.xml log4j.properties

Restart service Edit configuration

| File | Parameter | Value | Description |
|------------------|---|---------------------|------------------|
| log4j.properties | AuditLoggerLevel | WARN | AUDIT 日志级别 |
| log4j.properties | Logdir | /data/emr/hdfs/logs | 日志目录 |
| log4j.properties | RootLoggerLevel | WARN | root.logger 日志级别 |
| log4j.properties | SecurityLoggerLevel | WARN | Security 日志级别 |
| log4j.properties | mkimpalpath | /var/lib/hadoop | - |
| core-site.xml | delegation.token.identifier.serialization.version | 0 | - |
| core-site.xml | delegation.token.identifier.transition.version | 0 | - |
| core-site.xml | emr.cfs.group.id.map | root:0/hadoop500 | cfs user map |
| core-site.xml | emr.cfs.io.blocksize | 1048576 | cfs块大小 |
| core-site.xml | emr.cfs.user.id.map | root:0/hadoop500 | cfs user map |

- If you wish to search for a specific configuration item or narrow down the search range for configuration items, you can filter through the left-hand side filter.
- Select the configuration file as needed, click on **Edit Configuration** to enter the editing state, and perform operations such as adding, editing, and deleting configuration items as required.
 - Select the parameter you wish to modify and input the new parameter value. If necessary, you can click on **Restore** to revert to the original value, or click on the default value to revert to the system recommended default value.
 - Some parameters support deletion. If you need to delete this configuration, select **Delete > Confirm**.
 - If the file does not contain the parameter you wish to configure, you can click on **Add Configuration Item** to open the new configuration item pop-up page, and fill in the parameter name and parameter value for the new configuration item.
- After confirming that everything is correct, click on **Save Configuration**. Once the configuration has been successfully issued, click on **Restart Service** to complete the modification of the configuration item settings.

Note

- If you are modifying the service process configuration, you will need to restart the service after saving for the configuration to take effect.
- If you are modifying the client configuration, the changes will take effect immediately after saving, without the need to restart the service.

Add configuration file

- Log in to the [EMR Console](#), select the corresponding cluster from the cluster list, and click on **Services** to access the cluster service list.
- In the cluster service list, select **Operations > Configuration Management** in the upper right corner of the service panel you wish to modify.
- If the configuration file you wish to configure is not available, you can click on **Add Configuration File** on the right to access the configuration file setup page and fill in the configuration file.
- After clicking on **Save Configuration**, the parameters will be distributed and the configuration file name in the configuration file list will be updated.

Add configuration file ✕

File name * Configuration path *

Permission property * User group *

Username *

Content *

5. Once the custom configuration file has been distributed and activated, it supports modification and deletion operations.

Configuration Dimension Comparison

After modifying the configuration across multiple dimensions, if you need to compare the differences in configuration parameters between the cluster dimension, configuration group dimension, and node dimension, you can click on **Switch to Dimension Comparison List** to view the configuration differences for each dimension. If you need a certain dimension's configuration to be consistent with its parent dimension, you can overwrite the difference value and save the changes.

Cluster services / HDFS More ▾ | Help documentation 📄

Service status File storage analysis Roles Client management **Configurations**

🔔 For configuration adjustment in the console, 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 to dimension-based comparisons

Configuration filter

Parameter name or value

Dimension 🔍

Cluster level

Range Clear all

NameNode

DataNode

Client

Category Clear all

About Address Or Port

Advanced

Basic

CHdfs

Compression

All core-site.xml hadoop-env.sh hdfs-site.xml https-env.sh https-site.xml log4j.properties +

Restart service

| File | Parameter | Value | Description |
|------------------|---|---------------------|------------------|
| log4j.properties | AuditLoggerLevel | WARN | AUDIT 日志级别 |
| log4j.properties | Logdir | /data/emr/hdfs/logs | 日志目录 |
| log4j.properties | RootLoggerLevel | WARN | root logger 日志级别 |
| log4j.properties | SecurityLoggerLevel | WARN | Security 日志级别 |
| log4j.properties | mkimppalpath | /var/lib/hadoop | - |
| core-site.xml | delegation.token.identifier.serialization.version | 0 | - |
| core-site.xml | delegation.token.identifier.transm | 0 | - |

© 2013–2023 Tencent Cloud. All rights reserved.

Page 105 of 295

Configuration Status

Last updated: 2023-12-21 16:18:17

Feature Overview

Upon modification and distribution of component configurations, one can discern whether there are instances of service configuration distribution failures or expired configurations through the configuration status. It supports the viewing of detailed information regarding configuration failures and expired configurations, thereby facilitating an understanding of specific configuration changes for each role.

Nomenclature Explanation

Synchronized: No configuration modifications, or the service has been restarted after configuration changes.

Configuration Expired: Configuration changes have been made and distributed, but the service has not been restarted, resulting in the configuration not taking effect.

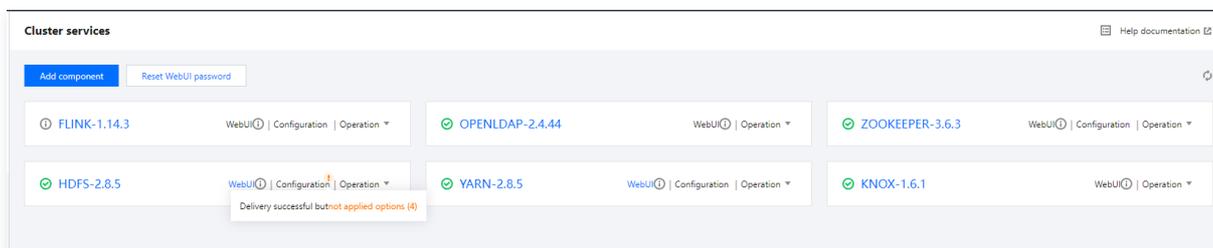
Configuration Failed: After configuration changes, there are nodes where the configuration distribution has failed.

Description

- The configuration status of role instances can be categorized into three types: Synchronized, Configuration Failed, and Configuration Expired.
 - The initial configuration status is Synchronized. However, if the distribution of role process configurations fails, the role configuration status transitions to Configuration Failed.
 - Upon successful distribution of the configuration, the role configuration status transitions to Configuration Expired.
 - Provided that the configuration has not failed, upon successful service restart, the role configuration status transitions to Synchronized.
- The configuration status of the client can be categorized into two types: Synchronized and Configuration Failed.
 - The initial configuration status is Synchronized. However, if the distribution of client configurations or role process configurations fails, the client configuration status transitions to Configuration Failed.
 - Upon successful distribution of the configuration, the client configuration status transitions to Synchronized.
- The configuration status of the component is displayed in two forms: Configuration Failed and Configuration Expired.
 - The initial status is not displayed. However, if the role instance or client configuration fails, the component configuration status is displayed as Configuration Failed.
 - When the role instance configuration expires, the component configuration status is displayed as Configuration Expired.
 - When both Configuration Failed and Configuration Expired exist simultaneously for the role instance and client, the component configuration status prioritizes displaying Configuration Failed.

Instructions

- Log in to the [EMR Console](#), and in the cluster list, click on the corresponding **Cluster ID/Name** to enter the cluster details page.
- In the cluster details page, select **Cluster Services**, and choose the component card with a configuration status of Configuration Failed or Configuration Expired.



- For instance, in the case of an expired configuration that has not taken effect, click on Not Effective to view the details of the ineffective configuration.

Cluster services

- FLINK-1.14.3
- HDFS-2.8.5
- OPENLDAP-2.4.44
- YARN-2.8.5

Configuration status

HDFSDelivery successful but not applied items (4)

- To apply successfully delivered configurations, restart nodes by role in [Role management](#)

Expired configurations

| Role | Configuratio... | Parameter | Node IP | Current applied ... | Value after change |
|----------|-----------------|------------------|---------|---------------------|--------------------|
| NameNode | hadoop-env.sh | test | | | test |
| DataNode | hadoop-env.sh | test | | | test |
| DataNode | hadoop-env.sh | test | | | test |
| NameNode | hadoop-env.sh | AuditLoggerLevel | | WARN | ERROR |

4. After confirming the expired configuration that needs to take effect, you can restart the corresponding service in the respective role management.

Cluster services / HDFS

Service status | File storage analysis | **Roles** | Client management | Configurations

Restart service | Start | Pause | Enter maintenance | Exit maintenance

| Role | Health status | Operation status | Configuration sta... | Configuration |
|--|---------------|------------------|----------------------|----------------|
| <input checked="" type="checkbox"/> NameNode | Good | Started | Expired | hdfs-master-de |
| <input type="checkbox"/> DataNode | Good | Started | Expired | hdfs-core-defa |
| <input type="checkbox"/> DataNode | Good | Started | Expired | hdfs-core-defa |

Total 3 items

Configuration status

HDFSDelivery successful but not applied items (4)

- To apply successfully delivered configurations, restart nodes by role in [Role management](#)

Expired configurations

| Role | Configuratio... | Parameter | Node IP | Current applied ... | Value after change |
|----------|-----------------|------------------|----------------|---------------------|--------------------|
| NameNode | hadoop-env.sh | test | 192.168.88.53 | | test |
| DataNode | hadoop-env.sh | test | 192.168.88.120 | | test |
| DataNode | hadoop-env.sh | test | 192.168.88.9 | | test |
| NameNode | hadoop-env.sh | AuditLoggerLevel | 192.168.88.53 | WARN | ERROR |

Configuration Rollback

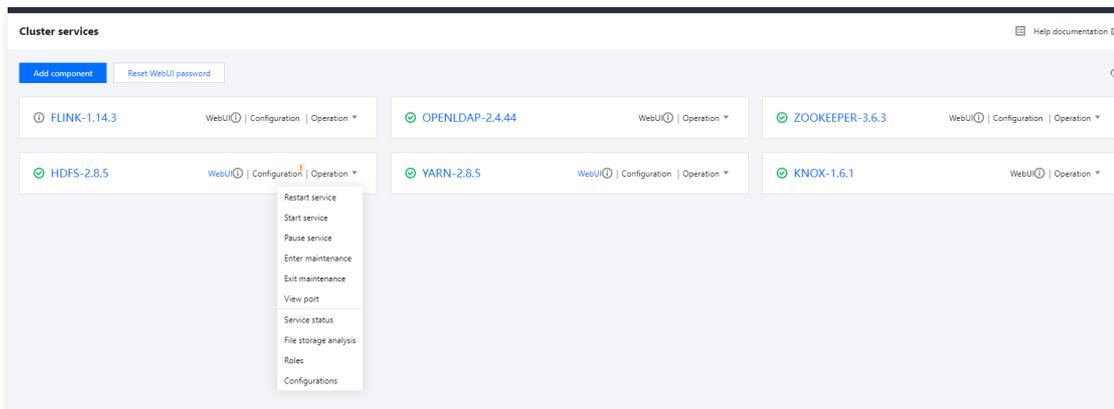
Last updated: 2023-12-21 16:18:31

Feature Overview

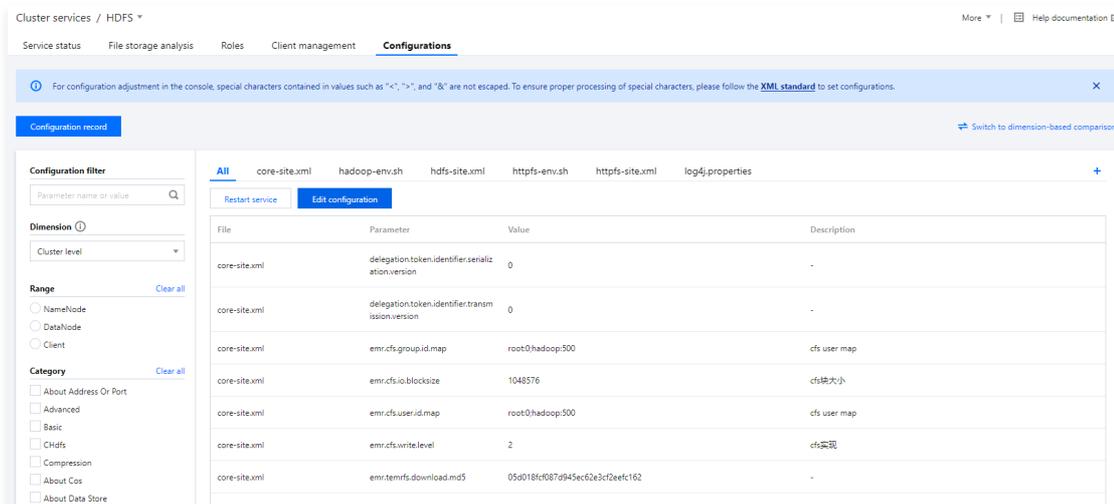
Elastic MapReduce facilitates the addition, modification, and deletion of component parameters via the console, allowing for configuration rollbacks. This document provides a comprehensive guide on how to rollback component parameter configurations through the console.

Instructions

1. Log in to the [EMR Console](#), select the appropriate cluster from the cluster list, and click on **Details** to access the cluster details page.
2. In the cluster details page, select **Cluster Services** and click on the top right corner of the corresponding component card **Operations > Configuration Management**, using HDFS as an example.



Navigate to the Configuration Management page:



3. On the Configuration Management page, select **Configuration History**. By default, the current component's configuration history is displayed, but you can switch to view the configuration history of all components. Click on **Details** to compare the parameter values before and after the configuration changes. Click on **Rollback** to rollback the parameter configuration changes for that record. Select **Rollback > Confirm Rollback**. After a successful rollback and component restart, wait a moment for the rollback to take effect.

Note

Adding, modifying, and deleting configuration items support rollback, while adding new configuration files and deleting added configuration files do not support rollback.

Configuration record ×

| Service ▾ | Configuration ... ▾ | Level ▾ | Modified at | Reason | Operation |
|-----------|---------------------|---------------|---------------------|---------|--|
| HDFS | hadoop-env.sh | Cluster level | 2023-11-23 15:22:04 | test | Details Rollback |
| HDFS | hadoop-env.sh | Cluster level | 2023-11-23 15:12:40 | test啊啊啊 | Details Rollback |

Total 2 items Lines per page 10 ▾ ⏪ ⏩ 1 / 1 page ⏪ ⏩

Configuration Group Management

Last updated: 2023-12-21 16:18:45

Feature Overview

Configuration groups are utilized to manage and group nodes of varying specifications or purposes that deploy the same service. The management rules for service configurations using configuration groups are as follows:

- Modifying configuration items at the cluster level will, by default, override the configuration items at the group and node levels, provided these items have not been independently modified at these levels.
- When modifying configuration items at the configuration group level, if the configuration item has not been independently modified within the group nodes, it will, by default, override the configuration items of all nodes within the group. Subsequent modifications at the cluster level will not override this configuration group's settings.
- When modifying configuration items at the node level, only the configuration items of that particular node are updated. Subsequent modifications disseminated from the cluster level and configuration group level will not override this node's configuration.

Note

A node can only belong to one configuration group, ensuring no overlap of nodes between different groups.

Instructions

1. Log in to the [EMR Console](#), select the appropriate cluster from the cluster list, and click on **Details** to access the cluster details page.
2. In the cluster details page, select **Cluster Services** and click on the top right corner of the corresponding component card **Operations > Configuration Management**, using HDFS as an example.
3. Click on **Configuration Management**, set the "Dimension" to Configuration Group, select the corresponding **Configuration Group**, and click on **Manage Configuration Groups** to access the "Manage Configuration Group" page.

Note

Upon the completion of cluster creation, each component will generate several default configuration groups divided by node type. These default groups only support modification and cannot be deleted. The default configuration groups inherit the cluster dimension configuration upon initialization. If needed, new configuration groups can be added, ensuring that each component does not exceed 15 configuration groups.

The screenshot shows the 'Configuration record' page in the EMR console. It features a 'Configuration filter' on the left with a search bar and a 'Dimension' dropdown set to 'Configuration group'. Below the filter is a 'Range' section with a radio button for 'NameNode'. The main area displays a table of configuration groups for the 'core-site.xml' file. The table has columns for 'File', 'Parameter', 'Value', and 'Description'. The parameters listed are 'delegation.token.identifier.serialization.version' and 'delegation.token.identifier.transition.version', both with a value of '0'. A third parameter, 'emr.cfs.group.id.map', has a value of 'root0/hadoop:500' and a description of 'cfs user map'.

| File | Parameter | Value | Description |
|---------------|---|------------------|--------------|
| core-site.xml | delegation.token.identifier.serialization.version | 0 | - |
| core-site.xml | delegation.token.identifier.transition.version | 0 | - |
| core-site.xml | emr.cfs.group.id.map | root0/hadoop:500 | cfs user map |

4. On the Configuration Group Management page, you can view all configuration groups of the current component. Additionally, you can add, modify, and delete configuration groups.

Note

- When adding a new configuration group, you can choose to inherit the configuration group settings. If you do not select to inherit from a configuration group, the cluster dimension configuration parameters will be inherited by default.
- Upon deletion of a custom configuration group, the nodes of that group will be relocated to the default configuration group of the same node type.

Cluster services / HDFS / Configuration group Help documentation

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, so that nodes in the same configuration group have the same configuration parameters. After a cluster is created, a default configuration group for each component is generated and initialized to inherit the cluster-level configuration. You can add configuration groups as needed, with a maximum of 15 configuration groups allowed for each component.

[Add configuration group](#)

- hdfs-core-defaultGroup** [Modify](#)
Configuration group description: Default configuration group for hdfs core nodes
Node: [Show More](#)
- hdfs-master-defaultGroup** [Modify](#)
Configuration group description: Default configuration group for hdfs master nodes
Node: [Show More](#)
- hdfs-task-defaultGroup** [Modify](#)
Configuration group description: default configuration group for hdfs task nodes
Node: [Show More](#)
- hdfs-core-defaultGroup-322** [Modify](#) [Delete](#)
Configuration group description:
Node: [Show More](#)

5. If the configuration group is inconsistent with the cluster dimension configuration, you can directly click on **Save** in the cluster dimension to forcibly overwrite the configuration group settings. Similarly, if there are node configurations in the configuration group that are inconsistent with the current configuration group settings, you can directly click on **Save** in the configuration group dimension to forcibly overwrite the corresponding node configurations.

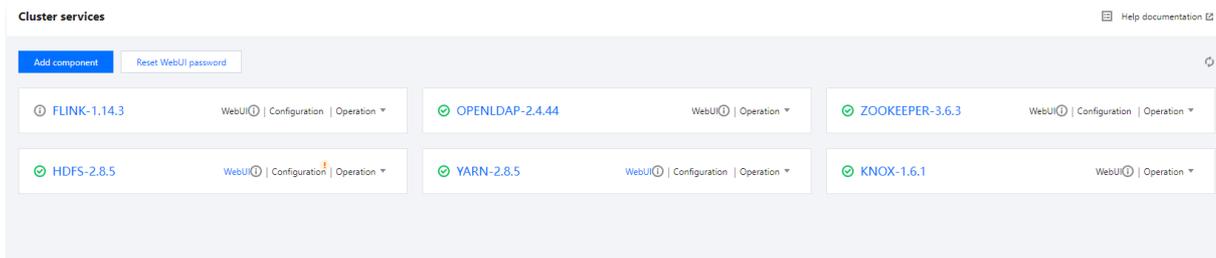
Services

Last updated: 2023-12-21 16:19:00

The service list displays the services installed in the cluster, as well as the health status, configuration status, version information, etc. of the services. It also provides convenient service operation management tools, including general service operations and some service-specific command-type operation management.

Service Health Status

The Service Health Status indicates whether the current service is operating normally, which is aggregated from the health status of various roles.



There are primarily four types of Service Health Status: Good, At Risk, Unavailable, and Unknown or Undetected. Different status types correspond to different color displays.

| Component Health Status | Health status description | Status Aggregation Rules |
|---------------------------------|--|---|
| Green: Good | The service is operating normally. | The health status of all role instances is good. |
| Orange: Potential risks present | The service is available, but some role instances are in an unhealthy state or have potential risks that require attention and resolution. | The health status of some instances of a certain role in this component is unavailable or has potential risks. For instance, HDFS has one NameNode role instance and two DataNode role instances. One of the DataNode role instances is in an unhealthy state, while the other DataNode role instance and the NameNode role instance are in good health. Therefore, the health status of HDFS is considered to have potential risks. |
| Red: Unavailable | The service is unavailable, and the health status of all instances of a certain role is unavailable. Please address this issue promptly. | The health status of all instances of a certain role in this component is unavailable. For example, HDFS has one NameNode role instance and two DataNode role instances. Both DataNode role instances are in an unhealthy state, while the NameNode role instance is in good health. Therefore, the health status of HDFS is considered unavailable. |
| Grey: Unknown or Undetected | The health status of the service is unknown or undetected. For components without a process, the health status is undetected. For components with a process, the health status is undetected if they are in maintenance mode or their operational status has stopped. If the role instance health status information cannot be correctly obtained, the health status is unknown. If the business operations are found to be problem-free upon investigation, there is no need for concern. | <ol style="list-style-type: none"> The health status of all role instances in this component does not pose a risk or is unavailable, and at least one role instance's health status is unknown. For example, HDFS has one NameNode role instance and two DataNode role instances. The health status of one DataNode role instance is unknown, while the other DataNode role instance and the NameNode role instance are in good health. Therefore, the health status of HDFS is considered unknown. The health status of all role instances in this service is undetected. When all role instances of a service enter maintenance mode or their operational status has stopped, their health status is not detected. For components without a process, their health status is not detected, such as Iceberg, Hudi, Flink, and so on. |

Service operation

Common service operations include service restart, start, pause, enter/exit maintenance mode, and port viewing. Directive-based service operations include HDFS NameNode primary-secondary switch, HDFS data balancing, YARN ResourceManager primary-secondary switch, YARN queue refresh, etc. The operation instructions are as follows:

| Service operation | Note |
|---|---|
| HDFS NameNode Primary-Secondary Switch | Commonly referred to as NN primary-secondary switch, it involves transitioning the currently active NameNode to a standby state, and converting the originally standby NameNode to an active state. |
| HDFS Data Balancing | This operation is typically executed when a new DataNode is added. It ensures a uniform data distribution, prevents hotspot issues, and balances the read-write load across the cluster. |
| Yarn ResourceManager Primary-Secondary Switch | Commonly referred to as RM primary-secondary switch, it involves transitioning the currently active ResourceManager to a standby state, and converting the originally standby ResourceManager to an active state. The RM primary-secondary switch is only permissible when <code>yarn.resourcemanager.ha.automatic-failover.enabled</code> is disabled. If the RM primary-secondary switch does not appear in the Yarn card operation dropdown menu, please locate <code>yarn.resourcemanager.ha.automatic-failover.enabled</code> in the Yarn Configuration Management - Configuration file <code>yarn-site.xml</code> , and disable it. |
| Yarn Queue Refresh | When new content is added or existing content is updated in <code>capacity-scheduler.xml</code> or <code>fair-scheduler.xml</code> , this operation enables these changes to take effect in the ResourceManager. Please be cautious not to delete the effective queues defined in <code>capacity-scheduler.xml</code> or <code>fair-scheduler.xml</code> . |
| Ranger Metadata Database Modification | When it is necessary to modify the underlying database of Ranger, the <code>conf/install.properties</code> file needs to be altered, followed by the execution of the <code>setup.sh</code> script locally. This operation offers a one-click configuration of the metadata database, preventing service anomalies that may occur due to missed configurations when users change the Ranger metadata database address. This operation currently only supports the MySQL database, and the test connection function is solely for testing the administrator user's connection. This operation synchronizes the database information to the local <code>ranger-admin-site.xml</code> configuration file, but it does not synchronize and modify the content of <code>ranger-admin-site.xml</code> in the configuration management. If users modify and issue <code>ranger-admin-site.xml</code> on the configuration management page due to additional requirements, it will result in the database information being overwritten, thereby leading to anomalies. |

To perform service operations, you can follow the steps below:

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding **Cluster ID/Name** to enter the cluster details page.
2. In the cluster details page, select **Cluster Services** and choose the component card that needs to be operated.
3. For instance, to switch between HDFS NN primary and secondary, select **HDFS Component Card Operation > NN Primary-Secondary Switch** in the cluster services to perform the primary-secondary switch operation.

NN主备切换
✕

📘 后台命令: `hdfs haadmin -failover`

服务名称 **HDFS**

服务角色 **NameNode**

操作参数

请选择NameService * HDFS17500

操作范围 默认节点

操作原因 *

请输入不超过200字

确定
取消

YARN Resource Scheduling Overview

Last updated: 2023-12-21 16:19:34

Feature Overview

YARN resource scheduling offers an interactive management capability for YARN resource queue scheduling, providing a more convenient operation than file-based configuration management. Currently, it supports two types of scheduling configurations: Fair Scheduler and Capacity Scheduler.

Supports and Limits

1. The resource scheduler defaults to the Fair Scheduler, maintaining consistency with the parameters in the fair-scheduler.xml configuration file of the YARN component found on the resource scheduling page. When switching to the Capacity Scheduler, the relevant configuration parameters in the capacity-scheduler.xml configuration file of the YARN component also remain consistent with those on the resource scheduling page.
2. After switching the scheduler type or setting the policy on the resource scheduling page, you must click on **Deploy to take effect** to initiate the corresponding policy configuration. This operation may restart the ResourceManager.
3. If users have already enabled label scheduling in configuration management, a synchronization operation will be performed when capacity scheduling is enabled.
4. Resource pools support nested sub-pools, with the sub-pools being subject to the rule restrictions set by their parent pool.
5. When disabling the resource scheduler, selecting **Synchronize** to modify the configuration will overwrite the corresponding scheduler's configuration file and parameters in its configuration management.

Configuring Capacity Scheduler

Last updated: 2023-12-21 16:19:50

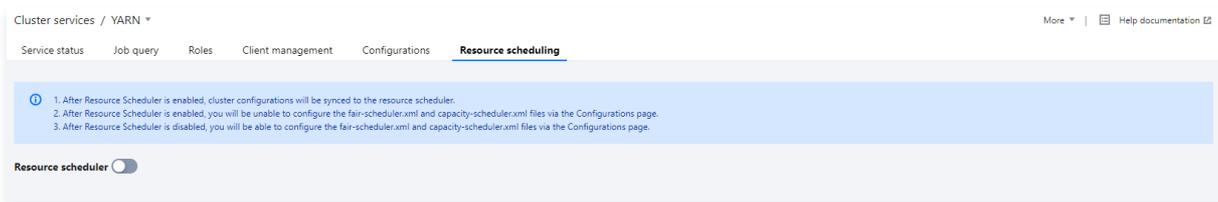
Feature Overview

The Capacity Scheduler is a hierarchical resource organizer that allows multiple users to share cluster resources through multi-level resource constraint conditions.

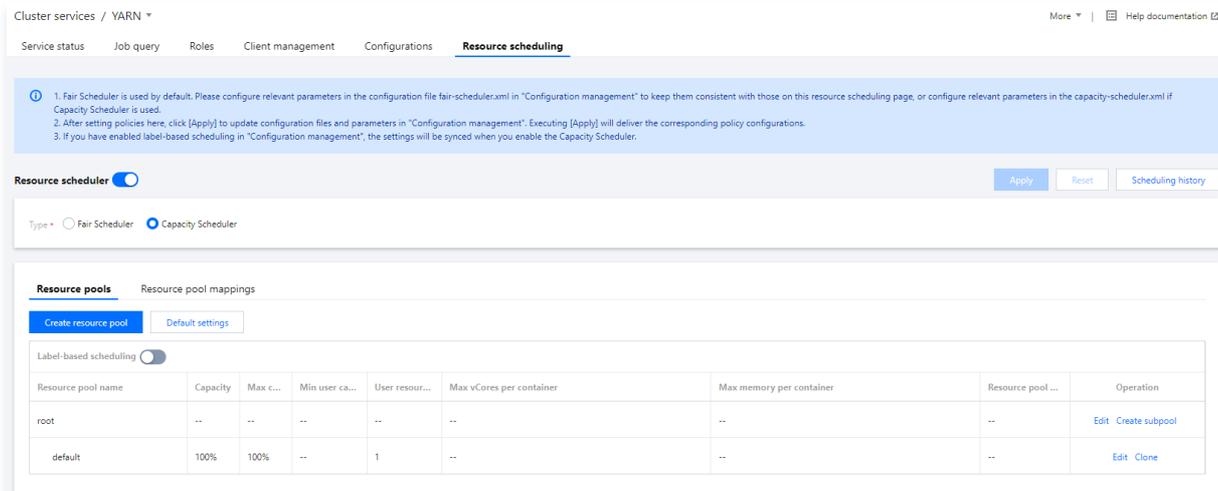
Instructions

Create resource pool

1. Log in to the [EMR Console](#), select the corresponding Hadoop cluster from the cluster list, and click on **Details** to access the cluster details page.
2. In the cluster details page, select **Cluster Services > Yarn Component Card** at the top right corner, then choose **Operations > Resource Scheduling** to access the resource scheduling page.



3. Click on the **Resource Scheduler Switch**. Once the switch is turned on, you can proceed with the related scheduler configurations.
4. To create a new Capacity Scheduler, select the Capacity Scheduler as the scheduling policy type to access its configuration page. Click on **Create Resource Pool** to create a new resource pool. You can edit existing resource pools, create sub-pools, clone, and perform other operations. You can also click on **Default Settings** to set the number of delay scheduling for capacity scheduling.



Create a new resource pool:

Create resource pool ✕

Basic settings

Resource pool name ⓘ

Resource limits

Label settings

| Tag | Capacity | Max capacity | Op... |
|----------------------|--------------------------------------|--------------------------------------|-------|
| <DEFAULT_PARTITION ▾ | <input type="text" value="Enter"/> % | <input type="text" value="Enter"/> % | |

Min user capacity ⓘ %

User resource factor ⓘ %

Max memory per container ⓘ MB

Max vCores per container ⓘ

Resource pool status ⓘ RUNNING STOPPED

App limits

Max apps ⓘ - 10000 +

Max resources for AM ⓘ %

Resource pool priority ⓘ

Access control

Submission ⓘ Allow all users to submit jobs to this resource pool
 Allow the following users to submit jobs to this resource pool

Management ⓘ Allow all users to manage this pool
 Allow the following users to manage this pool

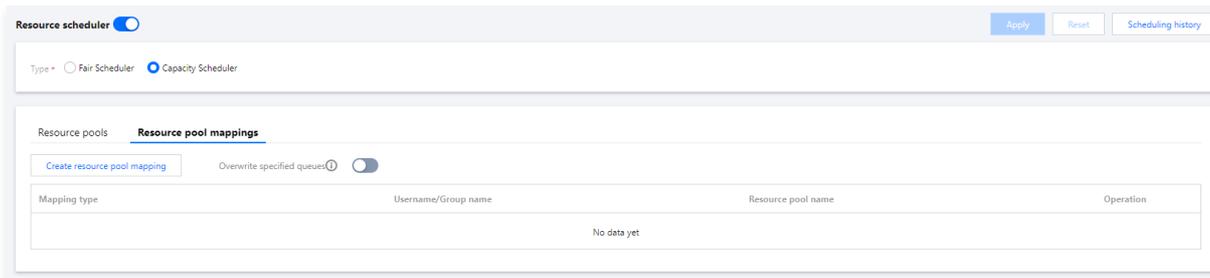
Field and Configuration Item Comparison Table:

| Field | Corresponding Parameter Name | Description |
|--------------------|--|---|
| Resource pool name | yarn.scheduler.capacity.<queue-path>.queues</queue-path> | The name of the resource pool or queue. |
| Label settings | - | Establish the specific labels that the queue can access. |
| Capacity | yarn.scheduler.capacity.<queue-path>.capacity</queue-path> | The size of the resources available, the total capacity of the sub-pools of the same parent resource pool is 100, and the resources that can be used = parent resource pool * capacity%. If this queue requires more resources than this ratio, and other queues have idle resources, it can occupy more resources than this ratio. |
| Max Capacity | yarn.scheduler.capacity.<queue-path>.maximum-capacity</queue-path> | The upper limit of resource usage (percentage) for the queue. Due to the existence of resource sharing, the amount of resources a queue uses may exceed its capacity, and the maximum amount of resources used can be limited by this parameter. |
| Default label | yarn.scheduler.capacity.<queue- | If no node label is specified in the resource request, the |

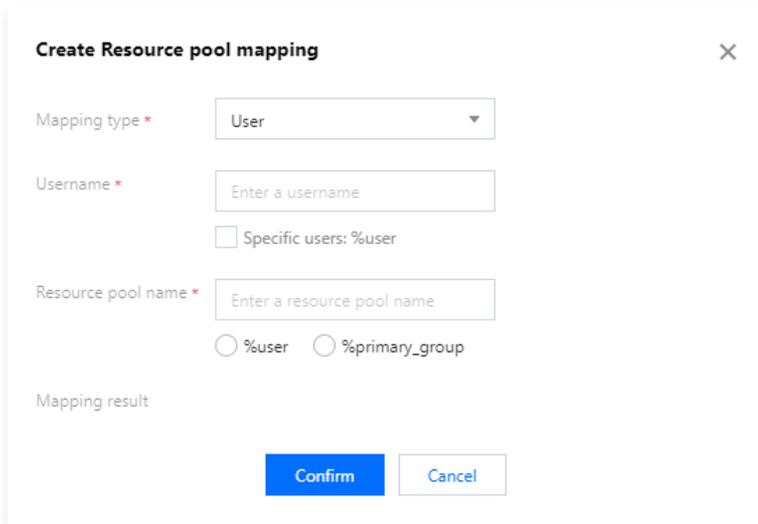
| | | |
|--|--|---|
| expression | <code>path>.default-node-label-expression</queue-path></code> | application will be submitted to the container this value corresponds to. This value defaults to null, indicating the application will be assigned with a container in an untagged node. |
| Min user capacity | <code>yarn.scheduler.capacity.<queue-path>.minimum-user-limit-percent</queue-path></code> | Minimum resource guarantee per user (percentage). At any given time, the amount of resources each user can use in a queue is limited. When multiple users' applications are running simultaneously in a queue, the amount of resources each user uses fluctuates between a minimum and maximum value. The minimum value depends on the number of applications running, while the maximum value is determined by the <code>minimum-user-limit-percent</code> . |
| User resource factor | <code>yarn.scheduler.capacity.<queue-path>.user-limit-factor</queue-path></code> | Maximum resource usage per user (percentage). For instance, if this value is set to 30, then at any given time, the resources used by each user cannot exceed 30% of the queue's capacity. |
| Maximum memory allocation for Container | <code>yarn.scheduler.capacity.<queue-path>.maximum-allocation-mb</queue-path></code> | The maximum memory value for each container. This configuration will override the <code>yarn.scheduler.maximum-allocation-mb</code> value, but it must be less than or equal to the system's <code>yarn.scheduler.maximum-allocation-mb</code> value. |
| Maximum vCore count for Container | <code>yarn.scheduler.capacity.<queue-path>.maximum-allocation-vcores</queue-path></code> | The maximum number of cores for each container. This configuration will override the <code>yarn.scheduler.maximum-allocation-vcores</code> value, but it must be less than or equal to the system's <code>yarn.scheduler.maximum-allocation-vcores</code> value. |
| Resource pool status | <code>yarn.scheduler.capacity.<queue-path>.state</queue-path></code> | The status of the queue. It can be either running or stopped. If the queue is in a stopped state, new applications cannot be submitted to it or any of its sub-queues. |
| Maximum Application Count – Max-Applications | <code>yarn.scheduler.capacity.<queue-path>.maximum-applications</queue-path></code> | The maximum number of applications that can be concurrently active (running and suspended) within the system. |
| Maximum AM Proportion | <code>yarn.scheduler.capacity.<queue-path>.maximum-am-resource-percent</queue-path></code> | The maximum percentage of resources in the cluster available for running application hosts – controlling the number of concurrently active applications. |
| Resource pool priority | <code>yarn.scheduler.capacity.root.<leaf-queue-path>.default-application-priority</leaf-queue-path></code> | Configuring the priority of the resource queue, default is 0, the larger the set value, the higher the priority. |
| Submission | <code>yarn.scheduler.capacity.root.<queue-path>.acl_submit_applications</code> | The list of users who can submit apps to the queue. |
| Management | <code>yarn.scheduler.capacity.root.<queue-path>.acl_administer_queue</code> | The list of users who can manage the queue. |
| Delayed Scheduling | <code>yarn.scheduler.capacity.node-locality-delay</code> | Ensures task localization execution, with the number of times scheduling can be delayed. If the value is -1, delayed scheduling will be disabled. |

Configuring Resource Pool Mapping

1. Click on **Resource Pool Mapping** in the Policy Settings to access the Resource Pool Mapping page. Click on **Create New Resource Pool Mappings** to create a new resource pool mapping.



Creating a New Resource Pool Mapping:

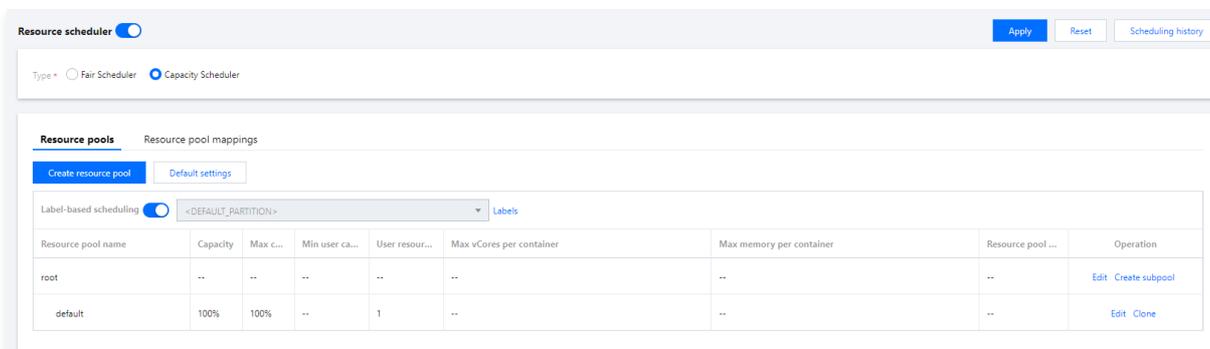


2. Overriding User-Specified Queues

By default, this is turned off. If a user has defined a mapped queue in the resource pool mapping, and the user specifies a queue when submitting a task, but the queue is different from the mapped queue: If the user-specified queue is 'default' or if overriding is enabled, the mapped queue will be used. Otherwise, the user-specified queue will be used.

Example of Tag Scheduling

1. Log in to the [EMR Console](#), select the corresponding Hadoop cluster from the cluster list, and click on **Details** to access the cluster details page.
2. In the Cluster Details page, select **Cluster Services > Yarn Component Card** at the top right corner, then choose **Operations > Resource Scheduling** to access the Resource Scheduling page.
3. Click on **Resource Scheduler Switch** and select Capacity Scheduler as the scheduler type.
4. Click on the **Tag Scheduling Switch**. After enabling tag scheduling, click on **Tag Management** to enter the Tag Management page.



5. Click on **Create New Tag**, enter the tag name, and set the tag type and the node bound to this tag as needed.

Create tag ✕

Label name

Label type Exclusive Non-exclusive

Node to be bound

6. Upon completion of the tag settings, click on **Activate Directive** to edit and view the resource queue of this tag in the resource pool.

The screenshot shows the 'Resource scheduling' page. At the top, there are navigation tabs: Service status, Job query, Roles, Client management, Configurations, and Resource scheduling. Below this is a blue information box with three numbered instructions. The 'Resource scheduler' section has a toggle switch for 'Resource scheduler' and buttons for 'Apply', 'Reset', and 'Scheduling history'. Underneath, there are radio buttons for 'Fair Scheduler' and 'Capacity Scheduler'. The 'Resource pools' section has a 'Create resource pool' button and a 'Default settings' button. A 'Label-based scheduling' toggle is also present. A table lists resource pools with columns for name, capacity, max capacity, min user capacity, user resources, max vCores per container, max memory per container, and operations.

| 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 subpool |
| default | 100% | 100% | .. | 1 | .. | .. | .. | Edit Clone |

7. In the Resource Scheduling page, click on **Create New Resource Pool** and select tags, capacity, maximum capacity, etc., according to business requirements.

Note
The capacity and maximum capacity of the resource pool in different tags are independent of each other, meaning they can be configured separately according to business needs without any mutual interference.

This screenshot is identical to the previous one, showing the 'Resource scheduling' page with the 'Resource scheduler' and 'Resource pools' sections. The table of resource pools is the same as in the previous screenshot.

8. After setting up the resource pool, click on **Deploy to Take Effect** to submit the deployment activation task to the backend.

Note
Since restarting ResourceManager is a high-risk operation, if you are prompted to restart ResourceManager when clicking on **Deploy to Take Effect**, please check in **Scheduling History** to see if the operation was successful, and verify in **Role**

Management whether the health status of ResourceManager is satisfactory.

Tag Management

Last updated: 2023-12-21 16:20:00

Feature Overview

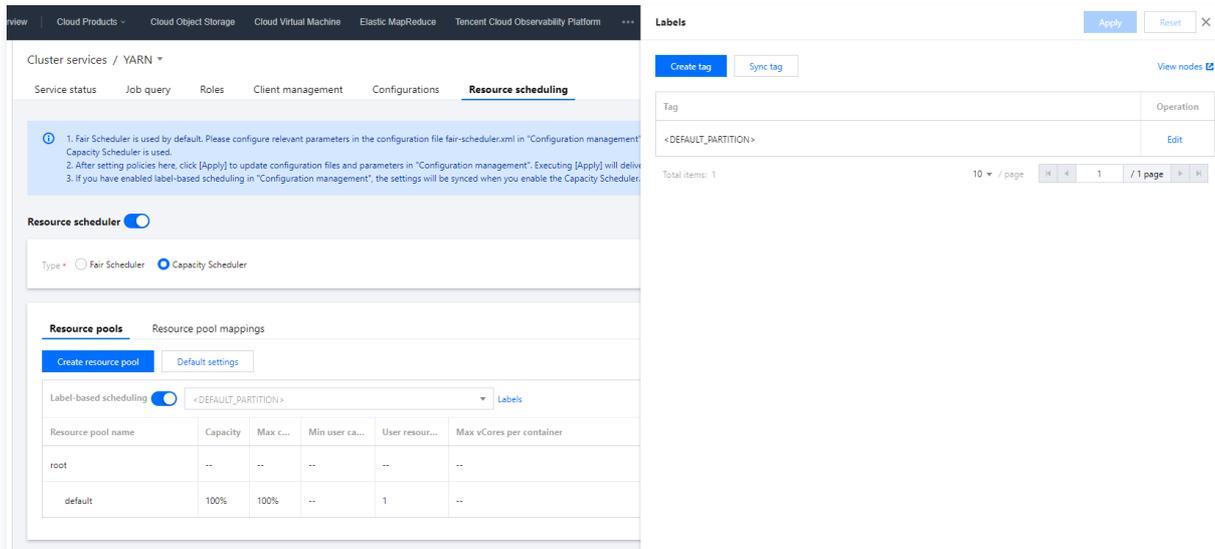
Tag Management offers the creation, editing, and deletion of tags, as well as node binding capabilities. Assigning distinct tags to various nodes within a cluster facilitates a more granular resource allocation atop the Capacity Scheduler, and supports the specification of execution locations for applications.

Preparations

Activate Yarn resource scheduling, switch the scheduler type to Capacity Scheduler, and enable tag scheduling.

Instructions

1. Log in to the [EMR Console](#), select the corresponding Hadoop cluster from the cluster list, and click on **Details** to access the cluster details page.
2. In the Cluster Details page, select **Cluster Services > Yarn Component Card** at the top right corner, then choose **Operations > Resource Scheduling** to access the Resource Scheduling page.
3. Click on **Resource Scheduler Switch** and select Capacity Scheduler as the scheduler type.
4. Click on the **Tag Scheduling Switch**. After enabling tag scheduling, click on **Tag Management** to enter the Tag Management page.



5. On the Tag Management page, you can view all tags for the current cluster. Additionally, you can add, edit, delete, and synchronize tags, as well as navigate to the WebUI to view nodes bound to tags.



Viewing Scheduling History

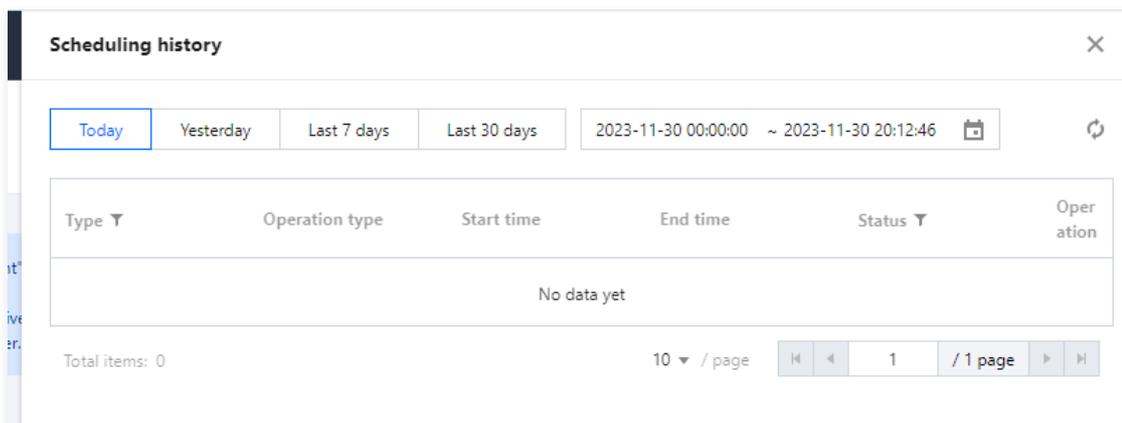
Last updated: 2023-12-21 16:20:11

Feature Overview

The Scheduling History provides visibility into the operational records of resource queue configurations, task statuses, and other pertinent information.

Instructions

1. Log in to the [EMR Console](#), select the corresponding Hadoop cluster from the cluster list, and click on **Details** to access the cluster details page.
2. In the Cluster Details page, select **Cluster Services > Yarn Component Card** at the top right corner, then choose **Operations > Resource Scheduling** to access the Resource Scheduling page.
3. On the Resource Scheduling page, click on **Scheduling History** to view details such as task start time, end time, status, and operational information.



4. You can filter scheduling records by execution time period. Click on **Details** under Operation Type to view more detailed information.

HBase RIT Fixing

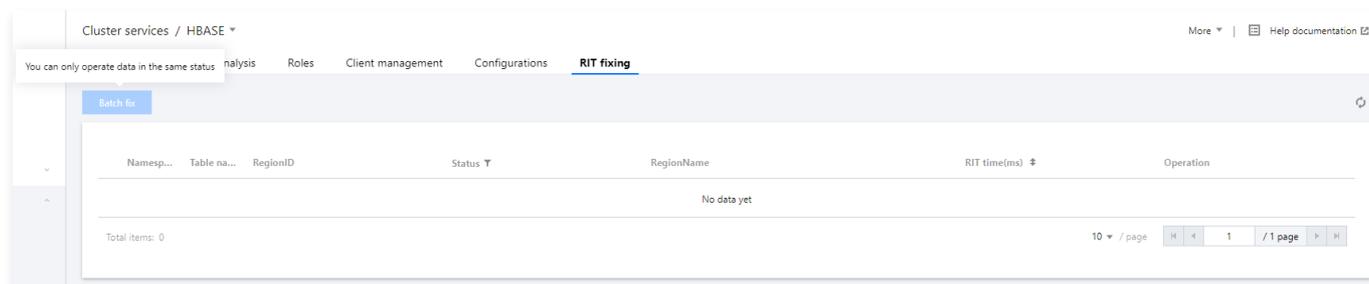
Last updated: 2024-01-12 11:18:00

Feature Overview

The HBase RIT Repair is applicable when encountering RIT (Region-In-Transition) issues within an HBase cluster (HBase version 2.2.0 and above), and when the Region is in a prolonged RIT state. It facilitates the repair of Regions in FAILED_OPEN, FAILED_CLOSED, and CLOSING states.

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding **Cluster ID/Name** to enter the cluster details page.
2. Within the cluster details page, click on **Cluster Services**, then select **Operations > RIT fixing** in the upper right corner of the HBase component. This allows you to view and repair Regions in RIT state.



3. Click on **Repair** in the list operations, or select the Regions that need repair and click on **Batch Repair**. In the RIT Repair pop-up page, confirm the relevant information.
4. After ensuring the information is accurate, click on **Confirm** to proceed with the RIT repair.
5. The progress and results of the RIT repair can be viewed in the Task Center.

Component port

Last updated: 2023-12-21 16:20:40

This document presents information on commonly used component ports.

Frequently Used HDFS Ports

| Component | Port | Note |
|-----------|------|---|
| HDFS | 4007 | NameNode RPC Port. This port is utilized for: <ol style="list-style-type: none"> 1. Communication between the HDFS client and the NameNode. 2. Connection between the DataNode and the NameNode. |
| HDFS | 4008 | HDFS HTTP Port (NameNode). This port is employed for: <ol style="list-style-type: none"> 1. Peer-to-peer NameNode checkpoint operations. 2. Remote web client connections to the NameNode UI. |
| HDFS | 4009 | HDFS HTTPS Port (NameNode). This port is utilized for: <ol style="list-style-type: none"> 1. Peer-to-peer NameNode checkpoint operations. 2. Remote web client connections to the NameNode UI. |
| HDFS | 4004 | DataNode IPC Server Port. This port is designated for: Client connections to the DataNode for the execution of RPC operations. |
| HDFS | 4001 | DataNode Data Transfer Port. This port is designated for: <ol style="list-style-type: none"> 1. HDFS client transmitting data from, or to, the DataNode. 2. Peer-to-peer data transmission between DataNodes. |
| HDFS | 4002 | DataNode HTTP Port. This port is employed for: In secure mode, remote web client connections to the DataNode UI. |
| HDFS | 4003 | DataNode HTTPS Port. This port is utilized for: In secure mode, remote web client connections to the DataNode UI. |
| HDFS | 4005 | JournalNode RPC Port. This port is utilized for: Client communication for accessing a variety of information. |
| HDFS | 4006 | JournalNode HTTP Port. This port is employed for: In secure mode, remote web client connections to the JournalNode. |
| HDFS | 4032 | HTTPS HTTP Server Listening Port. This port is utilized for: Remote REST interface connections to HTTPS. |

Commonly Used YARN Ports

| Component | Port | Note |
|-----------|------|--|
| YARN | 5024 | Job History Server Web HTTP Port. This port is employed for: Viewing the Web pages of the Job History Server. |
| YARN | 5022 | Job History Server Port. This port is employed for: <ol style="list-style-type: none"> 1. Utilized for the recovery of task data by MapReduce clients. 2. Employed for Job clients to retrieve task reports. |
| YARN | 5004 | Web HTTP Port for the ResourceManager Service. |
| YARN | 5002 | Web HTTPS Port for the ResourceManager Service. This port is utilized for: Accessing the ResourceManager Web application in secure mode. |
| YARN | 5004 | NodeManager Web HTTP Port. |
| YARN | 5005 | NodeManager Web HTTPS Port. This port is employed for: Accessing the NodeManager web application in secure mode. |
| YARN | 5001 | Scheduler Port for ResourceManager. |

| | | |
|------|------|---|
| YARN | 5000 | ResourceManager rpc interface, employed for client task submission. |
|------|------|---|

Commonly Used Ports for Hive.

| Component | Port | Note |
|-----------|------|--|
| Hive | 7010 | Port for WebHCat providing REST services. This port is utilized for: Communication between the WebHCat client and the WebHCat server. |
| Hive | 7001 | Port for HiveServer providing Thrift services. This port is utilized for: Communication between the HiveServer client (beeline) and HiveServer. |
| Hive | 7004 | Port for MetaStore providing Thrift services. This port is utilized for: Communication between the MetaStore client and MetaStore, specifically, communication between HiveServer and MetaStore. |
| Hive | 7003 | Web UI port for Hive. This port is used for HTTPS/HTTP communication with the Hive UI server via web requests. |

Commonly Used Ports for Spark.

| Component | Port | Note |
|-----------|-------|-----------|
| Spark | 10000 | HTTP Port |

Commonly Used Ports for Presto.

| Component | Port | Note |
|-----------|------|--|
| Presto | 9000 | HTTP port for Presto Coordinator or Worker to provide external services. |

Commonly Used Ports for PrestoSQL.

| Component | Port | Note |
|-----------|------|---|
| PrestoSQL | 9000 | HTTP port for PrestoSQL Coordinator or Worker to provide external services. |

Commonly Used Ports for Trino.

| Component | Port | Note |
|-----------|------|---|
| Trino | 9000 | HTTP port for TRINO Coordinator or Worker to provide external services. |

Commonly Used Ports for Impala.

| Component | Port | Note |
|-----------|-----------------|---|
| Impala | 27000、 27009 | Port provided for communication with Impala applications. The former is used by Impala 2.x, while the latter is utilized by Impala 3.x. |
| Impala | 27001 | Port provided for communication with Impala-shell. |
| Impala | 27004 | Web port for ImpalaServer. |
| Impala | 27007 | Web port for ImpalaCatalog. |
| Impala | 27005 | Web port for ImpalaStateStore. |

Commonly used ports for Kudu.

| Component | Port | Note |
|-----------|------|----------------------|
| Kudu | 7051 | KuduMaster PRC port. |
| Kudu | 7050 | KuduServer RPC port. |

| | | |
|------|------|-----------------------|
| Kudu | 8051 | KuduMaster HTTP port. |
| Kudu | 8050 | KuduServer HTTP port. |

Commonly used ports for ClickHouse.

| Component | Port | Note |
|------------|------|------------------------------------|
| ClickHouse | 9000 | Business client TCP access port. |
| ClickHouse | 8123 | Business client HTTP access port. |
| ClickHouse | 9009 | Business client HTTPS access port. |

Commonly used ports for Kylin.

| Component | Port | Note |
|-----------|-------|------------------|
| Kylin | 16500 | Kylin HTTP port. |

Commonly used ports for Doris.

| Component | Port | Note |
|-----------|------|--|
| Doris | 8000 | Thrift server port on Broker. |
| Doris | 9060 | Thrift server port on BE, used to receive requests from FE. |
| Doris | 8060 | Brpc port on BE, utilized for communication between BEs. |
| Doris | 9050 | Heartbeat service port (thrift) on BE, designed to receive heartbeats from FE. |
| Doris | 8040 | Thrift server on Broker, designed to receive requests. |
| Doris | 9010 | Port on FE used for communication between metadata management modules (bdbje). |
| Doris | 8030 | HTTP server port on FE. |
| Doris | 9020 | Thrift server port on FE. |
| Doris | 9030 | Port on FE for receiving access from Mysql client. |

Commonly used ports in StarRocks.

| Component | Port | Note |
|-----------|------|--|
| StarRocks | 8000 | Thrift server port on Broker. |
| StarRocks | 9060 | Thrift server port on BE, used to receive requests from FE. |
| StarRocks | 8060 | Brpc port on BE, utilized for communication between BEs. |
| StarRocks | 9050 | Heartbeat service port (thrift) on BE, designed to receive heartbeats from FE. |
| StarRocks | 8040 | Thrift server on Broker, designed to receive requests. |
| StarRocks | 9010 | Port on FE used for communication between metadata management modules (bdbje). |
| StarRocks | 8030 | HTTP server port on FE. |
| StarRocks | 9020 | Thrift server port on FE. |
| StarRocks | 9030 | Port on FE for receiving access from Mysql client. |

Commonly used ports in Druid.

| Component | Port | Note |
|-----------|------|------|
|-----------|------|------|

| | | |
|-------|-------|--|
| Druid | 8082 | Broker service listening port (broker-runtime.properties), utilized for receiving client queries. |
| Druid | 8081 | Coordinator service listening port (coordinator-runtime.properties), employed for communication with other components. |
| Druid | 8083 | Historical service listening port (historical-runtime.properties), employed for communication with other components. |
| Druid | 8091 | MiddleManager service listening port (middleManager-runtime.properties), employed for communication with other components. |
| Druid | 8090 | Overlord service listening port (overlord-runtime.properties), employed for communication with other components. |
| Druid | 18888 | Router service listening port (router-runtime.properties), utilized for routing client queries to the broker. |

Commonly used HBase ports

| Component | Port | Note |
|-----------|------|--|
| HBase | 6000 | HMaster RPC port. This port is utilized for HBase client connections to HMaster. |
| HBase | 6001 | HMaster HTTPS port. This port is utilized for remote web client connections to the HMaster UI. |
| HBase | 6002 | RS (RegionServer) RPC port. This port is employed for HBase client connections to the RegionServer. |
| HBase | 6003 | RegionServer HTTPS port. This port is employed for remote web client connections to the RegionServer UI. |
| HBase | 6005 | Listening port for the Info Server of the HBase Thrift Server. |
| HBase | 6004 | Listening port for the HBase Thrift Server. |

Commonly used ports for Flink.

| Component | Port | Note |
|-----------|-------|---|
| Flink | 16001 | Flink's Web UI port. This port is utilized for HTTP/HTTPS communication between the client web requests and the Flink server. |

Frequently utilized ports for Storm.

| Component | Port | Note |
|-----------|-------|---------------------------------|
| Storm | 15002 | Service port for the Logviewer. |
| Storm | 15000 | Service port for Nimbus. |
| Storm | 15001 | Storm Web UI . |

Frequently utilized ports for Hue.

| Component | Port | Note |
|-----------|-------|---|
| Hue | 13000 | Hue's HTTPS service port. This port is employed for providing web services via HTTPS, with support for modifications. |

Commonly used ports for Oozie.

| Component | Port | Note |
|-----------|------|------|
| | | |

| | | |
|-------|-------|--|
| Oozie | 12000 | HTTP port, designated for client access. |
|-------|-------|--|

Commonly utilized ports for Superset.

| Component | Port | Note |
|-----------|-------|---|
| Superset | 18088 | Superset service port, designated for client connections. |

Frequently used ports for Zeppelin.

| Component | Port | Note |
|-----------|-------|---|
| Zeppelin | 18000 | Zeppelin service port, designated for client connections. |

Frequently utilized ports for Kafka.

| Component | Port | Note |
|-----------|------|--|
| Kafka | 9092 | The Broker provides services for data reception and acquisition. |

Commonly used ports for KafkaManager.

| Component | Port | Note |
|--------------|------|---|
| KafkaManager | 9000 | KafkaManager service port, designated for client connections. |

Commonly employed ports for Ranger.

| Component | Port | Note |
|-----------|------|---|
| Ranger | 6080 | Ranger Admin web UI. |
| Ranger | 5151 | Authentication service port for RangerUsersync. |

Commonly utilized ports for COSRanger.

| Component | Port | Note |
|-----------|------|--|
| COSRanger | 9999 | CosRangerServer service port, intended for client connections. |

Frequently used ports for KRB5.

| Component | Port | Note |
|-----------|------|---|
| KRB5 | 749 | Service port for Kadmin. |
| KRB5 | 754 | Service port for Kprop. |
| KRB5 | 88 | Service port for Kerberos. This port is utilized for component authentication to the Kerberos service. It may be employed in the configuration of cluster mutual trust. |

Commonly used ports for KNOX.

| Component | Port | Note |
|-----------|-------|--|
| KNOX | 30002 | HTTP port, typically employed for browser connections. |
| KNOX | 33389 | The port number for the built-in LDAP of Knox, used for Knox authentication. |

Commonly used ports for ZooKeeper.

| Component | Port | Note |
|-----------|------|---|
| ZooKeeper | 2181 | Client port for ZooKeeper. This port is utilized for the connection of ZooKeeper clients to the ZooKeeper server. |

Commonly used ports for OpenLDAP.

| Component | Port | Note |
|-----------|------|---------------------|
| OpenLDAP | 389 | Client access port. |

Commonly used ports for Tez.

| Component | Port | Note |
|-----------|------|----------------------|
| Tez | 2000 | Web UI port for Tez. |

Commonly used ports for Livy.

| Component | Port | Note |
|-----------|------|--|
| Livy | 8998 | The LivyServer service-side listening port, utilized for client connections. |

Commonly used ports for Kyuubi.

| Component | Port | Note |
|-----------|-------|--|
| Kyuubi | 10009 | The KyuubiServer service-side listening port, utilized for client connections. |

Commonly used ports for Alluxio.

| Component | Port | Note |
|-----------|-------|--|
| Alluxio | 2001 | The AlluxioJobMaster RPC port, through which the AlluxioMaster assigns tasks to the AlluxioJobMaster. |
| Alluxio | 3001 | The AlluxioJobWorker RPC port, through which the AlluxioJobMaster assigns tasks to the AlluxioJobWorker. |
| Alluxio | 19998 | The AlluxioMaster RPC port, through which clients establish connections with the AlluxioMaster. |
| Alluxio | 29998 | The AlluxioWorker RPC port, through which the AlluxioMaster assigns read and write tasks to the AlluxioWorker. |

Commonly used ports for GooseFS.

| Component | Port | Note |
|-----------|------|--|
| GooseFS | 9206 | The AlluxioJobMaster RPC port, through which the AlluxioMaster assigns tasks to the AlluxioJobMaster. |
| GooseFS | 9210 | The AlluxioJobWorker RPC port, through which the AlluxioJobMaster assigns tasks to the AlluxioJobWorker. |
| GooseFS | 9201 | The AlluxioMaster RPC port, through which clients establish connections with the AlluxioMaster. |
| GooseFS | 9211 | The GooseFSProxy service listening port, utilized for proxying clients. |

| | | |
|---------|------|--|
| GooseFS | 9204 | The AlluxioWorker RPC port, through which the AlluxioMaster assigns read and write tasks to the AlluxioWorker. |
|---------|------|--|

Commonly used ports for Ganglia.

| Component | Port | Note |
|-----------|------|---|
| Ganglia | 1602 | The Gmetad service port. |
| Ganglia | 1603 | The Gmetad service port, employed for receiving data queries from HTTPd. |
| Ganglia | 1601 | The communication port between gmond processes, also available for gmetad access. |

Monitoring and Alarming

Node Status

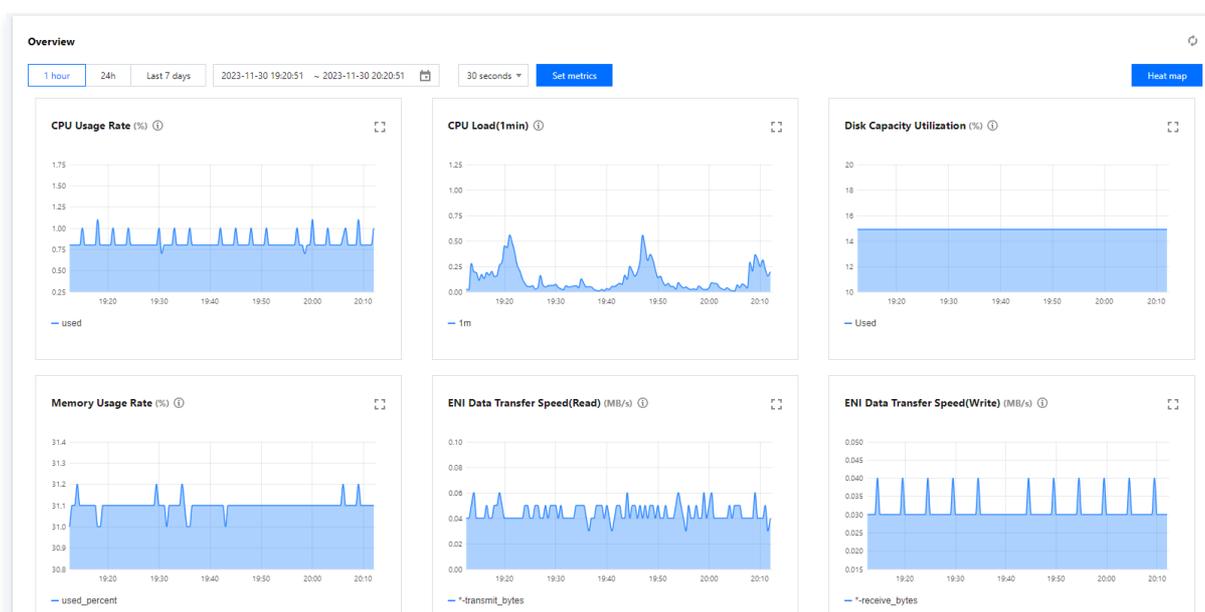
Last updated: 2023-12-21 16:21:23

Feature Overview

The Node Status page provides a comprehensive display of the monitoring overview and list of all nodes in the current cluster, along with the capability to view all node heatmaps. Users can manage the status and metric information of nodes through the Elastic MapReduce console in their daily usage.

Instructions

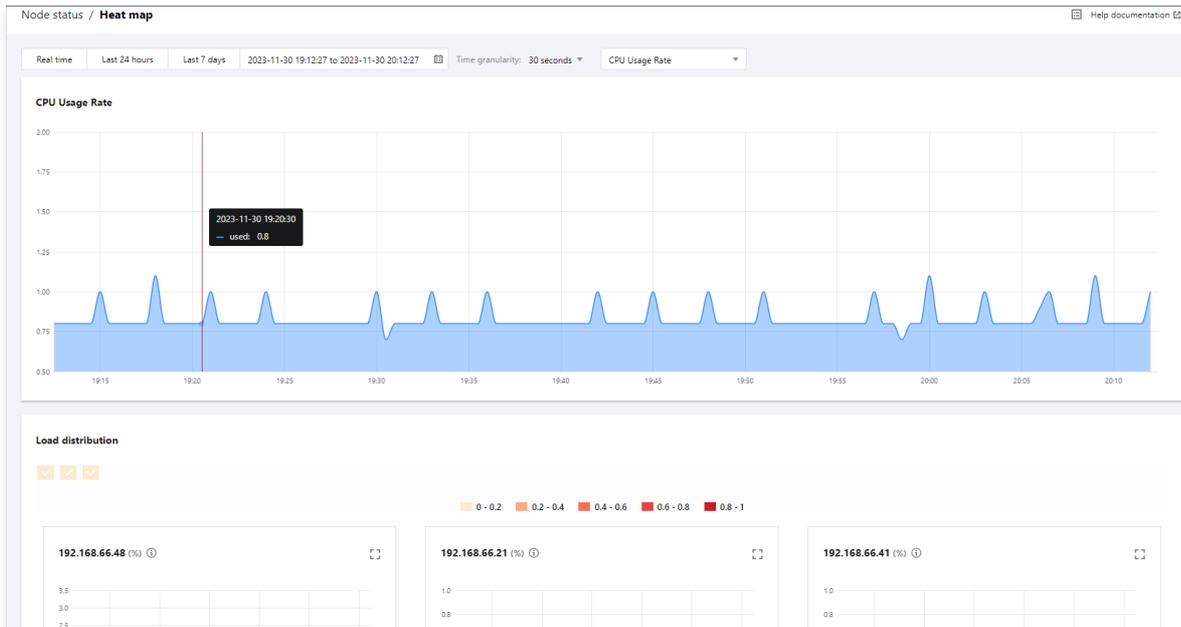
1. Log in to the [Elastic MapReduce Console](#). In the cluster list, click on the corresponding cluster **ID/Name** to access the cluster details page.
2. Select **Cluster Resources > Node Status** to view the monitoring information of all nodes within the cluster.
3. Within node monitoring, you can view an aggregated overview of monitoring metrics and a list of all nodes for the current cluster.



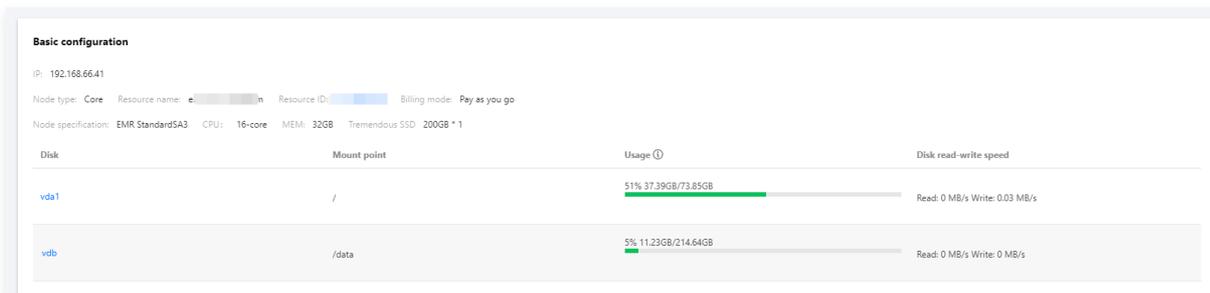
- **Overview:** Provides an intuitive view of the aggregated monitoring metrics and statistical rules for all nodes over a specific time period. Click on **Set Metrics** to customize the display metrics.



- **Heatmap:** The load heatmap provides a more intuitive display of the node's load status, with the ability to specify a time period and load conditions for viewing. The load heatmap is primarily divided into two sections: one section presents an aggregated load diagram of all nodes in the current cluster, while the other section displays individual heatmaps for all nodes, allowing for an intuitive view of the load status of all nodes.



- **Node List:** Displays all current nodes along with their deployment types, CPU utilization, memory utilization, and disk utilization. Click on the corresponding node IP to view the basic configuration, deployment status, load status, and node monitoring of an individual node.
- **Basic Configuration:** Allows you to view the basic information of the current node, such as node type, resource name, resource ID, billing type, specification size, and so forth.



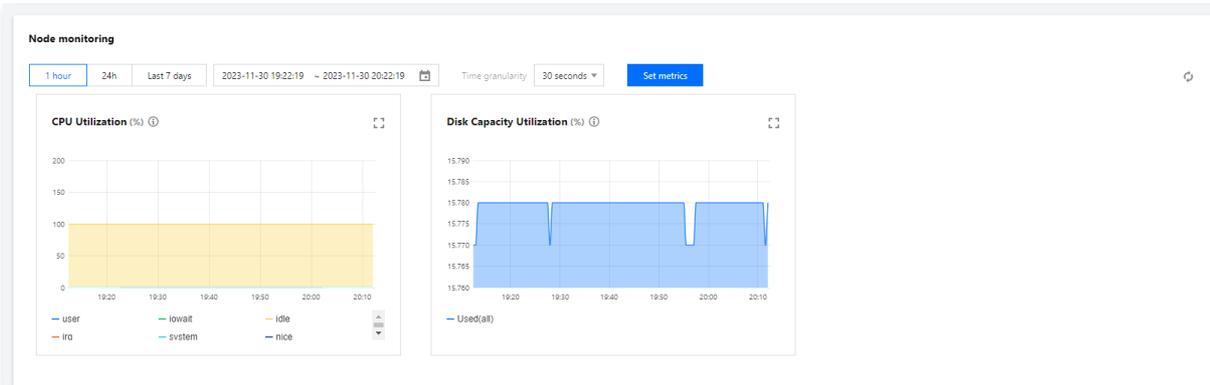
- **Deployment Status:** Allows you to view the service deployment status of the current node, such as whether it is a standard process and whether the process status is running normally.

| Process name | Process type | Process status |
|---------------|------------------|----------------|
| AlluxioWorker | Standard process | Running |
| DataNode | Standard process | Running |
| NodeManager | Standard process | Running |
| RegionServer | Standard process | Running |
| Trino-Worker | Standard process | Running |

- **Load Status:** Allows you to view the TOP N process status in the current node dimension, with the ability to view based on a specified time.



- **Node Monitoring:** Allows you to view the load trend charts for each group indicator of the current node. By default, six indicators are displayed, with a maximum of twelve indicators that can be displayed. You can click on **Set Indicators** to customize the display indicators.



Service Status

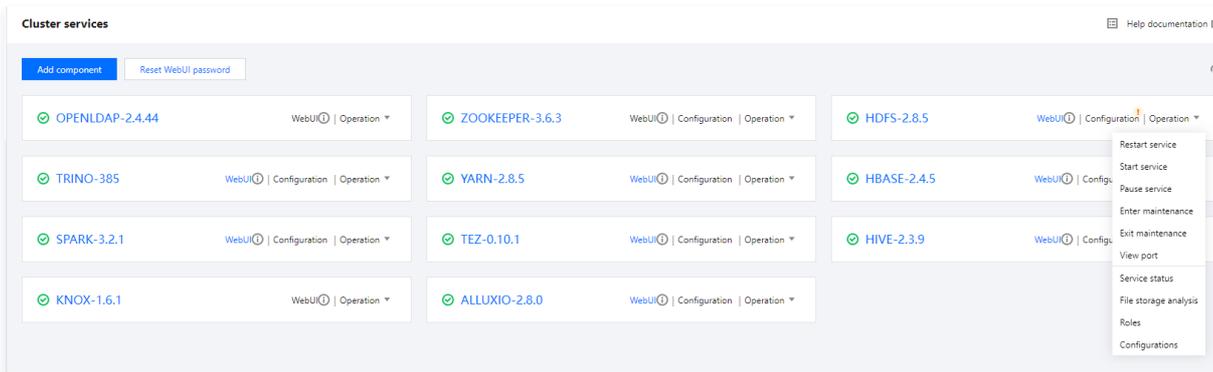
Last updated: 2023-12-21 16:21:44

Feature Overview

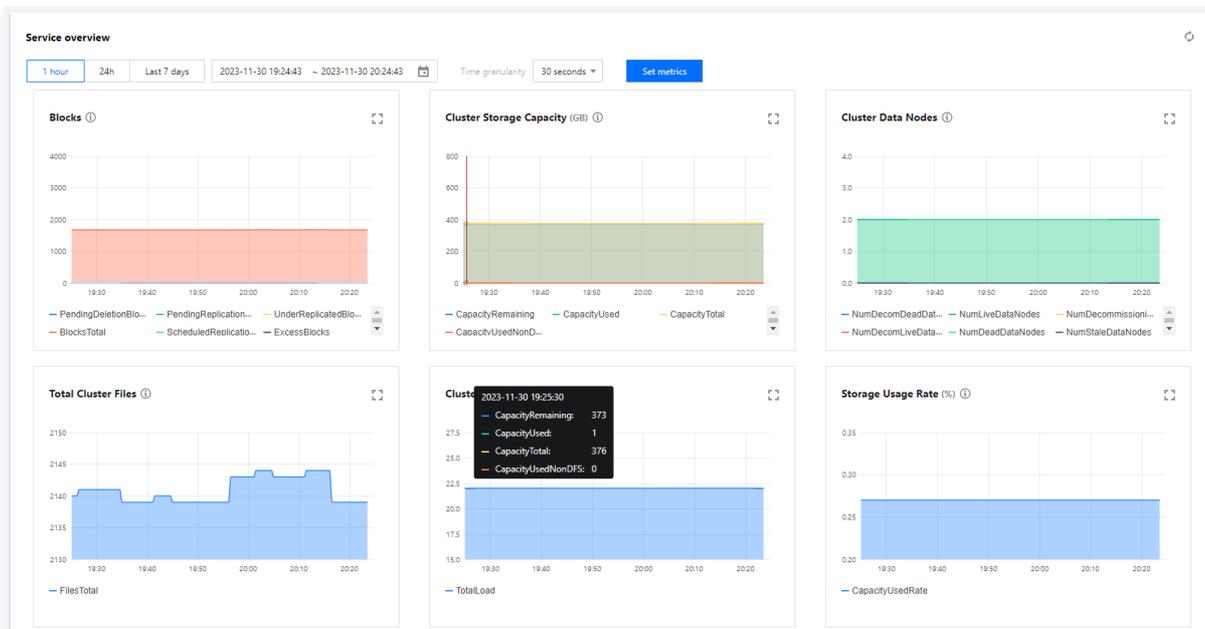
The Service Status provides comprehensive monitoring capabilities for the primary services installed on the cluster, including HDFS, YARN, HIVE, ZOOKEEPER, SPARK, HBase, and PRESTO, among others. This document presents the procedure for viewing the cluster service status via the console.

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding cluster ID/Name to access the cluster details page.
2. In the cluster details page, select **Cluster Services**, and click on **Operation > Service Status** in the top right corner of the corresponding component, using HDFS as an example.



3. The Service Status page offers three monitoring views based on service dimensions, namely **Service Summary**, **Health Status**, and **Service Overview**. As each service component varies, the displayed dimensions will differ accordingly.
4. The **Service Overview** allows for an intuitive examination of various metrics and their statistical rules for service components within a specific time frame. The system defaults to displaying six metric items, but you can customize the display metrics by clicking on **Set Metrics**.



Setting Metrics:

Set metrics ✕

Cluster Overview

- Snapshots [Preview](#)
- Block Capacity [Preview](#)
- Blocks [Preview](#)
- Cluster Storage Capacity [Preview](#)
- Cluster Data Nodes [Collapse](#)

- Total Cluster Files [Preview](#)
- Cluster Load [Preview](#)
- Master/slave Conditions [Preview](#)
- Storage Usage Rate [Preview](#)

Disk

- Disk Failures [Preview](#)

Confirm
Cancel

5. The Service Summary displays the current overall usage status of the service.
6. The Health Status displays the operational overview of each service component. Clicking on the role name or operational overview will redirect you to the Role Management or Role Status page.

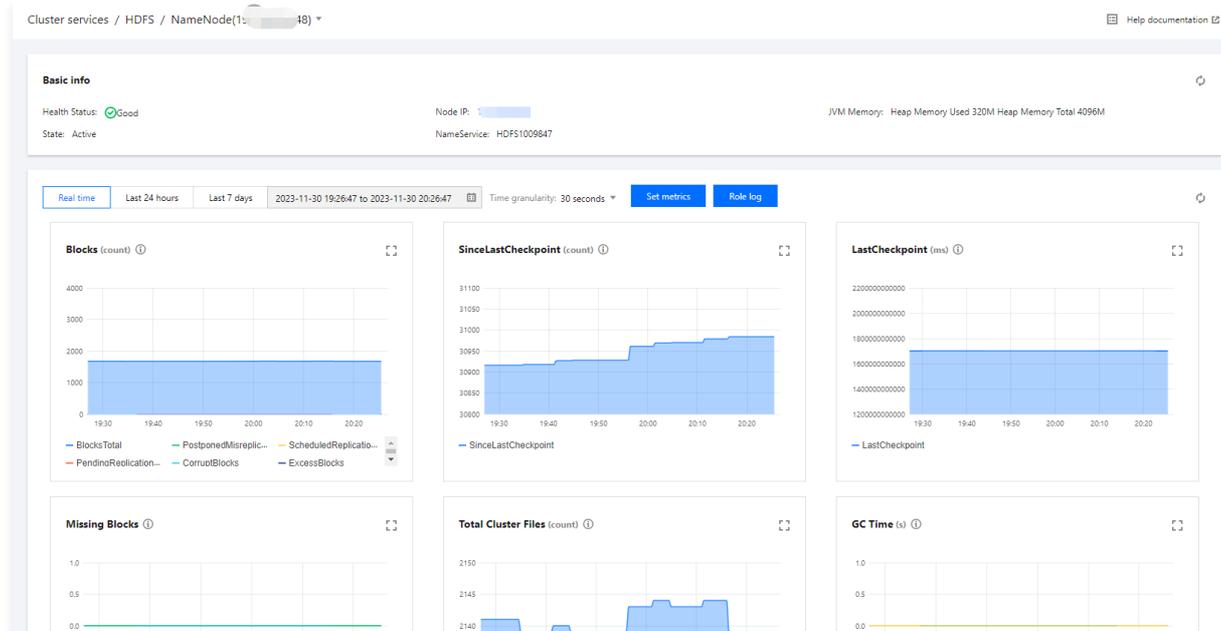
Service summary 🔄

NameNode HA State: 192.168.66.48(Active) -(Standby) DataNode Status: 2 Alive 0 Dead 0 Decommissioning Missing Report: 0 Blocks 0 BlocksWithReplicationFactorOne

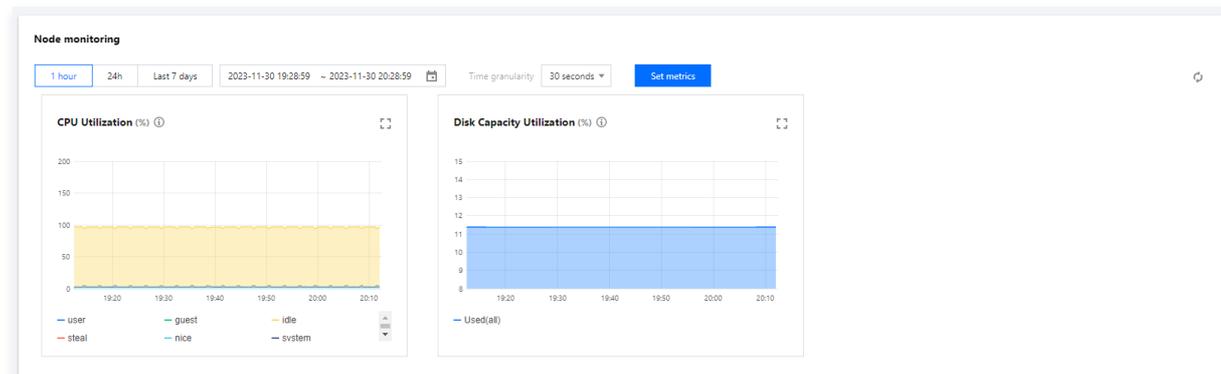
Health status Create instant inspection 🔄

NameNode(1) ✔️ 1 runningGood DataNode(2) ✔️ 2 runningGood

Upon entering a specific role status page, you can customize the display metrics by clicking on **Set Metrics**.



7. The Service Overview allows for the examination of statistical metrics at the cluster dimension. You can customize the display metrics by clicking on **Set Metrics**.



Note

- By default, the Service Monitoring displays the HDFS service components. You have the option to manually adjust and view other service components.
- Given the varying nature of each service component, the service monitoring dimensions differ accordingly. For instance, HBASE supports table-level monitoring dimensions.

Cluster Event

Last updated: 2023-12-21 16:22:02

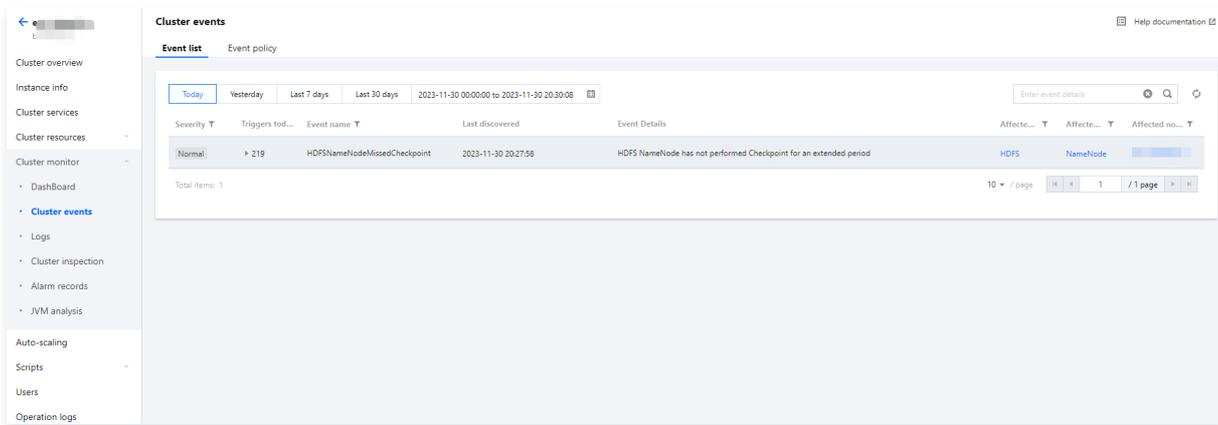
Feature Overview

The cluster events encompass both the event list and the event policy.

- **Event List:** A record of key changes or abnormal events occurring within the cluster.
- **Event Policy:** Supports the customization of event monitoring trigger strategies based on business conditions. Monitored events can be set as cluster inspection items.

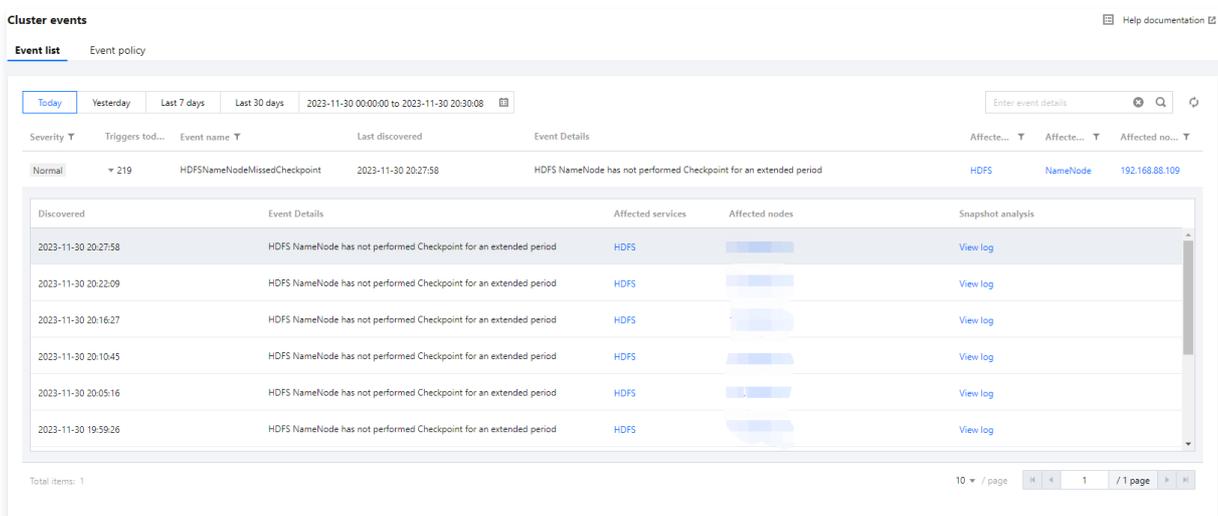
View event list

1. Log in to the [Elastic MapReduce Console](#). In the cluster list, click on the corresponding cluster **ID/Name** to access the cluster details page.
2. In the cluster details page, select **Cluster Monitoring > Cluster Events > Event List** to directly view all operation events of the current cluster.



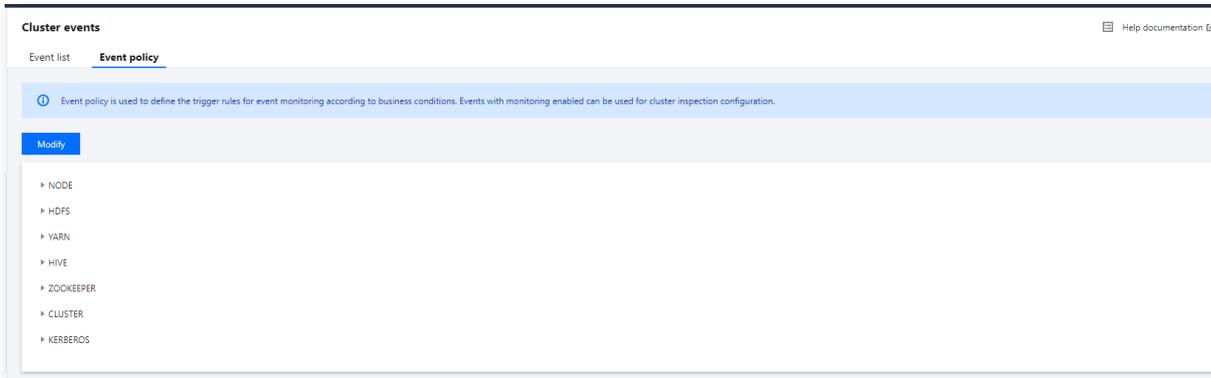
The severity levels are explained as follows:

- **Fatal:** Abnormal events of nodes or services that require manual intervention. Otherwise, the service will be unavailable. Such events may persist for a period of time.
 - **Severe:** Temporarily has not caused service or node unavailability issues. It is a warning category. If left untreated, it may lead to fatal events.
 - **General:** Records regular events occurring in the cluster, generally requiring no special handling.
3. Click on the **Number of Triggers Today** column value to view the event trigger records, as well as related metrics, logs, or scenes of the event records.

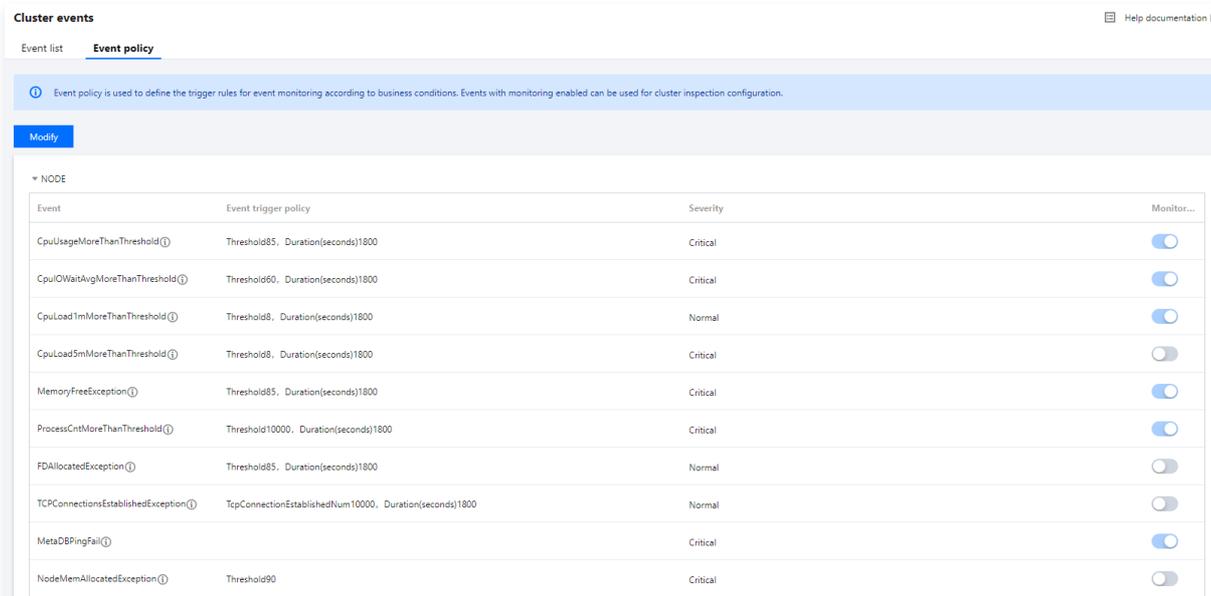


Setting Event Policies

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding cluster ID/Name to access the cluster details page.
2. In the cluster details page, select **Cluster Monitoring > Cluster Events > Event Policy** to customize and set event monitoring trigger policies.
3. The event configuration list includes: Event Name, Event Discovery Policy, Severity Level (Fatal/Serious/General), and Monitoring Activation. It supports modifications and saving.



4. Event discovery policies fall into two categories: one type of event is a system-fixed policy event, which does not support user modifications; the other type of event varies according to different customer business standards, and supports user settings.



5. Event policies can be customized to enable or disable event monitoring. Only events with activated monitoring can be selected in the inspection items of cluster inspection. Some events are enabled by default, while others are enabled by default and cannot be disabled. The specific rules are as follows:

| Category | Event name | Event Significance | Recommendations & Measures | Default value | Severity | Permitted to close | Enabled by Default |
|----------|--|--|----------------------------|-------------------|----------|--------------------|--------------------|
| Node | Continuous CPU Utilization Exceeds Threshold | Machine CPU Utilization $\geq m$, for a duration of t seconds ($300 \leq t \leq 2592000$) | Node expansion or upgrade | $m=85$, $t=1800$ | Critical | Supported | Supported |
| | Average CPU IO Wait | Average CPU iowait utilization $\geq m$ within t seconds ($300 \leq t \leq 2592000$) | Manually check | $m=60$, $t=1800$ | Critical | Supported | Supported |

| | | | | | | | |
|--|--|--|-------------------|----------|-----------|--------------|--|
| Exceeds Threshold | | | | | | | |
| Continuous CPU Load Exceeds Threshold for One Minute | CPU Load $\geq m$ for a duration of t seconds ($300 \leq t \leq 2592000$) | Node expansion or upgrade | $m=8, t=1800$ | Average | Supported | Not required | |
| Continuous CPU Load Exceeds Threshold for Five Minutes | CPU Load $\geq m$ for Five Minutes, Continuously for t seconds ($300 \leq t \leq 2592000$) | Node expansion or upgrade | $m=8, t=1800$ | Critical | Supported | Not required | |
| The memory usage consistently goes over the limit. | Memory Usage $\geq m$, Continuously for t seconds ($300 \leq t \leq 2592000$) | Node expansion or upgrade | $m=85, t=1800$ | Critical | Supported | Supported | |
| The total number of system processes consistently goes over the limit. | Total System Processes $\geq m$, Continuously for t seconds ($300 \leq t \leq 2592000$) | Manually check | $m=10000, t=1800$ | Critical | Supported | Supported | |
| Node file handle usage consistently goes over the limit. | Node File Handle Usage $\geq m$, Continuously for t seconds ($300 \leq t \leq 2592000$) | Manually check | $m=85, t=1800$ | Average | Supported | Not required | |
| Node TCP connection count consistently exceeds the threshold. | Node TCP Connection Count $\geq m$, Continuously for t seconds ($300 \leq t \leq 2592000$) | Inspect for potential connection leaks. | $m=10000, t=1800$ | Average | Supported | Not required | |
| Node memory usage exceeds the limit | The cumulative memory usage configuration of all roles on the node surpasses the physical memory threshold of the node. | Adjust the allocation of heap memory for the node process. | 90% | Critical | Supported | Not required | |
| Ping to the metadata database failed. | CDB heartbeat has not been reported on schedule. | - | - | - | - | - | |
| The utilization of single disk capacity keeps exceeding the threshold. | The utilization of single disk capacity is greater than or equal to 'm', persisting for 't' seconds ($300 \leq t \leq 2592000$). | Node expansion or upgrade | $m=0.85, t=1800$ | Critical | Supported | Supported | |

| | | | | | | | |
|------|---|---|--|------------------------------|----------|--------------|--------------|
| | The utilization rate of single disk IO devices consistently surpasses the threshold. | The utilization rate of single disk IO devices is greater than or equal to 'm', persisting for 't' seconds (300<=t<=2592000). | Node expansion or upgrade | m=0.85, t=1800 | Critical | Supported | Supported |
| | The usage rate of single disk INODES consistently exceeds the threshold. | The usage rate of single disk INODES is greater than or equal to 'm', persisting for 't' seconds (300<=t<=2592000). | Node expansion or upgrade | m=0.85, t=1800 | Critical | Supported | Supported |
| | The difference between the sub-machine's UTC time and NTP time exceeds the threshold. | The discrepancy between the sub-machine's UTC time and NTP time surpasses the threshold (measured in milliseconds). | <ol style="list-style-type: none"> 1. Ensure the NTP daemon is in operational status. 2. Ensure that network communication with the NTP server is functioning properly. | The difference equals 30000. | Critical | Supported | Supported |
| | Automatic compensation for faulty nodes. | Upon enabling the automatic compensation feature, should task nodes or router nodes malfunction, the system will automatically purchase replacements of the same model and configuration to compensate. | <ol style="list-style-type: none"> 1. Compensation replacement has been successfully executed, no further attention is required. 2. Compensation replacement has failed. Please proceed to the console to manually terminate and repurchase nodes for replacement. | - | Average | Supported | Supported |
| | Node failure | There are faulty nodes within the cluster. | Please proceed to the console for resolution or submit a ticket to liaise with a specialist for assistance. | - | Critical | Not required | Supported |
| | Node disk IO is exhibiting abnormalities. | Node disk IO is exhibiting abnormalities (Detection is based on device IOPS and IO settings usage rate, applicable to certain IO abnormal situations). | The node's disk IO is exhibiting abnormalities. Resolution method: The issue may be due to IO Hang or disk abnormalities. | - | Critical | Supported | Not required |
| HDFS | The total number of HDFS files consistently surpasses | The total number of files in the cluster is greater than or equal to m, for a duration of t seconds (where 300 <= t <= 2592000). | Increasing the memory allocation for the namenode. | m=50,000,000, t=1800 | Critical | Supported | Not required |

| | | | | | | | |
|---|--|--|----------------------|----------|-----------|--------------|--|
| the threshold. | | | | | | | |
| The total number of HDFS blocks consistently exceeds the threshold. | The total number of blocks in the cluster is greater than or equal to m, for a duration of t seconds (where $300 \leq t \leq 2592000$). | Increase the memory allocation for the namenode or enlarge the block size. | m=50,000,000, t=1800 | Critical | Supported | Not required | |
| The quantity of HDFS data nodes marked as 'Dead' persistently surpasses the threshold. | The quantity of data nodes marked as 'Dead' is greater than or equal to m, for a duration of t seconds (where $300 \leq t \leq 2592000$). | Manually check | m=1, t=1800 | Average | Supported | Not required | |
| The utilization rate of HDFS storage space persistently surpasses the threshold. | The utilization rate of HDFS storage space is greater than or equal to m, for a duration of t seconds (where $300 \leq t \leq 2592000$). | Clean up files within HDFS or expand the cluster capacity. | m=85, t=1800 | Critical | Supported | Supported | |
| A primary-secondary switch has occurred in the NameNode. | A primary-secondary switch has occurred in the NameNode. | Investigate the cause of the NameNode switch. | - | Critical | Supported | Supported | |
| The delay in processing NameNode RPC requests consistently exceeds the threshold. | The delay in processing RPC requests is greater than or equal to m milliseconds, for a duration of t seconds (where $300 \leq t \leq 2592000$). | Manually check | m=300, t=300 | Critical | Supported | Not required | |
| The current number of connections to the NameNode consistently surpasses the threshold. | The current number of NameNode connections is $\geq m$, with a duration of t seconds ($300 \leq t \leq 2592000$). | Manually check | m=2000, t=1800 | Average | Supported | Not required | |
| A full GC event has occurred in the NameNode. | A full GC event has occurred in the NameNode. | Parameter optimization. | - | Critical | Supported | Supported | |

| | | | | | | |
|---|--|--|------------------------|----------|-----------|--------------|
| The JVM memory usage of the NameNode consistently exceeds the threshold. | The JVM memory usage of the NameNode consistently remains $\geq m$, for a duration of t seconds ($300 \leq t \leq 2592000$). | Adjusting the heap memory size of the NameNode. | $m=85$, $t=1800$ | Critical | Supported | Supported |
| The delay in processing DataNode RPC requests consistently exceeds the threshold. | The delay in processing RPC requests is greater than or equal to m milliseconds, for a duration of t seconds (where $300 \leq t \leq 2592000$). | Manually check | $m=300$, $t=300$ | Average | Supported | Not required |
| The current number of DataNode connections consistently exceeds the threshold. | The current number of DataNode connections remains $\geq m$, for a duration of t seconds ($300 \leq t \leq 2592000$). | Manually check | $m=2000$, $t=1800$ | Average | Supported | Not required |
| DataNode experiences full GC | A full GC event has occurred in the NameNode. | Parameter optimization. | - | Average | Supported | Not required |
| The JVM memory usage of the DataNode consistently exceeds the threshold. | The JVM memory usage of the NameNode consistently remains $\geq m$, for a duration of t seconds ($300 \leq t \leq 2592000$). | Adjusting the heap memory size of the DataNode. | $m=85$, $t=1800$ | Average | Supported | Supported |
| Both NameNode services in HDFS are in Standby status. | Both NameNode roles are concurrently in Standby status. | Manually check | - | Critical | Supported | Supported |
| The number of HDFS MissingBlocks goes over the limit consistently. | The number of MissingBlocks in the cluster is greater than or equal to ' m ', persisting for ' t ' seconds ($300 \leq t \leq 604800$). | It is recommended to investigate the occurrence of data block corruption in HDFS, using the command 'hadoop fsck /' to check the distribution of HDFS files. | $m=1$, $t=1800$ | Critical | Supported | Supported |
| The HDFS NameNode enters Safe Mode. | The NameNode enters Safe Mode, persisting for 300 seconds. | It is recommended to investigate the occurrence of data block corruption in HDFS, using the | - | Critical | Supported | Supported |

| | | | | | | | |
|------|--|--|---|--------------|----------|-----------|--------------|
| | | | command 'hadoop fsck /' to check the distribution of HDFS files. | | | | |
| | The HDFS NameNode has not performed a Checkpoint for an extended period of time. | The HDFS NameNode has not performed a Checkpoint for an extended period of time. | <ol style="list-style-type: none"> Inspect the status of the SecondaryNameNode (Standby NameNode). Inspect the parameters 'dfs.namenode.checkpoint.period' and 'dfs.namenode.checkpoint.txns' in the HDFS configuration file 'hdfs-site.xml'. Examine the log information of the HDFS cluster. | m=24 | Average | Supported | Supported |
| YARN | The current number of lost NodeManagers in the cluster consistently exceeds the threshold. | The current number of lost NodeManagers in the cluster is $\geq m$, persisting for t seconds ($300 \leq t \leq 2592000$). | Examine the status of the NM process and verify the network connectivity. | m=1, t=1800 | Average | Supported | Not required |
| | The number of Pending Containers consistently surpasses the threshold. | The number of pending Containers is $\geq m$, persisting for t seconds ($300 \leq t \leq 2592000$). | Appropriately allocate available resources for YARN tasks. | m=90, t=1800 | Average | Supported | Not required |
| | Cluster memory usage consistently goes over the limit. | Memory Usage $\geq m$, Continuously for t seconds ($300 \leq t \leq 2592000$) | Cluster scale-out | m=85, t=1800 | Critical | Supported | Supported |
| | The cluster's CPU usage consistently exceeds the threshold. | CPU usage rate is $\geq m$, persisting for t seconds ($300 \leq t \leq 2592000$). | Cluster scale-out | m=85, t=1800 | Critical | Supported | Supported |
| | The number of available CPU cores in each queue consistently falls below the threshold. | The number of available CPU cores in any queue is $\leq m$, persisting for t seconds ($300 \leq t \leq 2592000$). | Allocate additional resources to the queue. | m=1, t=1800 | Average | Supported | Not required |

| | | | | | | |
|---|--|---|------------------------|----------|-----------|--------------|
| The available memory in queues consistently goes below the limit. | The available memory in any queue is $\leq m$, persisting for t seconds ($300 \leq t \leq 2592000$). | Allocate additional resources to the queue. | $m=1024$, $t=1800$ | Average | Supported | Not required |
| A primary-secondary switch has occurred in the ResourceManager. | A primary-secondary transition has occurred in the ResourceManager. | Inspect the RM process status and review the standby RM log to determine the cause of the primary-secondary switch. | - | Critical | Supported | Supported |
| The ResourceManager has undergone a full Garbage Collection (GC). | The ResourceManager has undergone a comprehensive Garbage Collection (GC). | Parameter optimization. | - | Critical | Supported | Supported |
| The JVM memory usage of the ResourceManager consistently exceeds the threshold. | The JVM memory usage of the RM has consistently remained greater than or equal to 'm' for a duration of 't' seconds (where $300 \leq t \leq 2592000$). | Adjusting the heap memory size of the ResourceManager. | $m=85$, $t=1800$ | Critical | Supported | Supported |
| The NodeManager has undergone a full Garbage Collection (GC). | The NodeManager has undergone a full Garbage Collection (GC). | Parameter optimization. | - | Average | Supported | Not required |
| The available memory of the NodeManager consistently falls below the threshold. | The available memory of a single NodeManager has consistently remained less than or equal to 'm' for a duration of 't' seconds (where $300 \leq t \leq 2592000$). | Modifying the heap memory size of the NodeManager. | $m=1$, $t=1800$ | Average | Supported | Not required |
| The JVM memory usage of the NodeManager consistently exceeds the threshold. | The JVM memory usage of the NodeManager has consistently been greater than or equal to 'm' for a duration of 't' seconds (where $300 \leq t \leq 2592000$). | Modifying the heap memory size of the NodeManager. | $m=85$, $t=1800$ | Average | Supported | Not required |

| | | | | | | | |
|-------|---|--|--|----------------|----------|-----------|--------------|
| | YARN ResourceManager has no active status | YARN ResourceManager has no active status | Manually check | t=90 | Critical | Supported | Supported |
| | The Yarn Application job has failed to execute. | The Yarn Application job has failed to execute. | Manually check | m=1, t=300 | Average | Supported | Not required |
| HBase | The cluster persistently exceeds the threshold for the number of Regions in Transition (RIT). | The cluster is in a state where the number of RIT Regions is $\geq m$, persisting for t seconds ($300 < t \leq 2592000$). | For HBase versions 2.0 and below, use <code>hbase hbck -fixAssignment</code> . | m=1, t=60 | Critical | Supported | Supported |
| | The number of dead RegionServers in the cluster consistently exceeds the threshold. | The number of dead RegionServers in the cluster is $\geq m$, persisting for t seconds ($300 < t \leq 2592000$). | Manually check | m=1, t=300 | Average | Supported | Supported |
| | The average number of REGIONS per RS in the cluster consistently exceeds the threshold. | The average number of REGIONS per RegionServer in the cluster is $\geq m$, persisting for t seconds ($300 < t \leq 2592000$). | Node expansion or upgrade | m=300, t=1800 | Average | Supported | Supported |
| | HMaster undergoes a full Garbage Collection (GC). | HMaster has undergone a full Garbage Collection (GC). | Parameter optimization. | m=5, t=300 | Average | Supported | Supported |
| | The JVM memory usage of the HMaster consistently exceeds the threshold. | The JVM memory usage of the HMaster is $\geq m$, persisting for t seconds ($300 < t \leq 2592000$). | Adjusting the heap memory size of HMaster. | m=85, t=1800 | Critical | Supported | Supported |
| | The current number of connections to HMaster | The current number of connections to HMaster is $\geq m$, persisting for t seconds ($300 < t \leq 2592000$). | Manually check | m=1000, t=1800 | Average | Supported | Not required |

| | | | | | | | |
|------|---|---|--|----------------|----------|-----------|--------------|
| | consistently exceeds the threshold. | | | | | | |
| | RegionServer undergoes a full Garbage Collection (GC). | RegionServer undergoes a full Garbage Collection (GC). | Parameter optimization. | m=5, t=300 | Critical | Supported | Not required |
| | The JVM memory usage of the RegionServer consistently exceeds the threshold. | The JVM memory usage of the RegionServer is $\geq m$, persisting for t seconds ($300 \leq t \leq 2592000$). | Modifying the heap memory size of RegionServer. | m=85, t=1800 | Average | Supported | Not required |
| | The current number of RPC connections to RegionServer consistently surpasses the threshold. | The current number of RPC connections to RegionServer is $\geq m$, persisting for t seconds ($300 \leq t \leq 2592000$). | Manually check | m=1000, t=1800 | Average | Supported | Not required |
| | The number of RegionServer Storefiles consistently surpasses the threshold. | The number of RegionServer Storefiles is $\geq m$, persisting for t seconds ($300 \leq t \leq 2592000$). | It is recommended to execute a major compaction. | m=5000, t=1800 | Average | Supported | Not required |
| | Both HMaster services in HBASE are in Standby status. | Both HMaster roles are concurrently in Standby status. | Manually check | - | Critical | Supported | Supported |
| | A primary-secondary switch has occurred in HMaster. | A primary-secondary switch has occurred in HMaster. | Investigate through the HMaster service logs. | - | Critical | Supported | Supported |
| Hive | HiveServer2 has undergone a full Garbage Collection (GC). | HiveServer2 has undergone a full Garbage Collection (GC). | Parameter optimization. | m=5, t=300 | Critical | Supported | Supported |

| | | | | | | | |
|-----------|---|--|---|-----------------|----------|-----------|--------------|
| | The JVM memory usage of HiveServer2 consistently exceeds the threshold. | The HiveServer2 JVM memory usage rate is $\geq m$, persisting for t seconds ($300 \leq t \leq 2592000$). | Adjust the heap memory size of HiveServer2. | m=85, t=1800 | Critical | Supported | Supported |
| | A full GC event has occurred in HiveMetaStore. | A full GC event has occurred in HiveMetaStore. | Parameter optimization. | m=5, t=300 | Average | Supported | Supported |
| | A full GC event has occurred in HiveWebHcat. | A full GC event has occurred in HiveWebHcat. | Parameter optimization. | m=5, t=300 | Average | Supported | Supported |
| Zookeeper | The number of Zookeeper connections consistently exceeds the threshold. | The number of Zookeeper connections is $\geq m$, persisting for t seconds ($300 \leq t \leq 2592000$). | Manually check | m=65535, t=1800 | Average | Supported | Not required |
| | The number of ZNode nodes consistently surpasses the threshold. | The number of ZNode nodes is $\geq m$, persisting for t seconds ($300 \leq t \leq 2592000$). | Manually check | m=2000, t=1800 | Average | Supported | Not required |
| | A leader switch has occurred in Zookeeper. | A leader switch has occurred in Zookeeper. | Investigate through the Zookeeper service logs. | - | Critical | Supported | Supported |
| Impala | The ImpalaCatalog JVM memory usage consistently exceeds the threshold. | The ImpalaCatalog JVM memory usage rate is $\geq m$, persisting for t seconds ($300 \leq t \leq 604800$). | Adjust the heap memory size of ImpalaCatalog. | m=0.85, t=1800 | Average | Supported | Not required |
| | The ImpalaDaemon JVM memory usage consistently surpasses the threshold. | The ImpalaDaemon JVM memory usage rate is $\geq m$, persisting for t seconds ($300 \leq t \leq 604800$). | Modify the heap memory size of ImpalaDaemon. | m=0.85, t=1800 | Average | Supported | Not required |
| | The number of client | The number of client connections to the Impala | Adjust the number of fs_service_threads in | m=64, t=120 | Critical | Supported | Support |

| | | | | | | | |
|-----------|---|--|--|------------------|----------|-----------|--------------|
| | connections to the Impala Beeswax API exceeds the threshold. | Beeswax API is $\geq m$. | the <code>impalad.flgs</code> configuration via the console. | | | d | ed |
| | Number of Impala HS2 client connections exceeds the limit | The number of Impala HS2 client connections is $\geq m$. | Adjust the number of <code>fs_service_threads</code> in the <code>impalad.flgs</code> configuration via the console. | $m=64, t=120$ | Critical | Supported | Supported |
| | The runtime of the Query surpasses the threshold. | The runtime of the Query surpasses the threshold of $\geq m$ seconds. | Manually check | - | Critical | Supported | Not required |
| | The total number of failed Query executions exceeds the threshold. | The failure rate of Query execution surpasses the threshold of $\geq m$, with a statistical time granularity of t seconds ($300 \leq t \leq 604800$). | Manually check | $m=1, t=300$ | Critical | Supported | Not required |
| | The total number of submitted Queries exceeds the threshold. | The total number of failed Query executions surpasses the threshold of $\geq m$, with a statistical time granularity of t seconds ($300 \leq t \leq 604800$). | Manually check | $m=1, t=300$ | Critical | Supported | Not required |
| | The failure rate of executing Query exceeds the threshold. | The total number of submitted Queries surpasses the threshold of $\geq m$, with a statistical time granularity of t seconds ($300 \leq t \leq 604800$). | Manually check | $m=1, t=300$ | Critical | Supported | Not required |
| PrestoSQL | Number of failed nodes of PrestoSQL consistently goes over the limit | The number of current failed nodes in PrestoSQL is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 604800$). | Manually check | $m=1, t=1800$ | Critical | Supported | Supported |
| | Number of queued resources in the PrestoSQL resource group consistently goes over the limit | The number of queued tasks in the PrestoSQL resource group is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 604800$). | Parameter optimization. | $m=5000, t=1800$ | Critical | Supported | Supported |
| | Number of | The number of failed | Manually check | $m=1,$ | Crit | Supp | Not |

| | | | | | | | |
|--------|--|--|---|------------------------|--------------|-------------------|-------------------------|
| | failed queries per minute of PrestoSQL goes over the limit | PrestoSQL queries is $\geq m$. | | $t=1800$ | ical | orted | req uire d |
| | Full GC happened in PrestoSQL Coordinator | Full GC happened in PrestoSQLCoordinator | Parameter optimization. | - | Ave rage | Supp orte d | Not req uire d |
| | The JVM memory usage of PrestoSQL Coordinator consistently exceeds the threshold. | The JVM memory usage rate of PrestoSQLCoordinator is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 604800$). | Adjust the heap memory size of PrestoSQLCoordinator . | $m=0.85$, $t=1800$ | Crit ical | Supp orte d | Sup port ed |
| | Full GC occurred in PrestoSQL Worker. | Full GC occurred in PrestoSQLWorker. | Parameter optimization. | - | Ave rage | Supp orte d | Not req uire d |
| | The JVM memory usage of PrestoSQL Worker consistently surpasses the threshold. | The JVM memory usage rate of PrestoSQLWorker is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 604800$). | Adjust the heap memory size of PrestoSQLWorker. | $m=0.85$, $t=1800$ | Crit ical | Supp orte d | Not req uire d |
| Presto | Number of failed nodes of Presto consistently goes over the limit | The number of failed nodes in Presto is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 604800$). | Manually check | $m=1$, $t=1800$ | Crit ical | Supp orte d | Sup port ed |
| | Number of queued resources in the Presto resource group consistently goes over the limit | The number of queued tasks in the Presto resource group is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 604800$). | Parameter optimization. | $m=5000$, $t=1800$ | Crit ical | Supp orte d | Sup port ed |
| | Number of failed queries per minute of Presto goes over the limit | Number of failed queries in Presto is $\geq m$ | Manually check | $m=1$, $t=1800$ | Crit ical | Supp orte d | Not req uire d |
| | Full GC event | Full GC event occurred in PrestoCoordinator | Parameter optimization. | - | Ave rag | Supp orte | Not req |

| | | | | | | | |
|---------|---|--|---|------------------------|----------|-----------|--------------|
| | occurred in PrestoCoordinator | | | | e | d | uire d |
| | Continuous usage rate of JVM memory in PrestoCoordinator exceeds the threshold | The usage rate of JVM memory in PrestoCoordinator is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 604800$). | Adjust the heap memory size of PrestoCoordinator. | $m=0.85$, $t=1800$ | Average | Supported | Supported |
| | Full GC event occurred in PrestoWorker. | Full GC event occurred in PrestoWorker. | Parameter optimization. | - | Average | Supported | Not required |
| | Continuous usage rate of JVM memory in PrestoWorker exceeds the threshold. | The usage rate of JVM memory in PrestoWorker is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 604800$). | Adjust the heap memory size of PrestoWorker. | $m=0.85$, $t=1800$ | Critical | Supported | Not required |
| Alluxio | The total number of current Alluxio Workers consistently falls below the threshold. | The total number of current Alluxio Workers consistently falls below the threshold $\leq m$, persisting for a duration of t seconds ($300 \leq t \leq 604800$). | Manually check | $m=1$, $t=1800$ | Critical | Supported | Not required |
| | Alluxio worker layer resource usage consistently goes over the limit. | The capacity usage rate of the current Alluxio Worker layer is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 604800$). | Parameter optimization. | $m=0.85$, $t=1800$ | Critical | Supported | Not required |
| | A full GC event has occurred in AlluxioMaster. | A full GC event has occurred in AlluxioMaster. | Manually check | - | Average | Supported | Not required |
| | The JVM memory usage rate in AlluxioMaster consistently exceeds the threshold. | The JVM memory usage rate in AlluxioMaster is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 604800$). | Adjust the heap memory size of the AlluxioWorker. | $m=0.85$, $t=1800$ | Critical | Supported | Supported |
| | A full GC event has | A full GC event has occurred in AlluxioWorker. | Manually check | - | Average | Supported | Not req |

| | | | | | | | |
|------|--|--|--|-------------------------|----------|-----------|--------------|
| | occurred in AlluxioWorker. | | | | e | d | uire d |
| | The JVM memory usage rate in AlluxioWorker consistently exceeds the threshold. | The JVM memory usage rate in AlluxioWorker is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 604800$). | Adjust the heap memory size of the AlluxioMaster. | $m=0.85$, $t=1800$ | Critical | Supported | Supported |
| kudu | The degree of imbalance of cluster replicas exceeds the limit. | The degree of imbalance of cluster replicas is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 3600$). | Implement balance among replicas using the rebalance command. | $m=100$, $t=300$ | Average | Supported | Supported |
| | Number of hybrid clock errors exceeds the limit | The number of hybrid clock errors is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 3600$). | Ensure the NTP daemon is operational and network communication with the NTP server is functioning properly. | $m=500000$, $t=300$ | Average | Supported | Supported |
| | The number of tablets in operation exceeds the threshold. | The number of tablets in operation is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 3600$). | An excessive number of tablets on a single node can impact performance. It is advisable to clean unnecessary tables and partitions, or consider appropriate expansion. | $m=1000$, $t=300$ | Average | Supported | Supported |
| | The number of tablets in a failed state exceeds the threshold. | The number of tablets in a failed state is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 3600$). | Verify whether any disks are unavailable or data files are damaged. | $m=1$, $t=300$ | Average | Supported | Supported |
| | Number of failed data directories exceeds the limit. | The number of data directories in a failed state is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 3600$). | Verify whether the paths configured in the <code>fs_data_dirs</code> parameter are available. | $m=1$, $t=300$ | Critical | Supported | Supported |
| | Number of fully-occupied data directories exceeds the limit | The number of fully-occupied data directories is $\geq m$, persisting for a duration of t seconds ($120 \leq t \leq 3600$). | Purge unnecessary data files, or appropriately expand capacity. | $m=1$, $t=120$ | Critical | Supported | Supported |
| | Number of write requests rejected due to | The number of write requests rejected due to queue overload is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 3600$). | Inspect for the presence of write hotspots or a disproportionately low | $m=10$, $t=300$ | Average | Supported | Not required |

| | | | | | | | |
|----------|---|--|--|-----------------|----------|--------------|-----------|
| | queue overload exceeds the limit. | | number of working threads. | | | | |
| | The number of expired scanners exceeds the threshold. | The number of expired scanners is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 3600$). | Upon completion of data reading, remember to invoke the close method of the scanner. | $m=100, t=300$ | Average | Supported | Supported |
| | Number of error logs exceeds the limit | The number of error logs is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 3600$). | Manually check | $m=10, t=300$ | Average | Supported | Supported |
| | The number of RPC requests waiting in the queue that have exceeded the timeout threshold. | The number of RPC requests waiting in the queue that have exceeded the timeout threshold is $\geq m$, persisting for a duration of t seconds ($300 \leq t \leq 3600$). | Inspect whether the system load is excessively high. | $m=100, t=300$ | Average | Supported | Supported |
| Kerberos | The response time of Kerberos consistently exceeds the threshold. | The response time of Kerberos is $\geq m$ (measured in milliseconds), persisting for a duration of t seconds ($300 \leq t \leq 604800$). | Manually check | $m=100, t=1800$ | Critical | Supported | Supported |
| Clusters | The execution of the auto-scaling strategy has failed. | <ol style="list-style-type: none"> The execution of the expansion rule has failed due to an insufficient number of elastic IPs bound to the cluster subnet. The execution of the expansion rule has failed due to an insufficient stock of preset expansion resource specifications. The execution of the expansion rule has failed due to insufficient account balance. Internal error. Please check and try again. | <ol style="list-style-type: none"> Switch to another subnet within the same VPC. Consider switching to a resource specification with ample availability or submit a ticket to contact our internal development team. Recharge your account balance to ensure sufficient funds are available. Submit a ticket to get in touch with our internal development team. | - | Critical | Not required | Supported |
| | Execution of auto-scaling policy timed out | <ol style="list-style-type: none"> The cluster is currently in a cooling-off period, temporarily preventing any scaling operations. The current expiration retry time setting is too short, preventing the rule from triggering any scaling operations within the retry period. | <ol style="list-style-type: none"> Adjust the cooling-off period of the rule. It is recommended to extend the expiration retry period. Please retry later or submit a ticket to contact our | - | Critical | Not required | Supported |

| | | | | | | |
|---|--|---|---|---------|-----------|-----------|
| | 3. The cluster status is not in a state that prevents scaling. | internal development team. | | | | |
| The auto-scaling policy is not triggered. | <ol style="list-style-type: none"> Without setting the resource specifications for scaling, the scaling rule cannot be triggered. The elastic resources have reached the maximum node limit, preventing the triggering of scaling. The elastic resources have reached the minimum node limit, preventing the triggering of downscaling. The execution time range for time-based scaling has expired. Without elastic resources in the cluster, the downscaling rule cannot be triggered. | <ol style="list-style-type: none"> To add a scaling specification configuration, please set at least one elastic resource specification. Elastic resources have exceeded the maximum node limit. If further expansion is required, consider adjusting the maximum node limit. Elastic resources have reached the minimum node limit. If further contraction is required, consider adjusting the minimum node limit. If you wish to continue using this rule for automatic scaling, please modify the effective time range of the rule. Execute the downscaling rule after supplementing the elastic resources. | - | Average | Supported | Supported |
| The automatic scaling expansion was partially successful. | <ol style="list-style-type: none"> The resource inventory is less than the expansion quantity, thus only a portion of the resources has been supplemented. The expansion quantity exceeds the actual delivery quantity, thus only a portion of the resources has been supplemented. The expansion of elastic resources has reached the maximum node limit, thus the execution of the expansion rule was partially successful. The reduction of elastic resources has reached the minimum node limit, thus the execution of the reduction rule was partially successful. The elastic IP of the | <ol style="list-style-type: none"> Manually expand the inventory of sufficient resources to supplement the lack of required resources. Manually expand the inventory of sufficient resources to supplement the lack of required resources. Elastic resources have exceeded the maximum node limit. If further expansion is required, consider adjusting the maximum node limit. | - | Average | Supported | Supported |

| | | | | | | |
|---|---|--|---|----------|-----------|--------------|
| | <p>subnet bound to the cluster is insufficient, resulting in a failure to replenish resources.</p> <p>6. The inventory of the preset expansion resource specification is insufficient, resulting in a failure to replenish resources.</p> <p>7. The account balance is insufficient, resulting in a failure to replenish resources.</p> | <p>4. Elastic resources have reached the minimum node limit. If further contraction is required, consider adjusting the minimum node limit.</p> <p>5. Switch to another subnet within the same VPC.</p> <p>6. You may attempt to replace it with a more abundant resource specification or submit a ticket to contact the internal development team.</p> <p>7. Recharge your account balance to ensure sufficient funds are available.</p> | | | | |
| Anomaly detected in the JVM OLD region. | Anomaly detected in the JVM OLD region. | Manually check | <p>1. The OLD region has been at 80% capacity continuously for 5 minutes or more.</p> <p>2. The JVM memory usage has reached 90%.</p> | Critical | Supported | Supported |
| Service role health status has exceeded the timeout period. | The health status of the service role has exceeded the timeout period, with a duration of 't' seconds ($180 \leq t \leq 604800$). | <p>The health status of the service role has been exceeding the timeout period continuously on a minute-by-minute basis.</p> <p>Recommended action: Review the log information for the corresponding service role and take action based on the log details.</p> | t=300 | Average | Supported | Not required |

| | | | | | | |
|-----------------------------------|---|---|-------|----------|--------------|-----------|
| Service role status abnormal | The health status of the service role is abnormal, with a duration of 't' seconds (180 <= t <= 604800). | The health status of the service role has been continuously unavailable on a minute-by-minute basis. Recommended action: Review the log information for the corresponding service role and take action based on the log details. | t=300 | Critical | Supported | Supported |
| Auto-scaling policy expired | Auto-scaling policy expired | Manually check | / | Average | Not required | Supported |
| Node role process restarted | Node role process restarted | Manually check | / | Average | Not required | Supported |
| Bootstrap script execution failed | Bootstrap script execution failed | Manually check | / | Average | Not required | Supported |

Logs

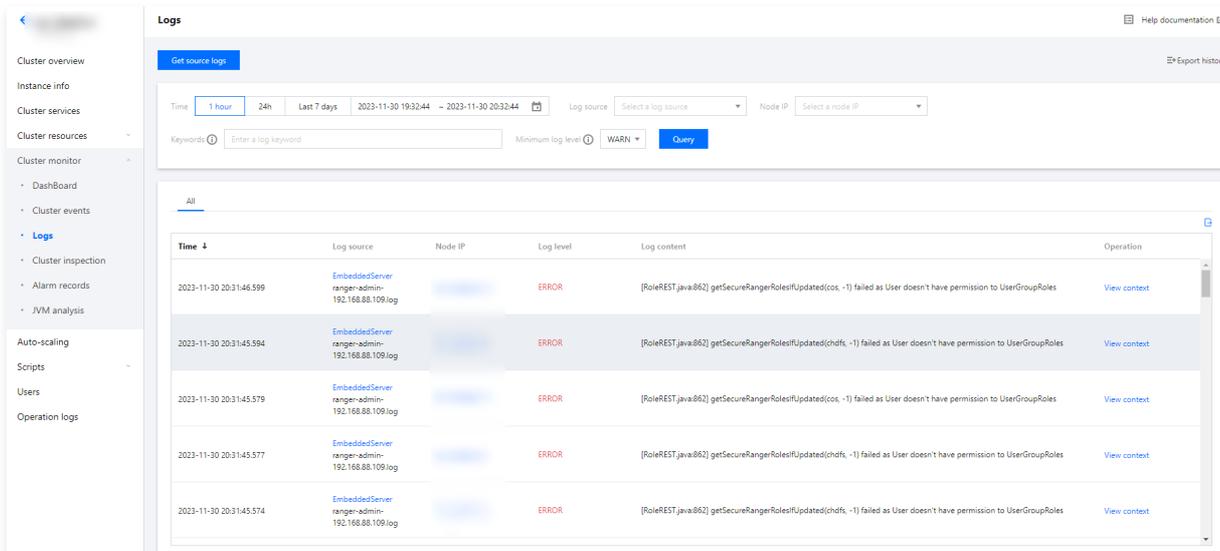
Last updated: 2023-12-21 16:22:13

Feature Overview

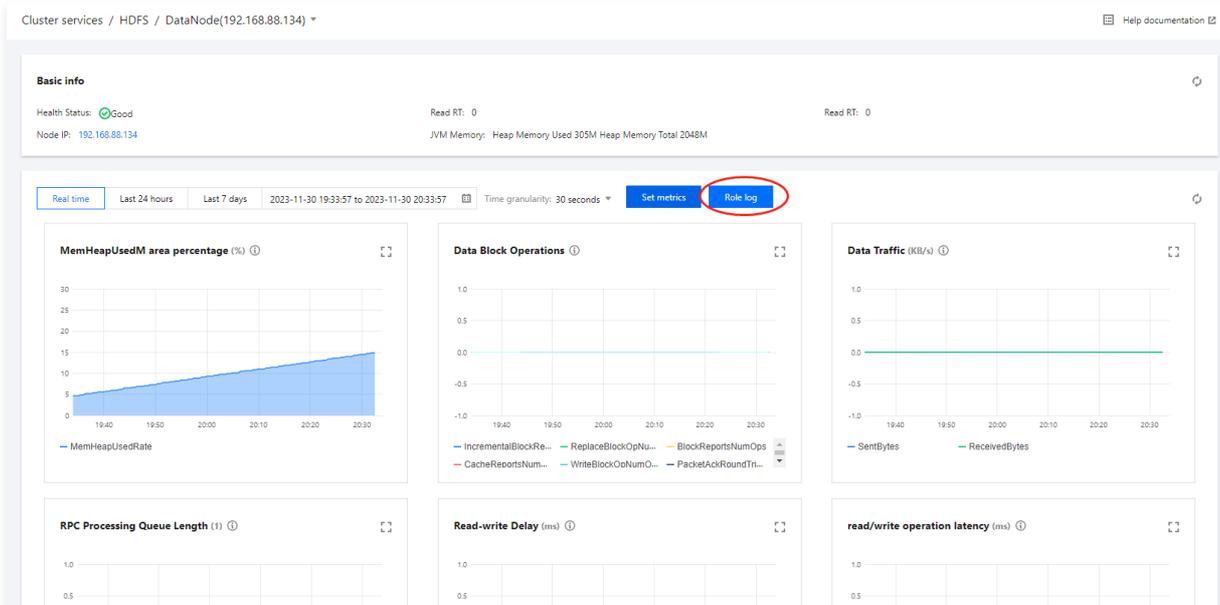
The log function offers the capability to gather, search, and export operational logs of components, supporting keyword searches for core service logs and node system logs within the current cluster. This allows for swift examination of crucial service logs without the necessity of logging into the node.

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding cluster ID/Name to access the cluster details page.
2. In the cluster details page, select **Cluster Monitoring > Logs** to filter and view log content based on the current cluster, log files, node IP, and time range.



Alternatively, in the cluster details page, select **Cluster Services>Component Card>Role Management**, view the role list, select a **specific node IP** and click to jump to the node monitoring metrics display page. Click on **Role Logs** to navigate to the log search page.



Click on **Node IP** to navigate to the corresponding node status page. Click on **Log Source** to navigate to the corresponding node monitoring metrics display page.

Keyword Explanation:

- Supports full-text keyword search.

- Supports search using special characters - . * > < = ! () { } / .
- Supports phrase search, for example: `address=/ip:port .`

3. When troubleshooting, it is often necessary to pay attention to the context of the keyword in the logs. On the log search page, click on **View Context** to access the log context page.

Log context
✕

Role EmbeddedServer Node IP 192.168.88.109

Log path /data/emr/ranger/logs/ranger-admin-192.168.88.109-user_IS_UNDEFINED.log

Time Go to latest entry

↻
↓

| | | | | | |
|----|------------|--------------|-------|---------------------|------------------------------|
| 1 | 2023-11-30 | 20:32:59.843 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 2 | 2023-11-30 | 20:32:59.843 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 3 | 2023-11-30 | 20:32:59.844 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 4 | 2023-11-30 | 20:33:00.854 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 5 | 2023-11-30 | 20:33:00.854 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 6 | 2023-11-30 | 20:33:00.854 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 7 | 2023-11-30 | 20:33:00.855 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 8 | 2023-11-30 | 20:33:01.866 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 9 | 2023-11-30 | 20:33:01.866 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 10 | 2023-11-30 | 20:33:01.866 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 11 | 2023-11-30 | 20:33:01.867 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 12 | 2023-11-30 | 20:33:01.883 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 13 | 2023-11-30 | 20:33:01.883 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 14 | 2023-11-30 | 20:33:01.884 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 15 | 2023-11-30 | 20:33:01.884 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 16 | 2023-11-30 | 20:33:02.894 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 17 | 2023-11-30 | 20:33:02.894 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 18 | 2023-11-30 | 20:33:02.895 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 19 | 2023-11-30 | 20:33:02.895 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 20 | 2023-11-30 | 20:33:03.906 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 21 | 2023-11-30 | 20:33:03.906 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 22 | 2023-11-30 | 20:33:03.907 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 23 | 2023-11-30 | 20:33:03.907 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 24 | 2023-11-30 | 20:33:04.918 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 25 | 2023-11-30 | 20:33:04.918 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 26 | 2023-11-30 | 20:33:04.919 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 27 | 2023-11-30 | 20:33:04.919 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 28 | 2023-11-30 | 20:33:04.938 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |
| 29 | 2023-11-30 | 20:33:04.938 | ERROR | [RoleREST.java:862] | getSecureRangerRolesIfUpdate |

4. Supports filtering log results and exporting source logs, which can be downloaded locally.

Supported service types for log search.

Note

- Currently, only a 30-day log search is supported.
- If the cluster has not enabled log collection and you wish to enable it, please contact your dedicated after-sales service.
- Only a select few services have audit log collection enabled by default. If your cluster requires the activation of audit log collection, please reach out to your dedicated after-sales support.

| Component | Role | Logs | Note |
|-----------|------------------|--|----------------------------------|
| HDFS | NameNode | /data/emr/hdfs/logs/hadoop-hadoop-namenode.log | Runtime logs of NameNode |
| | NNAudit | /data/emr/hdfs/logs/nn-audit.log | Audit logs of NameNode |
| | ZKFC | /data/emr/hdfs/logs/hadoop-hadoop-zkfc.log | Runtime logs of ZKFC |
| | ZKFCAudit | /data/emr/hdfs/logs/SecurityAuth-zkfc-audit.log | Audit logs of ZKFC |
| | DataNode | /data/emr/hdfs/logs/hadoop-hadoop-datanode.log | Runtime logs of DataNode |
| | DNAudit | /data/emr/hdfs/logs/SecurityAuth-datanode-audit.log | Audit logs of DataNode |
| | JournalNode | /data/emr/hdfs/logs/hadoop-hadoop-journalnode.log | Runtime logs of JournalNode |
| | JNAudit | /data/emr/hdfs/logs/SecurityAuth-journalnode-audit.log | Audit logs of JournalNode |
| YARN | DFSRouter | /data/emr/hdfs/logs/hadoop-hadoop-dfsrouter.log | Runtime logs of DFSRouter |
| | ResourceManager | /data/emr/yarn/logs/yarn-hadoop-resourcemanager.log | Runtime logs of ResourceManager |
| | RMAudit | /data/emr/yarn/logs/rm-audit.log | Audit logs of ResourceManager |
| | NodeManager | /data/emr/yarn/logs/yarn-hadoop-nodemanager.log | Runtime logs of NodeManager |
| | NMAudit | /data/emr/yarn/logs/nm-audit.log | Audit logs of NodeManager |
| | JobHistoryServer | /data/emr/yarn/logs/mapred-hadoop-historyserver.log | Runtime logs of JobHistoryServer |
| HBase | TimelineServer | /data/emr/yarn/logs/yarn-hadoop-timelineserver.log | Runtime logs of TimelineServer |
| | HMaster | /data/emr/hbase/logs/hbase-hadoop-master.log | Runtime logs of HMaster |

| | | | |
|------------|-------------------|---|-----------------------------------|
| | HMAudit | /data/emr/hbase/logs/SecurityAuth-hmaster-audit.log | Audit logs of HMaster |
| | ThriftServer | /data/emr/hbase/logs/hbase-hadoop-thrift.log | Runtime logs of ThriftServer |
| | RegionServer | /data/emr/hbase/logs/hbase-hadoop-regionserver.log | Runtime logs of RegionServer |
| | RSAudit | /data/emr/hbase/logs/SecurityAuth-regionserver-audit.log | Audit logs of RegionServer |
| ClickHouse | ClickHouse-server | /data/clickhouse/clickhouse-server/logs/clickhouse-server.log | Runtime logs of ClickHouse-server |
| Druid | Broker | /data/emr/druid/var/log/druid/broker.log | Runtime logs of Broker |
| | Coordinator | /data/emr/druid/var/log/druid/coordinator.log | Runtime logs of Coordinator |
| | Router | /data/emr/druid/var/log/druid/router.log | Runtime logs of Router |
| | Overload | /data/emr/druid/var/log/druid/overload.log | Runtime logs of Overload |
| | Historical | /data/emr/druid/var/log/druid/historical.log | Runtime logs of Historical |
| | MiddleManager | /data/emr/druid/var/log/druid/middleManager.log | Runtime logs of MiddleManager |
| Zookeeper | Zookeeper | /data/emr/zookeeper/logs/zookeeper-root-server.log | Runtime logs of Zookeeper |
| | ZKAudit | /data/emr/zookeeper/logs/zookeeper-audit.log | Audit logs of Zookeeper |
| Hive | HiveServer2 | /data/emr/hive/logs/hadoop-hive | Runtime logs of HiveServer2 |
| | HS2Audit | /data/emr/hive/logs/hadoop_hive_server2_audit | Audit logs of HiveServer2 |
| | HMSAudit | /data/emr/hive/logs/hadoop-hivemetastore | Audit logs of HiveMetaStore |
| KUDU | KuduMaster | /data/emr/kudu/logs/kudu-master.WARNING | Runtime logs of KuduMaster |
| | KuduServer | /data/emr/kudu/logs/kudu-tserver.WARNING | Runtime logs of KuduServer |
| Alluxio | AlluxioMaster | /data/emr/alluxio/logs/master.log | Runtime logs of AlluxioMaster |
| | AlluxioWorker | /data/emr/alluxio/logs/worker.log | Runtime logs of AlluxioWorker |
| Ranger | EmbeddedServer | /data/emr/ranger/logs/ranger-admin.log | Runtime logs of EmbeddedServer |
| CosRanger | CosRangerServer | /usr/local/service/cosranger/log/info.log | Runtime logs of CosRanger |
| Impala | Catalogd | /data/emr/impala/logs/catalogd.INFO | Runtime logs of Catalogd |
| | Statedored | /data/emr/impala/logs/statedored.INFO | Runtime logs of Statedored |
| | Impalad | /data/emr/impala/logs/impalad.INFO | Runtime logs of Impalad |
| Spark | HistoryServer | /data/emr/spark/logs/spark-hadoop.log | Runtime logs of HistoryServer |
| Kylin | Kylin | /data/emr/kylin/logs/kylin.log | Runtime logs of Kylin |
| Zeppelin | ZeppelinServer | /data/emr/zeppelin/logs/zeppelin-hadoop.log | Runtime logs of ZeppelinServer |
| Knox | Gateway | /data/emr/knox/logs/gateway.log | Runtime logs of Gateway |
| Doris | BrokerBootstrap | /data/emr/doris/broker/log/apache_hdfs_broker.log | Runtime logs of BrokerBootstrap |
| | PaloFe | /data/emr/doris/fe/log/fe.log | Runtime logs of PaloFe |
| | PaloBe | /data/emr/doris/be/log/be.INFO | Runtime logs of PaloBe |
| Kafka | Kafka | /usr/local/service/kafka/logs/server.log | Runtime logs of Kafka |
| | | | |

| | | | |
|-----------|--------------|--|------------------------------|
| Kyuubi | KyuubiServer | /data/emr/kyuubi/logs/kyuubi.log /data/emr/kyuubi/logs/kyuubi-hadoop-org.apache.kyuubi.server.KyuubiServer-hostip.out | Runtime logs of KyuubiServer |
| Trino | Coordinator | /data/emr/trino/logs/coordinator.log | Runtime logs of Coordinator |
| | Worker | /data/emr/trino/logs/worker.log | Runtime logs of the Worker |
| StarRocks | FE | /data/emr/starrocks/fe/log/fe.log | Runtime logs of FE |
| | BE | /data/emr/starrocks/be/log/be.INFO | Runtime logs of BE |
| | Broker | /data/emr/starrocks/broker/log/apache_hdfs_broker.log | Runtime logs of Broker |
| | CN | /data/emr/starrocks/be/log/cn.INFO | Runtime logs of CN |

The service supports the minimum log level.

| Service | Default collection of the lowest log level. |
|---------------------|---|
| Impala、Kudu | INFO |
| Additional Services | WARN |

Query rules for the lowest log level.

| Minimum log level | Queryable log levels. |
|-------------------|-----------------------|
| INFO | INFO、WARN、ERROR、FATAL |
| WARN | WARN、ERROR、FATAL |
| ERROR | ERROR、FATAL |
| FATAL | FATAL |

Application Analysis

HDFS File Storage Analysis

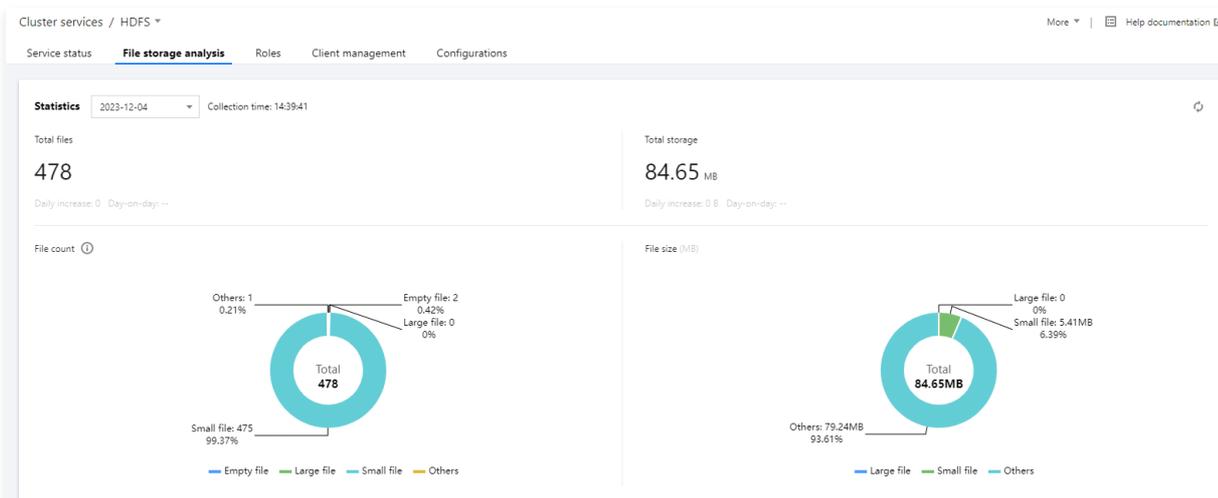
Last updated: 2023-12-21 16:22:46

Feature Overview

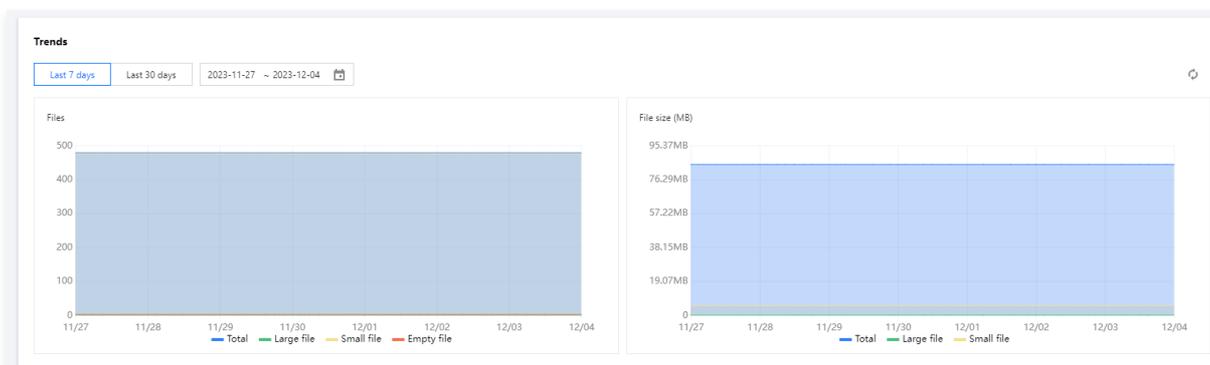
This feature allows for the creation of built-in NameService storage data based on clusters, enabling the viewing of total files, total storage volume, distribution information, and recent trends of HDFS files collected on T-1 day, as well as the top directory lists of large and small files.

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding **Cluster ID/Name** to enter the cluster details page.
2. From the cluster details page, click on **Cluster Services**, then select **Operations > File Storage Analysis** in the top right corner of the HDFS component. This provides information related to files and directories stored on HDFS up until the last collection time point.
3. Statistics
 - 3.1 You can view the total number of HDFS stored files, the daily increment of total storage volume, and the day-on-day comparison volume.
 - 3.2 Refer to the definitions of empty files (=0), small files (<=2M), others (2M<file storage volume<128M), and large files (>=128M) to provide views of file quantity distribution and file storage volume distribution.



4. Through the view, you can intuitively observe the recent historical trends in the quantity and storage volume of various types of files.



5. Query the relevant dimension information of the Top 1000 small/large files at the DayT-1 collection time point. This provides information for querying and downloading such as file name, path, user group, owner, size, and the most recent access time.

Small file Large file

Collection time 2023-12-04 14:39:41 Sort by size Top1000 Download Refresh

| FileName | Path | GROUP | Owner | Size | LastAccessTime | LastModifyTime |
|---------------------------------|--|------------|--------|---------|---------------------|---------------------|
| AMRMTokenSecretManagerNode | /data/emr/hdfs/tmp/yarn/system/mstore/F... | supergroup | hadoop | 23 B | 2023-08-22 18:49:25 | 2023-08-22 18:49:25 |
| AMRMTokenSecretManagerNode | /data/emr/hdfs/tmp/yarn/system/mstore/F... | supergroup | hadoop | 23 B | 2023-08-22 18:49:25 | 2023-08-22 18:49:25 |
| EpochNode | /data/emr/hdfs/tmp/yarn/system/mstore/F... | supergroup | hadoop | 2 B | 2023-07-26 18:34:25 | 2023-07-26 18:34:25 |
| EpochNode | /data/emr/hdfs/tmp/yarn/system/mstore/F... | supergroup | hadoop | 2 B | 2023-07-26 18:34:25 | 2023-07-26 18:34:25 |
| application_1689696505053_00... | /data/emr/hdfs/tmp/yarn/system/mstore/F... | supergroup | hadoop | 1.83 KB | 2023-07-26 17:42:51 | 2023-07-19 00:12:05 |
| application_1689696505053_00... | /data/emr/hdfs/tmp/yarn/system/mstore/F... | supergroup | hadoop | 1.83 KB | 2023-07-26 17:42:51 | 2023-07-19 00:12:05 |
| appattemp_1689696505053_0... | /data/emr/hdfs/tmp/yarn/system/mstore/F... | supergroup | hadoop | 370 B | 2023-07-26 17:42:51 | 2023-07-19 00:12:38 |
| appattemp_1689696505053_0... | /data/emr/hdfs/tmp/yarn/system/mstore/F... | supergroup | hadoop | 370 B | 2023-07-26 17:42:51 | 2023-07-19 00:12:38 |
| application_1689696505053_00... | /data/emr/hdfs/tmp/yarn/system/mstore/F... | supergroup | hadoop | 1.92 KB | 2023-07-26 17:42:51 | 2023-07-19 00:12:38 |
| application_1689696505053_00... | /data/emr/hdfs/tmp/yarn/system/mstore/F... | supergroup | hadoop | 1.92 KB | 2023-07-26 17:42:51 | 2023-07-19 00:12:38 |

Total items: 950 10 / page 1 / 95 pages

Export settings Refresh

File path: Default path / [Edit](#) Export files [Export latest files](#)

Risk Disclosure

The analysis data required for file storage analysis will begin to be collected at 14:00 Beijing time every day.

1. File storage analysis involves the collection and analysis of backup fsimage files. This analysis can increase the use of local memory (up to a maximum increase of 4G). If the total proportion of memory usage in the cluster is consistently high, you can [submit a work order](#) to disable this feature.
2. In an HA cluster, this analysis function is executed on the Standby Master node, while in a non-HA cluster, this analysis function is executed on the Master node.

StarRocks Query Management

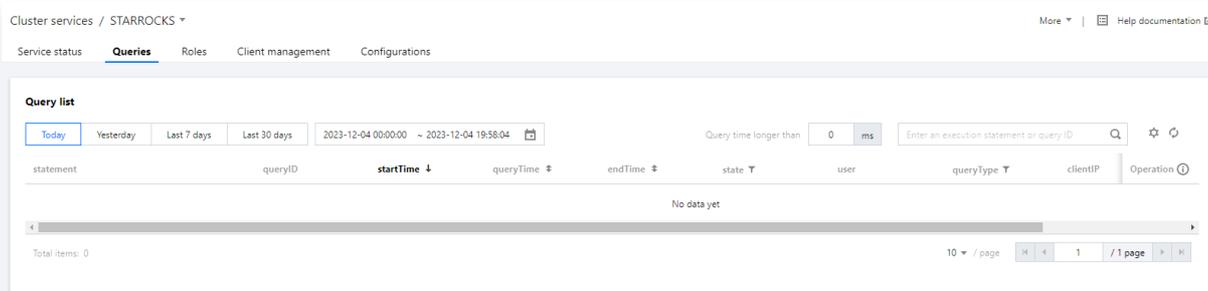
Last updated: 2023-12-21 16:22:57

Feature Overview

StarRocks provides a multi-dimensional metric and detail display for queries. The query list allows for swift examination of query statements, query start times, query statuses, query durations, users, number of rows retrieved, total CPU time, total memory consumption, and many other detailed metrics. Additionally, related SQL, query plans, and Profiles can be viewed through query details.

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding StarRocks **Cluster ID/Name** to access the cluster details page.
2. Within the cluster details page, click on **Cluster Services**, then select **Operations > Query Management** in the upper right corner of the StarRocks component. This allows for the viewing of related query lists. Some column header fields support filtering or sorting functions, support compound filtering operations across multiple dimensions, and allow for the customization of query duration for slow query filtering.



3. By clicking on **Operations Column > Details** at the far right of the query list, you can view the query statement, query plan, and Profile.

Note:

- StarRocks Query Management currently only supports StarRocks V1.4.0 and above. After September 12, 2023, the query management feature will be enabled by default for newly purchased StarRocks V1.4.0 and above clusters, and an `emr_admin` user will be automatically created for your StarRocks cluster for system query information collection. If you have a StarRocks V1.4.0 cluster created before September 12, 2023 and wish to use StarRocks Query Management, please [submit a ticket](#) to contact our after-sales team for evaluation and activation.
- Before executing a query, input `set enable_profile=true` in the command to print the Profile of the current session.

Hive Query Management

Last updated: 2023-12-21 16:23:09

Feature Overview

Hive Query Management offers a query overview and list, with support for the activation of Hive SQL Insights. This aids in multi-dimensional perception of query distribution and viewing multi-dimensional data information in the query list. Additionally, it supports the query insight function of HiveSQL, with various insight items effectively assisting in the analysis of potential query issues.

Note

- Clusters purchased after November 1, 2023, that select the Hive service will all support the query insight function. If clusters that have deployed the Hive component before November 1, 2023, need to use the Hive service query insight function for query optimization and cost reduction, please activate it by [submitting a work order](#). Activation note: An evaluation is required, as there may be situations such as restarting the HiveServer2 service due to an upgrade package.
- For clusters purchased before June 27, 2023, if you need to view additional percentile indicators and list information such as submitted IP, HS2 IP, compilation duration, query table, number of partitions, scanned data volume/number of rows, please activate it by [submitting a work order](#).
- Clusters purchased after June 8, 2022, that concurrently select Hive, all support the query management function. If clusters that have deployed the Hive component before June 8, 2022, need to use Hive query management, they must contact after-sales personnel for evaluation.

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding cluster **ID/Name** to access the cluster details page.
2. In the cluster details page, click on the **Cluster Service Directory**, then select the Hive service name or **Operations > Query Management** in the upper right corner of the Hive service. The query overview supports percentile query distribution observation of multi-dimensional indicators, and multi-dimensional analysis of stage indicators in the query cycle can be performed through filtering in the query list. Stage indicators include: submitted IP, HS2 IP, compilation duration, query table, number of partitions, scanned data volume/number of rows, execution duration, etc.

3. Clicking on the **Application ID** will redirect you to the Yarn application management and select the Yarn job set associated with the Hive query.
4. Clicking on **Details** allows you to view information such as the query statement, query plan, execution overview, Profile, and Error logs (in the case of a Failed status).
5. Activate the Query Insight feature, and you can view the Insight anomaly information in **Query Operation Column > Insight**. You can also view the anomaly distribution and anomaly information list of Insight items through the Insight overview.

HBase Table Analysis

Last updated: 2024-01-12 11:18:25

Feature Overview

Data Table Analysis offers dimensional information such as read/write request volume and storage status for Hbase table-level, within-table Regions, and RegionServers. Concurrently, it provides Region analysis, supporting read QPS, write QPS information, and historical trend analysis for the corresponding table or RegionServer, in conjunction with practical scenarios.

Instructions

1. Log into the [EMR Console](#), and in the cluster list, click on the corresponding **Cluster ID/Name** to access the cluster details page.
2. Within the cluster details page, click on **Cluster Services**, then select **Operations > Data Table Analysis** in the upper right corner of the HBase component to conduct related HBase data table load queries.

Tables

The Hbase data table list allows for the viewing of table-level request QPS, write QPS, MemStore storage volume, StoreFile size, and other information. By using the sorting button in the title column, the top data tables in the cluster can be located.

Cluster services / HBASE ▾ More ▾ | Help documentation 📄

Service status **Table analysis** Roles Client management Configurations RIT fixing

Tables

🔍 ↻

| Table | Read Requests/s ↕ | Write Requests/s ↕ | Memstore Size ↕ | Storefile Size ↕ | Operation |
|---------|-------------------|--------------------|-----------------|------------------|---|
| hbase_* | 0.04 | 0 | 768 | 11923 | Regions RegionServers |
| hbase_* | 0 | 0 | 256 | 0 | Regions RegionServers |
| hbase_* | 0 | 0 | 256 | 4926 | Regions RegionServers |

Total 3 items Lines per page 10 ▾ | 1 / 1 page

Statistics by Region

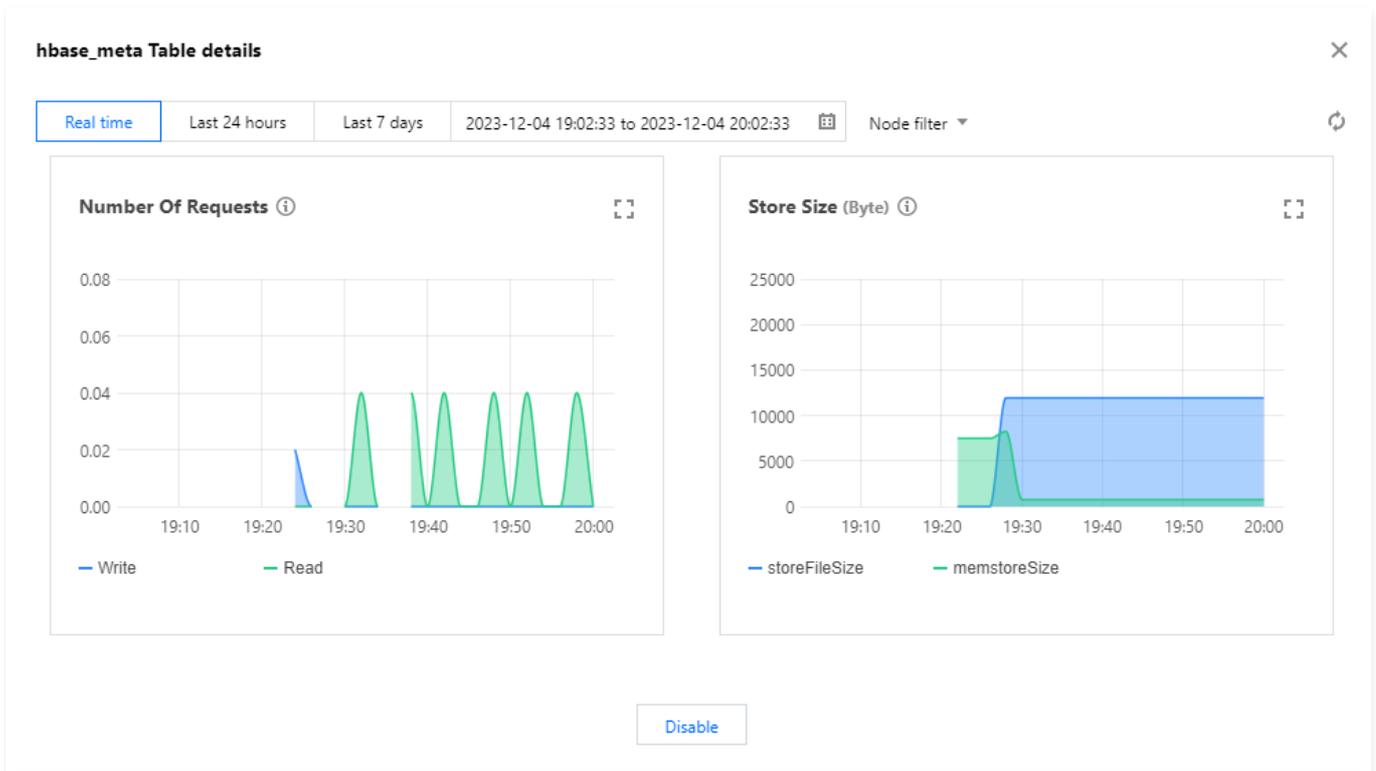
1 hour 6 hours 12h 2023-12-04 19:01:19 ~ 2023-12-04 20:01:19 📅 Search by table name 🔍 ↻

| Region | TableName | RegionServer ▾ | ReadRequestRate 📉 ↓ | WriteRequestRate 📉 ↕ |
|--------|-----------------|----------------|---------------------|----------------------|
| ... | hbase_meta | 186:6003 | 0.01 | 0 |
| ... | hbase_acl | 242:6003 | 0 | 0 |
| ... | hbase_namespace | 160:6003 | 0 | 0 |

Total items: 3 10 ▾ / page | 1 / 1 page

View Table Details

Click on the corresponding table name to display the table details. The details page can show the request volume (including reads and writes) and store size (including memstore and storeFile) of the selected table by the entire table and node dimensions. Select the node filter in the upper right corner to switch nodes for viewing.



Regions Operations

Click on **Regions Operations** to view the read and write request volume of each Region contained in the table, and to locate the hotspot situations within the table's regions.

Cluster services / HBASE / hbase_meta / **Region list** Help documentation

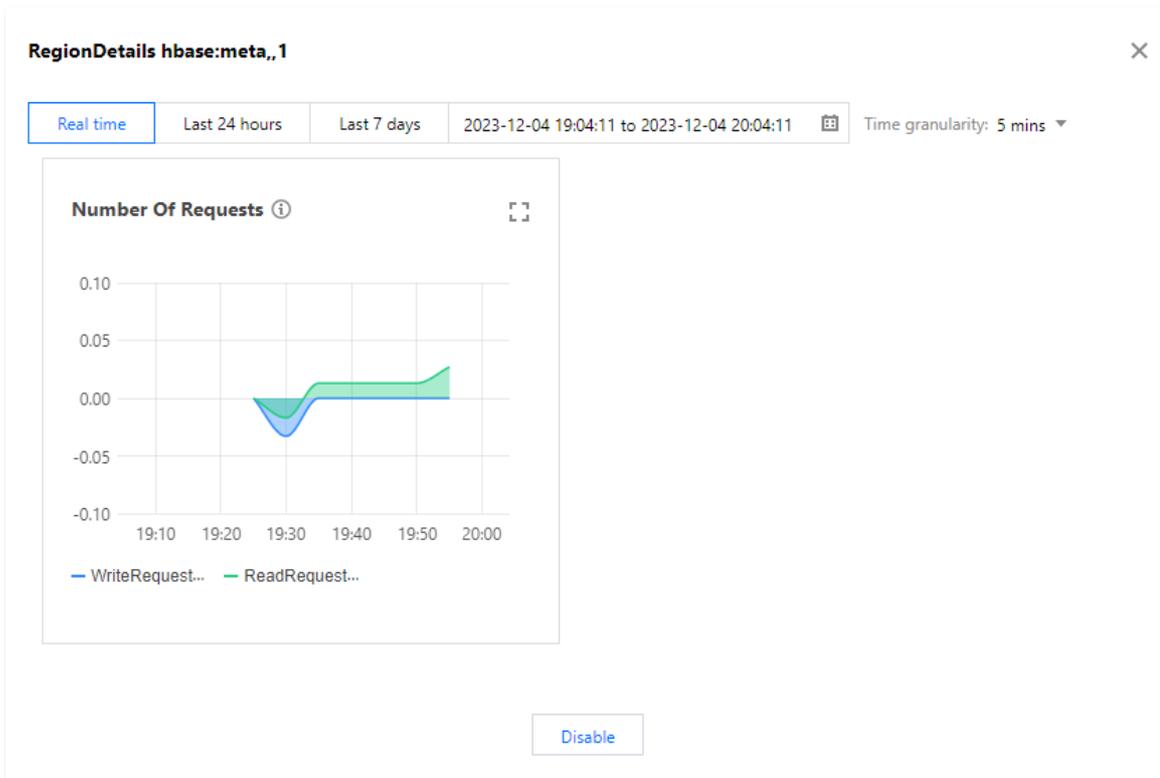
Refresh

| Region Name | RegionServer | Start Key | End Key | Read Request Count/s | Write Request Count/s |
|-------------|--------------|-----------|---------|----------------------|-----------------------|
| | 86:6003 | - | - | 0.015 | 0 |

Total 1 item Lines per page 10 1 / 1 page

Region Details

Click on the corresponding Region name to display the Region details and view the metric trends. The details page can show the request volume (including reads and writes) metric data of the selected table by different time granularities. Select the time granularity in the upper right corner to switch granularities for viewing.



RegionServers Operations

Click on **RegionServers Operations** to view the request latency of each RegionServer where the table is distributed.

Cluster services / HBASE / hbase_namespace / **Regionserver list** Help documentation

Refresh

| RegionServer | GetTimeTp99 | ScanTimeTp99 | PutTimeTp99 | IncrementTimeTp99 | AppendTimeTp99 | DeleteTimeTp99 |
|--------------|-------------|--------------|-------------|-------------------|----------------|----------------|
| 160 | 0 | 0 | 0 | 0 | 0 | 0 |

Total 1 item Lines per page 10

Region Analysis

Region Analysis allows for the retrieval of the associated table or filtering of the associated RegionServer, pinpointing the distribution of hotspot requests in the cluster through average request QPS and average read-write QPS information.

Statistics by Region

1 hour | 6 hours | 12h | 2023-12-04 19:05:47 ~ 2023-12-04 20:05:47

Search by table name

| Region | TableName | RegionServer | ReadRequestRate | WriteRequestRate |
|--------|-----------------|--------------|-----------------|------------------|
| | hbase_meta | 86:6003 | 0.01 | 0 |
| | hbase_acl | 42:6003 | 0 | 0 |
| | hbase_namespace | 60:6003 | 0 | 0 |

Total items: 3 10 / page

Clicking on the view button for the column titles Average Read QPS and Average Write QPS allows for the observation of the historical trends of the Region record metrics on the current page, monitoring abrupt request information, and supports the selection of time intervals.

Kyuubi Query Management

Last updated: 2023-12-21 16:23:23

Feature Overview

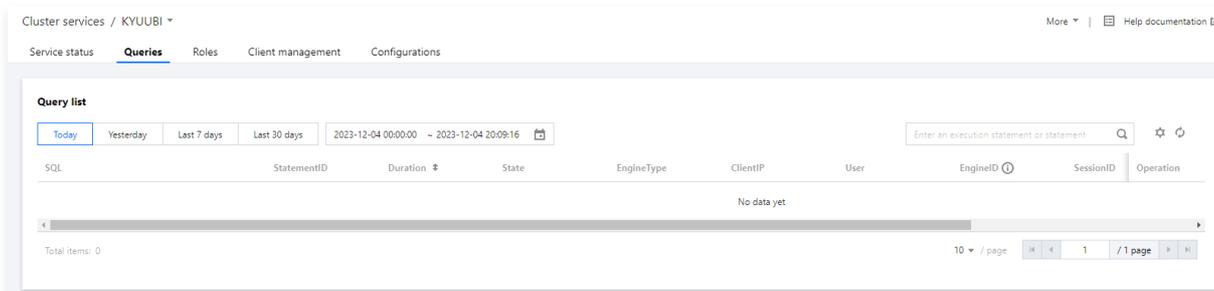
Kyuubi Query Management furnishes a multi-dimensional information array, inclusive of executed statements and execution engine IDs in the query list, thereby facilitating users with a swift comprehension of service query submissions and operational records.

Note:

Post November 1, 2023, all newly purchased clusters synchronously selecting Kyuubi will support the query management feature. For clusters that have deployed the Kyuubi component prior to November 1, 2023, and wish to utilize the Kyuubi query management, please initiate the process via [submitting a work order](#). Please note, an evaluation is required and situations may arise necessitating the restart of Kyuubi services due to upgrade packages.

Instructions

1. Log in to the [EMR Console](#), and within the cluster list, click on the corresponding cluster ID/Name to access the cluster details page.
2. Within the cluster details page, click on the **Cluster Service Directory**, then select the **Kyuubi Service Name** or **Operation > Queries** in the upper right corner of the Kyuubi service. Through the query list's filter, multi-dimensional analysis can be conducted, including the computation engine, submission IP, submission user, and more.



3. Click on Details to view information related to the query statement.

Cluster Inspection

Last updated: 2023-12-21 16:23:38

Feature Overview

Each cluster can conduct health checks on its nodes and services in real-time or on a scheduled basis (daily or weekly) based on the selected inspection items. Each cluster can only configure one regular inspection task to periodically grasp the health status of the cluster and promptly handle any anomalies or risk points.

When setting up an immediate inspection or scheduled inspection task, the standard inspection items are selected by default. In special cases where service functions need to be inspected, you can select the additional inspection items as needed. However, service function inspections consume cluster performance and it is not recommended to conduct resource-intensive inspections during peak business periods.

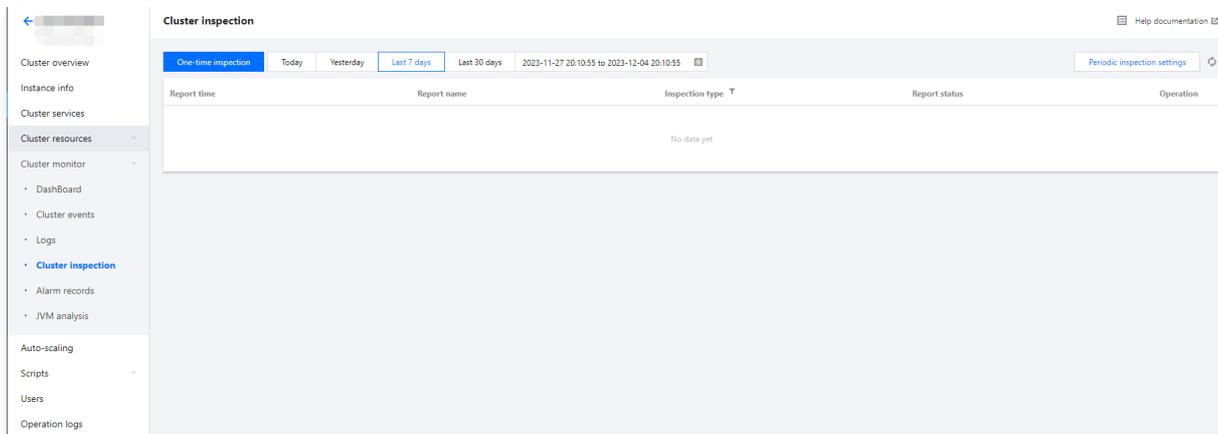
Upon completion of each inspection task, a PDF format inspection report is generated. Users can download or delete the inspection report. Each primary account can retain up to 50 inspection reports. Reports that exceed the maximum retention limit will result in the deletion of older reports.

Note

- The inspection system supports service inspection items, currently only supporting HDFS, Yarn, Hbase, Hive, Impala, Zookeeper, Alluxio, Kudu, and Kerberos.
- Unable to change the settings when the periodic inspection task is being executed.

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding cluster ID/Name to access the cluster details page.
2. In the cluster details page, select **Cluster Monitoring > Cluster Inspection** to conduct a health check based on the current cluster's nodes and services. Each cluster can only configure one regular inspection task, or you can click **Instant Inspection** to perform an inspection. To configure a regular inspection task, click **Regular Inspection Settings**.



- **One-time inspection:** Instant inspection involves checking the health status of the cluster's nodes and services from a certain moment to the current time and generating an inspection report.

Start one-time inspection ✕

i 1. 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.

2. 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.

3. All enabled event monitoring policies can be set as inspection options by default. You can adjust them on the [Event policies](#) tab.

i Unable to change the settings when the periodic inspection task is being executed.

Inspection interval: to date

Inspection item: Select all Invert all

- NODE
- HDFS
- YARN
- HIVE
- CLUSTER
- ZOOKEEPER

- **Periodic inspection:** Once the regular inspection policy is activated, the system will automatically check the health status of the cluster nodes and services within each inspection cycle and generate an inspection report. Each cluster can configure one regular inspection policy.

Periodic inspection settings ✕

i 1. 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.

2. 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.

3. All enabled event monitoring policies can be set as inspection options by default. You can adjust them on the [Event policies](#) tab.

i Unable to change the settings when the periodic inspection task is being executed.

Periodic inspection

- **Inspection Items:** By default, all activated event monitoring policies are supported. If you need to adjust the inspection items, refer to [Cluster Events – Setting Event Policies](#) for configuration. Initially, the system automatically selects all events with monitoring activated for inspection items. After modification, when setting inspection items for the second time, the system defaults to the inspection items selected in the previous round.

Periodic inspection settings ✕

i 1. 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.

2. 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.

3. All enabled event monitoring policies can be set as inspection options by default. You can adjust them on the [Event policies](#) tab.

! Unable to change the settings when the periodic inspection task is being executed.

Periodic inspection

Inspection cycle Daily 00:00

Inspection item

Select all Invert all

- ▶ NODE
- ▶ HDFS
- ▶ YARN
- ▶ HIVE
- ▶ CLUSTER
- ▶ ZOOKEEPER

Save configuration Cancel

Monitoring Metrics

Node Monitoring Metrics

Last updated: 2023-12-21 16:49:33

Node-CPU

| Title | Description | Metric Unit | Description |
|-----------------|-------------|-------------|--|
| CPU utilization | idle | % | Proportion of CPU IDLE Time |
| | irq | % | Interrupt Proportion |
| | nice | % | Proportion of CPU Usage by NICE Priority |
| | steal | % | Proportion of Time Virtual CPU Waits for Actual CPU |
| | softirq | % | Proportion of CPU Soft Interrupts |
| | guest | % | Percentage of Time Utilized by the Virtual Processor |
| | system | % | Kernel Mode CPU Utilization Ratio |
| | user | % | User Mode CPU Utilization Ratio |
| | iowait | % | Proportion of CPU Idle Time Waiting for IO Processes |
| Load | 1m | % | One-Minute Load Average |
| | 5m | % | Five-Minute Load Average |
| | 15m | % | Fifteen-Minute Load Average |
| Number of cores | cpu_count | Connections | CPU Cores |

Node - MEMORY

| Title | Description | Metric Unit | Description |
|--------------------|-------------------|-------------|---|
| Memory utilization | MemTotal | GB | Total Memory Capacity |
| | MemFree | GB | Total Free Memory |
| | MemAvailable | GB | Available Total Memory |
| | Buffers | GB | Total Memory Occupied by BUFFER Cache |
| | Cached | GB | Total Memory Occupied by File Cache |
| | SwapCached | GB | Total Memory Written to Swap Area from Anonymous Pages |
| | SwapFree | GB | Total Available Swap Space |
| | AnonPages | GB | Total Unmapped Memory Volume |
| | SwapTotal | GB | Total Swap Space Volume |
| | Dirty | GB | Total Memory Volume Requiring Disk Write |
| | Writeback | GB | Total Memory Volume Currently Being Written Back to Disk |
| | HardwareCorrupted | GB | Total Memory Volume Unavailable Due to Hardware Malfunction |
| | Shmem | GB | Total Memory Volume Occupied by Shared Memory |
| | MemUsed | GB | Total Memory Volume Utilized |

| | | | |
|---------------------------|-------------------|---|--|
| Proportional Memory Usage | available_percent | % | Percentage of Available Memory to Total Memory |
| | used_percent | % | Percentage of Used Memory to Total Memory |

Node – Disk

| Title | Description | Metric Unit | Description |
|--|-------------------------------------|--------------|--|
| Device Read/Write Speed Rate | Read | MB/s | Device Read/Write Speed, Data Read Volume Per Second |
| | Write | MB/s | Device Read/Write Speed, Data Written Per Second |
| Device IOPS | all | Reads/second | Device IOPS, Ongoing IO Operations on the Current Device |
| IO Operation Duration | Read | ms | Average Waiting Time for Each Device I/O Read Operation |
| | Write | ms | Average Waiting Time for Each Device I/O Write Operation |
| | IO | ms | Average Processing Time per IO Request |
| Device Read/Write Request QPS | Read | Reads/second | Read Operations QPS |
| | Write | Reads/second | Write Operations QPS |
| | Merge-Read | Reads/second | Merged Read Operations QPS |
| | Merge-Write | Reads/second | Merged Write Operations QPS |
| IO Device Utilization Rate | all | % | IO Device Utilization Rate, Disk Activity Intensity |
| Disk Space | Free | GB | Disk Free Storage Space |
| | Available | GB | Available Disk Storage Space (Non-privileged Users) |
| | Total | GB | Total Disk Storage Space |
| Disk space utilization | Used | % | Disk space utilization |
| INODES | Free | Connections | Remaining Quantity of Disk INODES |
| | Total | Connections | Total Quantity of Disk INODES |
| INODES Utilization Rate | Used | % | Disk INODES Utilization Rate |
| Single Disk Space Utilization Rate | Utilized(*) (*Denotes Disk Name) | % | Single Disk Space Utilization Rate |
| Single Disk IO Device Utilization Rate | * (*Denotes Disk Name) | % | Single Disk IO Device Utilization Rate |

Node – File Handle

| Title | Description | Metric Unit | Description |
|--------------------------|------------------------|--------------|-----------------------------------|
| File Handle | allocated | Connections | Number of Allocated File Handles |
| | maximum | Connections | Maximum Number of File Handles |
| System Interruption | intr_total | Reads/second | Number of System Interruptions |
| System Context Switching | context_switches_total | Reads/second | Number of System Context Switches |

| | | | |
|------------------|---------------|-------------|--|
| System Processes | forks_total | Connections | Number of Newly Created System Processes |
| | procs_running | Connections | Number of Running System Processes |
| | procs_blocked | Connections | Number of Blocked System Processes |
| | procs_total | Connections | Total Number of System Processes |
| | thrds_total | Connections | Total Number of System Threads |
| Agent Version | AgentVersion | version | Version of the Agent |

Node-NETWORK

| Title | Description | Metric Unit | Description |
|----------------------------------|-------------------|--------------|---|
| TCP LISTEN Anomalies | ListenDrops | Reads/second | The number of incoming connections (SYN packets) dropped due to any reason. |
| | ListenOverflows | Reads/second | The number of times the Accept queue exceeds its limit after the final step of the three-way handshake is completed. |
| TCPSyncookies | SyncookiesFailed | Reads/second | The number of packets received carrying invalid SYN Cookie information. |
| | SyncookiesRecv | Reads/second | The number of packets received carrying valid SYN Cookie information. |
| | SyncookiesSent | Reads/second | The number of SYN/ACK packets sent using SYN Cookie. |
| Abnormal TCP connection Abort | TCPAbortOnTimeout | Reads/second | The number of times the connection is closed due to exceeding the retransmission limit of various timers (RTO/PTO/keepalive). |
| | TCPAbortOnData | Reads/second | The number of times a socket was closed due to receiving unknown data. |
| | TCPAbortOnClose | Reads/second | The number of times a user-mode program closes a socket while there is still data in the buffer. |
| | TCPAbortOnMemory | Reads/second | The number of times a connection was closed due to memory issues. |
| | TCPAbortOnLinger | Reads/second | The number of times a connection was terminated while in a lingering state after closure. |
| | TCPAbortFailed | Reads/second | The number of unsuccessful attempts to terminate a connection. |
| Establishment of TCP connections | ActiveOpens | count/s | The quantity of actively established TCP connections. |
| | CurrEstab | count/s | The current number of established TCP connections. |
| | PassiveOpens | count/s | The number of passively established TCP connections. |
| | AttemptFails | count/s | The number of failed connection establishments. |
| | EstabResets | count/s | The number of connections that have been RESET. |
| TCP packet | InSegs | count/s | The number of data packets received, inclusive of erroneous packets. |
| | OutSegs | count/s | The quantity of data packets dispatched. |
| | RetransSegs | count/s | The quantity of TCP messages received. |
| | InErrs | count/s | The quantity of retransmitted packets. |
| | OutRsts | count/s | Number of RST packets dispatched. |

| | | | |
|-------------------------|----------------------|--------------|--|
| TCP Retransmission Rate | RetransSegsRate | % | TCP Layer Retransmission Rate. |
| | ResetRate | % | RESET Transmission Frequency |
| | InErrRate | % | Proportion of erroneous packets. |
| TCP TIME-WAIT | TW | count/s | Number of Sockets Terminating the TIME_WAIT State Through Standard Timeout |
| | TWKilled | count/s | Number of Sockets Terminating the TIME_WAIT State Through the tcp_tw_recycle Mechanism |
| | TCPTimeWaitOverflow | count/s | Number of TIME_WAIT Sockets Unable to be Allocated Due to Exceeding Limitations. |
| | TWRecycled | count/s | Number of Sockets Terminating the TIME_WAIT State Through the tcp_tw_reuse Mechanism |
| TCP RTO | TCPTimeouts | Reads/second | Number of Initial Timeouts by the RTO Timer |
| | TCPSpuriousRTOs | Reads/second | Number of False Timeouts Detected Through the F-RTO Mechanism |
| | TCPLossProbes | Reads/second | Number of Instances Where Probe Timeout (PTO) Triggers the Transmission of Tail Loss Probe (TLP) Packets |
| | TCPLossProbeRecovery | Reads/second | Number of Instances Where Lost Packets Were Precisely Rectified by TLP Probe Packets |
| | TCPRenoRecoveryFail | Reads/second | Number of Instances Entering the Recovery Phase Followed by RTO, Where the Counterpart Does Not Support the SACK Option |
| | TCPSackRecoveryFail | Reads/second | Number of Instances Entering the Recovery Phase Followed by RTO, Where the Counterpart Supports the SACK Option |
| | TCPRenoFailures | Reads/second | Number of Instances Entering the TCP_CA_Disorder Phase Followed by RTO Timeout, Where the Counterpart Does Not Support the SACK Option |
| | TCPSackFailures | Reads/second | Number of Instances Entering the TCP_CA_Disorder Phase Followed by RTO Timeout, Where the Counterpart Supports the SACK Option |
| | TCPLossFailures | Reads/second | Number of Instances Entering the TCP_CA_Loss Phase Followed by RTO Timeout |
| TCP RTO Constant | RtoAlgorithm | 1/s | Number of Algorithms for Calculating the Delay of Unanswered Forwarded Objects |
| | RtoMax | 1 | Maximum Value of TCP Delayed Retransmission |
| | RtoMin | 1 | Minimum Value of TCP Delayed Retransmission |
| TCP Retransmission | TCPLostRetransmit | Reads/second | Number of Lost Retransmissions of SKB |
| | TCPFastRetrans | Reads/second | Number of Fast Retransmissions of SKB |
| | TCPForwardRetrans | Reads/second | Number of General Retransmissions of SKB |
| | TCPSlowStartRetrans | Reads/second | Number of Successful Slow Start Retransmissions of SKB |
| | TCPRetransFail | Reads/second | Number of Failed Retransmission Attempts |
| UDP Datagram | OutDatagrams | count/s | Number of UDP Datagrams Sent |

| | | | |
|--|-------------------------|-------------|--|
| | InDatagrams | count/s | Number of UDP Datagrams Received |
| Network Interface Data Transmission Rate | eth0-receive_bytes | MB/s | Volume of Data Received by the Network Interface |
| | eth0-transmit_bytes | MB/s | Network Interface Data Transmission Volume |
| Network Interface Packet Rate | eth0-receive_drop | count/s | Network Interface Discarded Received Data Volume |
| | eth0-receive_errs | count/s | Volume of Abnormal Data Received by the Network Interface |
| | eth0-transmit_drop | count/s | Network Interface Discarded Sent Data Volume |
| | eth0-transmit_errs | count/s | Network Interface Anomalous Sent Data Volume |
| | eth0-transmit_packet_sl | count/s | Network Interface Sent Packet Quantity |
| TCP Socket | TCP_inuse | Connections | Quantity of TCP Sockets in Use (Currently Listening) |
| | TCP_orphan | Connections | Number of TCP Connections Awaiting Closure |
| | TCP_tw | Connections | Number of TCP Sockets Awaiting Termination |
| | TCP_alloc | Connections | Quantity of Allocated TCP Sockets (Established, Granted sk_buff) |
| | sockets_used | Connections | Total Quantity of Utilized Sockets |
| TCP Connection Status | ESTABLISHED | Connections | Quantity of TCP Connections in Established State |
| | SYN-SENT | Connections | Quantity of TCP Connections in SYN-SENT State |
| | SYN-RECV | Connections | Quantity of TCP Connections in SYN-RECV State |
| | FIN-WAIT1 | Connections | Quantity of TCP Connections in FIN-WAIT1 State |
| | FIN-WAIT2 | Connections | Quantity of TCP Connections in FIN-WAIT2 State |
| | TIME-WAIT | Connections | Quantity of TCP Connections in TIME-WAIT State |
| | CLOSE | Connections | Quantity of TCP Connections in CLOSE State |
| | CLOSE-WAIT | Connections | Quantity of TCP Connections in CLOSE-WAIT State |
| | LAST-ACK | Connections | Quantity of TCP Connections in LAST-ACK State |
| | LISTEN | Connections | Quantity of TCP Connections in LISTEN State |

| | | | |
|--|----------|-------------|--|
| | CLOSEING | Connections | Quantity of TCP Connections in CLOSING State |
|--|----------|-------------|--|

Node-Event

| Title | Description | Metric Unit | Description |
|---------------------------------|------------------|-------------|--|
| CPU utilization | used | % | 1 - (Proportion of CPU Idle Time) |
| CPU Load Over Fifteen Minutes | 15m | - | Fifteen-Minute Load Average |
| CPU Load Over One Minute | 1m | - | One-Minute Load Average |
| CPU Load Over Five Minutes | 5m | - | Five-Minute Load Average |
| Disk IOPS | all | - | Device IOPS, Ongoing IO Operations on the Current Device |
| Disk IO Operation Time | IO | - | Average Processing Time per IO Request |
| Disk space utilization | Used | - | Disk space utilization |
| Disk IO Device Utilization Rate | all | - | IO Device Utilization Rate, Disk Activity Intensity |
| Memory utilization | used_percent | - | Percentage of Used Memory to Total Memory |
| Outgoing Network Traffic Rate | *-transmit_bytes | - | Network Interface Data Transmission Volume |
| Incoming Network Traffic Rate | *-receive_bytes | - | Volume of Data Received by the Network Interface |
| TCP Connection Count | CurrEstab | - | The current number of established TCP connections. |

HDFS Monitoring Metrics

Last updated: 2023-12-21 16:49:43

HDFS-Overview

| Title | Description | Metric Unit | Description |
|--|------------------------------|-------------|--|
| Cluster Storage Capacity | CapacityTotal | GB | Total Cluster Storage Capacity |
| | CapacityUsed | GB | Utilized Cluster Storage Capacity |
| | CapacityRemaining | GB | Remaining Cluster Storage Capacity |
| | CapacityUsedNonDFS | GB | Cluster Non-HDFS Usage Capacity |
| Cluster Workload | TotalLoad | 1 | Number of current connections |
| Total Number of Files in the Cluster | FilesTotal | Connections | Total File Count |
| Number of BLOCKS | BlocksTotal | Connections | Total Block Count |
| | PendingReplicationBlocks | Connections | Number of Blocks Awaiting Backup |
| | UnderReplicatedBlocks | Connections | Number of Blocks with Insufficient Replicas |
| | CorruptBlocks | Connections | Number of Bad Blocks |
| | ScheduledReplicationBlocks | Connections | Number of Blocks Scheduled for Backup |
| | PendingDeletionBlocks | Connections | Number of Blocks Awaiting Deletion |
| | ExcessBlocks | Connections | Number of Excess Blocks |
| | PostponedMisreplicatedBlocks | Connections | Number of Exceptional Blocks Deferred for Processing |
| BLOCK Capacity | BlockCapacity | Connections | BLOCK Capacity |
| Cluster Data Nodes | NumLiveDataNodes | Connections | Number of Active Data Nodes |
| | NumDeadDataNodes | Connections | Number of Data Nodes Already Marked as Dead Status |
| | NumDecomLiveDataNodes | Connections | Number of Offline and Live Nodes |
| | NumDecomDeadDataNodes | Connections | Number of Offline and Dead Nodes |
| | NumDecommissioningDataNodes | Connections | Number of Nodes Currently Undergoing Decommissioning |
| | NumStaleDataNodes | Connections | Number of DataNodes Marked as Expired Status |
| Utilization Rate of HDFS Storage Space | CapacityUsedRate | Connections | Utilization Rate of HDFS Cluster Storage Space |

| | | | |
|------------------------|---------------------|---------|--|
| Pertaining to SNAPSHOT | Snapshots | Threads | Number of Snapshots |
| Disk Failure | VolumeFailuresTotal | Threads | Total Number of Volume Failures Across All Datanodes |

HDFS-NameNode

| Title | Description | Metric Unit | Description |
|----------------------------------|----------------------------|----------------------|---|
| Data Traffic | ReceivedBytes | Bytes/s | Data Reception Rate |
| | SentBytes | Bytes/s | Data Transmission Rate |
| QPS | RpcQueueTimeNumOps | 1/s | RPC Invocation Rate |
| Request Processing Latency | RpcQueueTimeAvgTime | ms | Average RPC Delay Time |
| | RpcProcessingTimeAvgTime | ms | Average RPC Request Processing Time |
| Authentication and Authorization | RpcAuthenticationFailures | Per Instance | RPC Authentication Failure Instances |
| | RpcAuthenticationSuccesses | Per Two Instance s | RPC Authentication Success Instances |
| | RpcAuthorizationFailures | Per Three Instance s | RPC Authorization Failure Instances |
| | RpcAuthorizationSuccesses | Per Four Instance s | RPC Authorization Success Instances |
| Number of current connections | NumOpenConnections | Per Single Instance | Current Connection Count |
| RPC Processing Queue Length | CallQueueLength | Per Single Instance | Current RPC Processing Queue Length |
| JVM Memory | MemNonHeapUsedM | MB | Current Size of NonHeapMemory Utilized by the JVM |
| | MemNonHeapCommittedM | MB | Configured Size of NonHeapCommittedM in JVM |
| | MemHeapUsedM | MB | Current Size of HeapMemory Utilized by the JVM |
| | MemHeapCommittedM | MB | Committed Size of JVM HeapMemory |
| | MemHeapMaxM | MB | Configured Size of JVM HeapMemory |
| | MemMaxM | MB | Maximum Memory Size Available for JVM Runtime |
| Heap Memory Utilization Rate | MemHeapUsedRate | % | Percentage of Configured HeapMemory Currently Utilized by the JVM |
| Data Block Reporting Latency | BlockReportAvgTime | Reads/s econd | Average Delay in Processing DataNode Block per Second |
| Number of JVM Threads | ThreadsNew | Connecti ons | Number of Threads in NEW State |

| | | | |
|---|---|--------------|--|
| | ThreadsRunnable | Connections | Number of Threads in RUNNABLE State |
| | ThreadsBlocked | Connections | Number of Threads in BLOCKED State |
| | ThreadsWaiting | Connections | Number of Threads in WAITING State |
| | ThreadsTimedWaiting | Connections | Number of Threads in TIMED WAITING State |
| | ThreadsTerminated | Connections | Number of Threads in Terminated State |
| Number of JVM Logs | LogFatal | Connections | Number of FATAL Level Logs |
| | LogError | Connections | Number of ERROR Level Logs |
| | LogWarn | Connections | Number of WARN Level Logs |
| | LogInfo | Connections | Number of INFO Level Logs |
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| Quantity of Storage Marked as Obsolete | NumStaleStorages | Connections | Total Storage Directory of All Obsolete DataNodes |
| Number of Messages Related to BLOCK Operations Suspended on the Backup NN | PendingDataNodeMessageCount | count/s | Number of DATANODE Requests Queued in the Standby Namenode |
| Missing Block Statistics | NumberOfMissingBlocks | Connections | Number of Missing Data Blocks |
| | NumberOfMissingBlocksWithReplicationFactorOne | Connections | Number of Missing Databases (rf = 1) |
| SNAPSHOT Operations | AllowSnapshotOps | Reads/second | Frequency of AllowSnapshot Operations Performed Per Second |

| | | | |
|--------------------------|---------------------------|--------------|---|
| | DisallowSnapshotOps | Reads/second | Frequency of DisallowSnapshot Operations Executed Per Second |
| | CreateSnapshotOps | Reads/second | Frequency of CreateSnapshot Operations Executed Per Second |
| | DeleteSnapshotOps | Reads/second | Frequency of DeleteSnapshot Operations Executed Per Second |
| | ListSnapshottableDirOps | Reads/second | Frequency of ListSnapshottableDir Operations Executed Per Second |
| | SnapshotDiffReportOps | Reads/second | Frequency of SnapshotDiffReportOps Executed Per Second |
| | RenameSnapshotOps | Reads/second | Frequency of RenameSnapshotOps Executed Per Second |
| File-Related Operations | CreateFileOps | Reads/second | Number of CreateFile Operations Executed Per Second |
| | GetListingOps | Reads/second | Number of GetListing Operations Executed Per Second |
| | TotalFileOps | Reads/second | Number of TotalFileOps Executed Per Second |
| | DeleteFileOps | Reads/second | Number of DeleteFile Operations Executed Per Second |
| | FileInfoOps | Reads/second | Number of FileInfo Operations Executed Per Second |
| | GetAdditionalDatanodeOps | Reads/second | Number of GetAdditionalDatanode Operations Executed Per Second |
| | CreateSymlinkOps | Reads/second | Number of CreateSymlink Operations Executed Per Second |
| | GetLinkTargetOps | Reads/second | Number of GetLinkTarget Operations Executed Per Second |
| | FilesInGetListingOps | Reads/second | Number of FilesInGetListing Operations Executed Per Second |
| File Statistics | FilesDeleted | count | Quantity of Files and Folders Deleted or Renamed |
| | FilesCreated | count | Quantity of Files and Folders Created |
| | FilesAppended | count | Quantity of Appended Files |
| Transactional Operations | TransactionsNumOps | Reads/second | Number of Journal Transaction Operations Processed Per Second |
| | TransactionsBatchedInSync | Reads/second | Number of Journal Transaction Operations Batch Processed Per Second |
| Mirror Operations | GetEditNumOps | Reads/second | Number of GetEditNumOps Executed Per Second |
| | GetImageNumOps | Reads/second | Number of GetImageNumOps Executed Per Second |
| | PutImageNumOps | Reads/second | Number of PutImageNumOps Executed Per Second |
| SYNC Operations | SyncsNumOps | Reads/second | Number of Journal Syncs Operations Processed Per Second |

| | | | |
|--|----------------------------|----------------------------------|--|
| Data Block Operations | BlockReceivedAndDeletedOps | Reads/second | Frequency of BlockReceivedAndDeletedOps Executed Per Second |
| | BlockOpsQueued | Reads/second | Number of DataNode Block Report Operations Processed |
| Cache Reporting | CacheReportNumOps | Reads/second | Number of CacheReport Operations Processed Per Second |
| Data Block Reporting | BlockReportNumQps | Reads/second | Number of Operations for Reporting DataNode Block per Second |
| SYNCS Operation Latency | SyncsAvgTime | ms | Average Latency of Processing Journal Syncs Operations |
| Cache Report Latency | CacheReportAvgTime | ms | Average Delay of Cache Report Actions |
| Latency of Mirror Operations | GetEditAvgTime | ms | Average Latency of Reading Edit File Operations |
| | GetImageAvgTime | ms | Average Latency of Reading Image Files |
| | PutImageAvgTime | ms | Average Latency of Writing to Image Files |
| Transaction Operation Latency | TransactionsAvgTime | ms | Average Latency of Processing Journal Transaction Operations |
| Initiation Time | StartTime | ms | Process Initiation Time |
| Primary and Secondary Status | State | 1 | NN HA Status |
| Primary and Secondary Status | State | 1: Primary 0: Secondary | NameNode Primary and Secondary Status |
| Number of Threads | PeakThreadCount | Connections | Peak Thread Count |
| | ThreadCount | Connections | Number of Threads |
| | DaemonThreadCount | Connections | Background Thread Count |
| Total Transactions Since Last Checkpoint | SinceLastCheckpoint | count | Total Number of Transactions Since the Last Checkpoint |
| Checkpoint Time | LastCheckpoint | time | Time Since the Last Checkpoint |
| Length of the Queue Waiting for File Locks | LockQueueLength | count | LockQueueLength: Length of the Queue Awaiting File Locks |
| Average RPC Time (1) | CompleteAvgTime | ms | Average Latency Time for Complete Requests |
| | CreateAvgTime | ms | Average Latency Time for Create Requests |
| | RenameAvgTime | ms | Average Latency Time for Rename Requests |
| | AddBlockAvgTime | ms | Average Latency Time for AddBlock Requests |
| | GetListingAvgTime | ms | Average Latency Time for GetListing Requests |

| | | | |
|--------------------------|--------------------------------|--------------|---|
| | GetFileInfoAvgTime | ms | Average Latency Time for GetFileInfo Requests |
| | SendHeartbeatAvgTime | ms | Average Latency Time for SendHeartbeat Requests |
| Average RPC Time (2) | RegisterDatanodeAvgTime | ms | Average Latency Time for RegisterDatanode Requests |
| | BlockReportAvgTime | ms | Average Latency Time for BlockReport Requests |
| | DeleteAvgTime | ms | Average Latency Time for Delete Requests |
| | RenewLeaseAvgTime | ms | Average Latency Time for RenewLease Requests |
| | BlockReceivedAndDeletedAvgTime | ms | Average Latency Time for BlockReceivedAndDeleted Requests |
| | FsyncAvgTime | ms | Average Latency Time for Fsync Requests |
| | VersionRequestAvgTime | ms | Average Latency Time for VersionRequest Requests |
| Average Time for RPC (3) | ListEncryptionZonesAvgTime | ms | Average Latency Time for ListEncryptionZones Requests |
| | SetPermissionAvgTime | ms | Average Latency Time for SetPermission Requests |
| | SetTimesAvgTime | ms | Average Latency Time for SetTimes Requests |
| | SetSafeModeAvgTime | ms | Average Latency Time for SetSafeMode Requests |
| | MkdirsAvgTime | ms | Average Latency Time for Mkdirs Requests |
| | GetServerDefaultsAvgTime | ms | Average Latency Time for GetServerDefaults Requests |
| | GetBlockLocationsAvgTime | ms | Average Latency Time for GetBlockLocations Requests |
| RPC Statistics (1) | CompleteNumOps | Reads/second | Number of Calls to Complete per Second |
| | CreateNumOps | Reads/second | Number of Calls to Create per Second |
| | RenameNumOps | Reads/second | Number of Calls to Rename per Second |
| | AddBlockNumOps | Reads/second | Number of Calls to AddBlock per Second |
| | GetListingNumOps | Reads/second | Number of Calls to GetListing per Second |
| | GetFileInfoNumOps | Reads/second | Number of Calls to GetFileInfo per Second |
| | SendHeartbeatNumOps | Reads/second | Number of Calls to SendHeartbeat per Second |
| RPC Statistics (2) | RegisterDatanodeNumOps | Reads/second | Number of Calls to RegisterDatanode per Second |

| | | | |
|--------------------|-------------------------------|--------------|---|
| | BlockReportNumOps | Reads/second | Number of Calls to BlockReport per Second |
| | DeleteNumOps | Reads/second | Number of Calls to Delete per Second |
| | RenewLeaseNumOps | Reads/second | Number of Calls to RenewLease per Second |
| | BlockReceivedAndDeletedNumOps | Reads/second | Number of Calls to BlockReceivedAndDeleted per Second |
| | FsyncNumOps | Reads/second | Number of Calls to Fsync per Second |
| | VersionRequestNumOps | Reads/second | Number of Calls to VersionRequest per Second |
| RPC Statistics (3) | ListEncryptionZonesNumOps | Reads/second | Number of Calls to ListEncryptionZones per Second |
| | SetPermissionNumOps | Reads/second | Number of Calls to SetPermission per Second |
| | SetTimesNumOps | Reads/second | Number of Calls to SetTimes per Second |
| | SetSafeModeNumOps | Reads/second | Number of Calls to SetSafeMode per Second |
| | MkdirsNumOps | Reads/second | Number of Calls to Mkdirs per Second |
| | GetServerDefaultsNumOps | Reads/second | Number of Calls to GetServerDefaults per Second |
| | GetBlockLocationsNumOps | Reads/second | Number of Calls to GetBlockLocations per Second |

HDFS-DataNode

| Title | Description | Metric Unit | Description |
|------------------------------|---------------------------|--------------|---|
| Quantity of Xceivers | XceiverCount | Connections | Quantity of Xceivers |
| Data Read/Write Speed | BytesWrittenMB | Bytes/s | Byte Rate Written to DN |
| | BytesReadMB | Bytes/s | Byte Rate Read from DN |
| | RemoteBytesReadMB | Bytes/s | Remote Client Byte Read Rate |
| | RemoteBytesWrittenMB | Bytes/s | Remote Client Byte Write Rate |
| Number of Client Connections | WritesFromRemoteClient | Connections | Write Operations QPS from Remote Client |
| | WritesFromLocalClient | Connections | Write Operations OPS from Local Client |
| | ReadsFromRemoteClient | Connections | Read Operations QPS from Remote Client |
| | ReadsFromLocalClient | Connections | Read Operations QPS from Local Client |
| Block Verification Failure | BlockVerificationFailures | Reads/second | Number of BLOCK Verification Failures |

| | | | |
|---|------------------------------------|------------------|---|
| Disk Failure | VolumeFailures | Reads/ second | Number of Disk Failures |
| Network Error | DatanodeNetworkErrors | Reads/ second | Network Error Statistics |
| Heartbeat Delay | HeartbeatsAvgTime | ms | Average Time of Heartbeat Interface |
| Heartbeat QPS | HeartbeatsNumOps | Reads/ second | Heartbeat Interface QPS |
| Packet Transmission Operation RT | SendDataPacketTransferNanosAvgTime | ms | Average Time for Sending Data Packets |
| Data Block Operations | ReadBlockOpNumOps | Reads/ second | Reading Block OPS from DataNode |
| | WriteBlockOpNumOps | Reads/ second | Writing Block OPS to DataNode |
| | BlockChecksumOpNumOps | Reads/ second | Checksum Operation OPS performed by DataNode |
| | CopyBlockOpNumOps | Reads/ second | OPS for Copying Block Operation |
| | ReplaceBlockOpNumOps | Reads/ second | OPS for Replace Block Operation |
| | BlockReportsNumOps | Reads/ second | OPS for Block Reporting Action |
| | IncrementalBlockReportsNumOps | Reads/ second | OPS for Incremental Block Reporting |
| | CacheReportsNumOps | Reads/ second | OPS for Cache Reporting |
| | PacketAckRoundTripTimeNanosNumOps | Reads/ second | ACK ROUND TRIP Processing Rate Per Second |
| FSYNC Operation | FsyncNanosNumOps | Reads/ second | FSYNC Operations Processed Per Second |
| FLUSH Operation | FlushNanosNumOps | Reads/ second | Flush Operation Processing Rate Per Second |
| Data Block Operation Latency Statistics | ReadBlockOpAvgTime | ms | Average Time for Block Read Operation |
| | WriteBlockOpAvgTime | ms | Average Time for Block Write Operation |
| | BlockChecksumOpAvgTime | ms | Average Time for Block Verification Operation |
| | CopyBlockOpAvgTime | ms | Average Time for Block Copy Operation |
| | ReplaceBlockOpAvgTime | ms | Average Time for Replace Block Operation |
| | BlockReportsAvgTime | ms | Average Time for Block Reporting |
| | IncrementalBlockReportsAvgTime | ms | Average Time for Incremental Block Reporting |
| | CacheReportsAvgTime | ms | Average Time for Cache Reporting |
| | PacketAckRoundTripTimeNanosAvgTime | ms | Average Time for Processing ACK ROUND TRIP |
| FLUSH Latency | FlushNanosAvgTime | ms | Average Time for Flush Operations |

| | | | |
|------------------------------|---|---------------|---|
| FSYNC Latency | FsyncNanosAvgTime | ms | Average Time for Fsync Operations |
| RAMDISK Blocks | RamDiskBlocksWrite | Blocks/second | Total Number of Blocks Written to Memory |
| | RamDiskBlocksWriteFallback | Blocks/second | Total Number of Blocks Written to Memory but Unsuccessful (Failure Shifted to Disk) |
| | RamDiskBlocksDeletedBeforeLazyPersisted | Blocks/second | Total Number of Blocks Deleted by the Application Prior to Being Saved to Disk |
| | RamDiskBlocksReadHits | Blocks/second | Total Number of Times Blocks in Memory Have Been Read |
| | RamDiskBlocksEvicted | Blocks/second | Total Number of Blocks Cleared in Memory |
| | RamDiskBlocksEvictedWithoutRead | Blocks/second | Total Number of Memory Blocks Retrieved from Memory |
| | RamDiskBlocksLazyPersisted | Blocks/second | Total Number of Writes to Disk by the Lazy Writer |
| | RamDiskBytesLazyPersisted | Bytes/s | Total Number of Bytes Written to Disk by the Lazy Writer |
| RAMDISK Write Speed | RamDiskBytesWrite | Bytes/s | Total Number of Bytes Written into Memory |
| JVM Memory | MemNonHeapUsedM | MB | Current Size of NonHeapMemory Utilized by the JVM |
| | MemNonHeapCommittedM | MB | Configured Size of NonHeapCommittedM in JVM |
| | MemHeapUsedM | MB | Current Size of HeapMemory Utilized by the JVM |
| | MemHeapCommittedM | MB | Committed Size of JVM HeapMemory |
| | MemHeapMaxM | MB | Configured Size of JVM HeapMemory |
| | MemMaxM | MB | Maximum Memory Size Available for JVM Runtime |
| Heap Memory Utilization Rate | MemHeapUsedRate | % | Percentage of Configured HeapMemory Currently Utilized by the JVM |
| Number of JVM Threads | ThreadsNew | Connections | Number of Threads in NEW State |
| | ThreadsRunnable | Connections | Number of Threads in RUNNABLE State |
| | ThreadsBlocked | Connections | Number of Threads in BLOCKED State |
| | ThreadsWaiting | Connections | Number of Threads in WAITING State |
| | ThreadsTimedWaiting | Connections | Number of Threads in TIMED WAITING State |
| | ThreadsTerminated | Connections | Number of Threads in Terminated State |
| Number of JVM Logs | LogFatal | Connections | Number of Fatal Logs |

| | | | |
|----------------------------------|----------------------------|--------------|---|
| | LogError | Connections | Number of Error Logs |
| | LogWarn | Connections | Number of Warn Logs |
| | LogInfo | Connections | Number of Info Logs |
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| Data Traffic | ReceivedBytes | Bytes/s | Data Reception Rate |
| | SentBytes | Bytes/s | Data Transmission Rate |
| QPS | RpcQueueTimeNumOps | Reads/second | RPC Invocation Rate |
| Request Processing Latency | RpcQueueTimeAvgTime | ms | Average RPC Delay Time |
| | RpcProcessingTimeAvgTime | Reads/second | Average RPC Request Processing Time |
| Authentication and Authorization | RpcAuthenticationFailures | Reads/second | RPC Authentication Failure Instances |
| | RpcAuthenticationSuccesses | Reads/second | RPC Authentication Success Instances |
| | RpcAuthorizationFailures | Reads/second | RPC Authorization Failure Instances |
| | RpcAuthorizationSuccesses | Reads/second | RPC Authorization Success Instances |
| Number of current connections | NumOpenConnections | Connections | Current Link Quantity |
| RPC Processing Queue Length | CallQueueLength | 1 | Current RPC Processing Queue Length |
| CPU Time | CurrentThreadSystemTime | ms | System time |
| | CurrentThreadUserTime | ms | User Time |
| Initiation Time | StartTime | s | Process Initiation Time |

| | | | |
|--------------------------------|--|--------------|--------------------------------|
| Number of Threads | PeckThreadCount | Connections | Peak Thread Count |
| | DaemonThreadCount | Connections | Background Thread Count |
| Read-Write Latency | write | ms | Write Operation Duration |
| | read | ms | Read Operation Duration |
| Packet Transfer Operations QPS | DataPacketOps | Reads/second | Packet Transfer Operations QPS |
| Block Quantity | Pertaining to disk information, for instance: /data/qcloud/data/hdfs | Connections | Block Quantity |
| Utilized Disk Capacity | Pertaining to disk information, for instance: /data/qcloud/data/hdfs | GB | Utilized Disk Capacity |
| Available Disk Capacity | Pertaining to disk information, for instance: /data/qcloud/data/hdfs | GB | Available Disk Capacity |
| Reserved Disk Capacity | Pertaining to disk information, for instance: /data/qcloud/data/hdfs | GB | Reserved Disk Capacity |

HDFS-JournalNode

| Title | Description | Metric Unit | Description |
|------------------------------|----------------------|-------------|---|
| JVM Memory | MemNonHeapUsedM | MB | Current Size of NonHeapMemory Utilized by the JVM |
| | MemNonHeapCommittedM | MB | Configured Size of NonHeapCommittedM in JVM |
| | MemHeapUsedM | MB | Current Size of HeapMemory Utilized by the JVM |
| | MemHeapCommittedM | MB | Committed Size of JVM HeapMemory |
| | MemHeapMaxM | MB | Configured Size of JVM HeapMemory |
| | MemMaxM | MB | Maximum Memory Size Available for JVM Runtime |
| Heap Memory Utilization Rate | MemHeapUsedRate | % | Percentage of Configured HeapMemory Currently Utilized by the JVM |
| Number of JVM Threads | ThreadsNew | Connections | Number of Threads in NEW State |
| | ThreadsRunnable | Connections | Number of Threads in RUNNABLE State |
| | ThreadsBlocked | Connections | Number of Threads in BLOCKED State |
| | ThreadsWaiting | Connections | Number of Threads in WAITING State |
| | ThreadsTimedWaiting | Connections | Number of Threads in TIMED WAITING State |
| | ThreadsTerminated | Connections | Number of Threads in Terminated State |
| Number of JVM Logs | LogFatal | Connections | Number of FATAL Level Logs |
| | LogError | Connections | Number of ERROR Level Logs |

| | | | |
|----------------------------------|----------------------------|--------------|---|
| | LogWarn | Connections | Number of WARN Level Logs |
| | LogInfo | Connections | Number of INFO Level Logs |
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| Data Traffic | ReceivedBytes | Bytes/s | Data Reception Rate |
| | SentBytes | Bytes/s | Data Transmission Rate |
| Request Processing Latency | RpcQueueTimeAvgTime | ms | Average RPC Delay Time |
| Authentication and Authorization | RpcAuthenticationFailures | Reads/second | RPC Authentication Failure Instances |
| | RpcAuthenticationSuccesses | Reads/second | RPC Authentication Success Instances |
| | RpcAuthorizationFailures | Reads/second | RPC Authorization Failure Instances |
| | RpcAuthorizationSuccesses | Reads/second | RPC Authorization Success Instances |
| Number of current connections | NumOpenConnections | Connections | Current Link Quantity |
| RPC Processing Queue Length | CallQueueLength | 1 | Current RPC Processing Queue Length |
| CPU Time | CurrentThreadSystemTime | ms | System time |
| | CurrentThreadUserTime | ms | User Time |
| Initiation Time | StartTime | s | Process Initiation Time |
| Number of Threads | PeckThreadCount | Connections | Peak Thread Count |
| | DaemonThreadCount | Connections | Background Thread Count |

HDFS-ZKFC

| Title | Description | Metric Unit | Description |
|-------|-------------|-------------|-------------|
| | | | |

| | | | |
|------------------------------|------|---------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |

HDFS-Router

| Title | Description | Metric Unit | Description |
|--|--|-------------|--|
| Alter Table Request Time | HIVE.HMS.API_ALTER_TABLE | ms | Average Time for Alter Table Request |
| Alter Table with Environmental Context Request Time | HIVE.HMS.API_ALTER_TABLE_WITH_ENV_CONTEXT | ms | Average Time for Alter Table Request with Environmental Context |
| Create Table Request Time | HIVE.HMS.API_CREATE_TABLE | ms | Average Time for Create Table Request |
| Create Table with Environmental Context Request Time | HIVE.HMS.API_CREATE_TABLE_WITH_ENV_CONTEXT | ms | Average Time for Create Table Request with Environmental Context |
| Drop Table Request Time | HIVE.HMS.API_DROP_TABLE | ms | Average Drop Table Request Time |
| Drop Table with Environmental Context Request Time | HIVE.HMS.API_DROP_TABLE_WITH_ENV_CONTEXT | ms | Average Drop Table with Environmental Context Request Time |
| Get Table Request Time | HIVE.HMS.API_GET_TABLE | ms | Average Get Table Request Time |
| Get Tables Request Time | HIVE.HMS.API_GET_TABLES | ms | Average Get Tables Request Time |
| Get Multi Table Request Time | HIVE.HMS.API_GET_MULTI_TABLE | ms | Average Get Multiple Tables Request Time |
| Get Table Request Query Time | HIVE.HMS.API_GET_TABLE_REQ | ms | Average Get Table Request Time |
| Get Database Request Time | HIVE.HMS.API_GET_DATABASE | ms | Average Get Database Request Time |
| Get Databases Request Time | HIVE.HMS.API_GET_DATABASES | ms | Average Get Databases Request Time |
| Get All Databases Request Time | HIVE.HMS.API_GET_ALL_DATABASES | ms | Average Get All Databases Request Time |
| Get All Functions Request | HIVE.HMS.API_GET_ALL_FUNCTIONS | ms | Average Get All Functions |

| Time | | | Request Time |
|---|--|-------------|---|
| Current Active Create Table Request Count | HIVE.HMS.ACTIVE_CALLS_API_CREATE_TABLE | Connections | Current Active Create Table Request Count |
| Current Active Drop Table Request Count | HIVE.HMS.ACTIVE_CALLS_API_DROP_TABLE | Connections | Current Active Drop Table Request Count |
| Current Active Alter Table Request Count | HIVE.HMS.ACTIVE_CALLS_API_ALTER_TABLE | Connections | Current Active Alter Table Request Count |

YARN Monitoring Metrics

Last updated: 2023-12-21 16:49:55

YARN – Overview

| Title | Description | Metric Unit | Description |
|---------------------------|------------------------------|--|--|
| Number of nodes. | NumActiveNMs | Connections | Current number of live NodeManagers. |
| | NumDecommissionedNMs | Connections | Current number of Decommissioned NodeManagers. |
| | NumLostNMs | Connections | Current number of Lost NodeManagers. |
| | NumUnhealthyNMs | Connections | Current number of Unhealthy NodeManagers. |
| CPU Cores | AllocatedVCores | Core | Current number of VCore allocations in the queue. |
| | ReservedVCores | Core | Current number of VCore reservations in the queue. |
| | AvailableVCores | Core | Current number of available VCore in the queue. |
| | PendingVCores | Core | Current number of pending VCore requests in the queue. |
| Total applications | AppsSubmitted | Connections | Historical number of jobs submitted in the current queue. |
| | AppsRunning | Connections | Current number of jobs running in the queue. |
| | AppsPending | Connections | Current number of pending jobs in the queue. |
| | AppsCompleted | Connections | Number of jobs completed in the current queue. |
| | AppsKilled | Connections | Number of jobs terminated in the current queue. |
| | AppsFailed | Connections | Number of failed jobs in the current queue. |
| | ActiveApplications | Connections | Number of active jobs in the current queue. |
| | running_0 | Connections | Number of jobs in the current queue with a runtime of less than 60 minutes. |
| | running_60 | Connections | Number of jobs in the current queue with a runtime between 60 and 300 minutes. |
| | running_300 | Connections | Number of jobs in the current queue with a runtime between 300 and 1440 minutes. |
| running_1440 | Connections | Number of jobs in the current queue with a runtime exceeding 1440 minutes. | |
| Memory Size | AllocatedMB | MB | Allocated memory size of the current queue. |
| | AvailableMB | MB | Available memory size of the current queue. |
| | PendingMB | MB | Pending memory size in the resource requests of the current queue. |
| | ReservedMB | MB | Reserved memory size in the current queue. |
| Number of Containers | AllocatedContainers | Connections | Number of containers allocated to the current queue. |
| | PendingContainers | Connections | Number of pending containers in the resource requests of the current queue. |
| | ReservedContainers | Connections | Number of reserved containers in the current queue. |
| Total Number of Container | AggregateContainersAllocated | Connections | Total Number of Containers Allocated to the Current Queue |

| | | | |
|---|-----------------------------|-------------|---|
| Allocations and Releases | AggregateContainersReleased | Connections | Total number of containers released from the current queue. |
| Users | ActiveUsers | Connections | Number of Active Users in the Current Queue |
| Memory | allocatedMB | MB | Allocated Memory Resources within the Cluster |
| | availableMB | MB | Available Memory Resources within the Cluster |
| | reservedMB | MB | Reserved Memory Resources within the Cluster |
| | totalMB | MB | Total Memory Resources within the Cluster |
| Applications | completed | Connections | Number of Jobs Completed within the Sampling Period in the Cluster |
| | failed | Connections | Number of Jobs Failed within the Sampling Period in the Cluster |
| | killed | Connections | Number of Jobs Terminated within the Sampling Period in the Cluster |
| | pending | Connections | Number of Jobs Awaiting Execution within the Sampling Period in the Cluster |
| | running | Connections | Number of Jobs in Progress within the Sampling Period in the Cluster |
| | submitted | Connections | Number of Jobs Submitted within the Sampling Period in the Cluster |
| Containers | containersAllocated | Connections | Number of Containers Allocated in the Cluster |
| | containersPending | Connections | Number of Containers Requested in the Cluster |
| | containersReserved | Connections | Number of Containers Reserved in the Cluster |
| Memory utilization | usageRatio | % | Current Utilization Rate of Memory Resources in the Cluster |
| Cores | allocatedVirtualCores | Connections | Allocated CPU Resources in the Cluster |
| | availableVirtualCores | Connections | Available CPU Resources in the Cluster |
| | reservedVirtualCores | Connections | Reserved CPU Resources in the Cluster |
| | totalVirtualCores | Connections | Total CPU Resources in the Cluster |
| CPU utilization | usageRatio | % | Current Utilization Rate of CPU Resources in the Cluster |
| Number of Application Masters Initiated | AMLlaunchDelayNumOps | Connections | Number of Application Masters Initiated |
| Average Time for Resource Manager to Initiate Application Master | AMLlaunchDelayAvgTime | ms | Average Time for Resource Manager to Initiate Application Master |
| Total Number of Registered Application Masters | AMRegisterDelayNumOps | Connections | Total Number of Registered Application Masters |
| Average Time for Application Master to Register with Resource Manager | AMRegisterDelayAvgTime | ms | Average Time for Application Master to Register with Resource Manager |
| Queue CPU | YARN.RM.QUEUE.VCO | Connections | Current Queue Allocation CPU Utilization Rate |

| | | | |
|-------------------------------|-------------------------|-------------|--|
| Utilization Rate | RES.RATIO | | |
| Queue Memory Utilization Rate | YARN.RM.QUEUE.MEM.RATIO | Connections | Current Queue Allocation Memory Utilization Rate |

YARN-ResourceManager

| Title | Description | Metric Unit | Description |
|-----------------------------------|----------------------------|-------------|---|
| RPC Authentication Authorizations | RpcAuthenticationFailures | Connections | Number of RPC Authentication Failures |
| | RpcAuthenticationSuccesses | Connections | Number of Successful RPC Authentications |
| | RpcAuthorizationFailures | Connections | Number of RPC Authorization Failures |
| | RpcAuthorizationSuccesses | Connections | Number of Successful RPC Authorizations |
| RPC Data Received and Sent | ReceivedBytes | bytes/s | Volume of Data Received via RPC |
| | SentBytes | bytes/s | Volume of Data Sent via RPC |
| RPC Connection Count | NumOpenConnections | Connections | Current Number of Open Connections |
| RPC Request Frequency | RpcProcessingTimeNumOps | Threads | RPC Request Frequency |
| | RpcQueueTimeNumOps | Threads | RPC Request Frequency |
| RPC Queue Length | CallQueueLength | Connections | Current RPC Queue Length |
| RPC Average Processing Time | RpcProcessingTimeAvgTime | s | Average RPC Request Processing Time |
| | RpcQueueTimeAvgTime | s | Average Time of RPC in Queue |
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| Heap Memory Utilization Rate | MemHeapUsedRate | % | Percentage of Configured HeapMemory Currently Utilized by the JVM |
| Number of JVM Threads | ThreadsNew | Connections | Number of Threads in NEW State |
| | ThreadsRunnable | Connections | Number of Threads in RUNNABLE State |
| | ThreadsBlocked | Connections | Number of Threads in BLOCKED State |
| | ThreadsWaiting | Connections | Number of Threads in WAITING State |

| | | | |
|----------------------------|-------------------------|--------------------|---|
| | ThreadsTimedWaiting | Connections | Number of Threads in TIMED WAITING State |
| | ThreadsTerminated | Connections | Number of Threads in Terminated State |
| Number of JVM Logs | LogFatal | Connections | Number of Fatal Logs |
| | LogError | Connections | Number of Error Logs |
| | LogWarn | Connections | Number of Warn Logs |
| | LogInfo | Connections | Number of Info Logs |
| JVM Memory | MemNonHeapUsedM | MB | Size of Non-Heap Memory Used by the Process |
| | MemNonHeapCommittedM | MB | Size of Non-Heap Memory Committed by the Process |
| | MemHeapUsedM | MB | Heap Memory Size Utilized by the Process |
| | MemHeapCommittedM | MB | Size of Heap Memory Committed by the Process |
| | MemHeapMaxM | MB | Maximum Heap Memory Size Utilized by the Process |
| | MemMaxM | MB | Maximum Memory Size Utilized by the Process |
| CPU utilization | ProcessCpuLoad | % | CPU utilization |
| Accumulated CPU Usage Time | ProcessCpuTime | ms | Accumulated CPU Usage Time |
| File Descriptor Count | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| | OpenFileDescriptorCount | Connections | Number of Open File Descriptors |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| Number of Working Threads | DaemonThreadCount | Connections | Number of Daemon Threads in the Process |
| | ThreadCount | Connections | Number of Threads in the Process |
| Node status | haState | 1:Active,0:Standby | ResourceManager Primary-Secondary Status |
| Active/Standby switch | switchOccurred | - | Switching between Primary and Secondary ResourceManager |

YARN-JobHistoryServer

| Title | Description | Metric Unit | Description |
|-----------------------|---------------------|-------------|--|
| Number of JVM Threads | ThreadsNew | Connections | Number of Threads in NEW State |
| | ThreadsRunnable | Connections | Number of Threads in RUNNABLE State |
| | ThreadsBlocked | Connections | Number of Threads in BLOCKED State |
| | ThreadsWaiting | Connections | Number of Threads in WAITING State |
| | ThreadsTimedWaiting | Connections | Number of Threads in TIMED WAITING State |
| | ThreadsTerminated | Connections | Number of Threads in Terminated State |
| Number of JVM Logs | LogFatal | Connections | Number of FATAL Level Logs |

| | | | |
|------------------------------|-------------------------|-------------|---|
| | LogError | Connections | Number of ERROR Level Logs |
| | LogWarn | Connections | Number of WARN Level Logs |
| | LogInfo | Connections | Number of INFO Level Logs |
| JVM Memory | MemNonHeapUsedM | MB | Size of Non-Heap Memory Used by the Process |
| | MemNonHeapCommittedM | MB | Size of Non-Heap Memory Committed by the Process |
| | MemHeapUsedM | MB | Heap Memory Size Utilized by the Process |
| | MemHeapCommittedM | MB | Size of Heap Memory Committed by the Process |
| | MemHeapMaxM | MB | Maximum Heap Memory Size Utilized by the Process |
| | MemMaxM | MB | Maximum Memory Size Utilized by the Process |
| Heap Memory Utilization Rate | MemHeapUsedRate | % | Percentage of Configured HeapMemory Currently Utilized by the JVM |
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| CPU utilization | ProcessCpuLoad | % | CPU utilization |
| Accumulated CPU Usage Time | ProcessCpuTime | ms | Accumulated CPU Usage Time |
| File Descriptor Count | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| | OpenFileDescriptorCount | Connections | Number of Open File Descriptors |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| Number of Working Threads | DaemonThreadCount | Connections | Number of Daemon Threads in the Process |
| | ThreadCount | Connections | Number of Threads in the Process |

YARN-NodeManager

| Title | Description | Metric Unit | Description |
|------------------------------|----------------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| Heap Memory Utilization Rate | MemHeapUsedRate | % | Percentage of Configured HeapMemory Currently Utilized by the JVM |
| Number of JVM Threads | ThreadsNew | Connections | Number of Threads in NEW State |
| | ThreadsRunnable | Connections | Number of Threads in RUNNABLE State |
| | ThreadsBlocked | Connections | Number of Threads in BLOCKED State |
| | ThreadsWaiting | Connections | Number of Threads in WAITING State |
| | ThreadsTimedWaiting | Connections | Number of Threads in TIMED WAITING State |
| | ThreadsTerminated | Connections | Current Number of Threads in TERMINATED State |
| Number of JVM Logs | LogFatal | Connections | Number of FATAL Level Logs |
| | LogError | Connections | Number of ERROR Level Logs |
| | LogWarn | Connections | Number of WARN Level Logs |
| | LogInfo | Connections | Number of INFO Level Logs |
| JVM Memory | MemNonHeapUsedM | MB | Size of Non-Heap Memory Used by the Process |
| | MemNonHeapCommittedM | MB | Size of Non-Heap Memory Committed by the Process |
| | MemHeapUsedM | MB | Heap Memory Size Utilized by the Process |
| | MemHeapCommittedM | MB | Size of Heap Memory Committed by the Process |

| | | | |
|---|--------------------------------|-------------|---|
| | MemHeapMaxM | MB | Maximum Heap Memory Size Utilized by the Process |
| | MemMaxM | MB | Maximum Memory Size Utilized by the Process |
| Total Number of Containers | ContainersLaunched | Connections | Number of Containers Launched |
| | ContainersCompleted | Connections | Number of Completed Containers |
| | ContainersFailed | Connections | Number of Failed Containers |
| | ContainersKilled | Connections | Number of Containers Killed |
| | ContainersIniting | Connections | Number of Containers in Initialization |
| | ContainersRunning | Connections | Number of Containers Currently Running |
| | AllocatedContainers | Connections | Number of Containers Allocated by NodeManager |
| Average Time Consumed in Container Initialization | ContainerLaunchDurationAvgTime | ms | Average Time Consumed in Container Initialization |
| Number of Container Initialization Operations | ContainerLaunchDurationNumOps | Connections | Number of Container Initialization Operations |
| CPU Cores | AvailableVCores | Core | Number of VCore Available to NodeManager |
| | AllocatedVCores | Core | Number of VCores Allocated by NodeManager |
| Memory Size | AllocatedGB | GB | Memory Size Allocated by NodeManager |
| | AvailableGB | GB | Memory Size Available to NodeManager |
| CPU utilization | ProcessCpuLoad | % | CPU utilization |
| Accumulated CPU Usage Time | ProcessCpuTime | ms | Accumulated CPU Usage Time |
| File Descriptor Count | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| | OpenFileDescriptorCount | Connections | Number of Open File Descriptors |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| Number of Working Threads | DaemonThreadCount | Connections | Number of Daemon Threads in the Process |
| | ThreadCount | Connections | Number of Threads in the Process |

YARN-Timeline

| Title | Description | Metric Unit | Description |
|---------------|-------------|-------------|--|
| Number of JVM | GcCount | count | Number of JVM Garbage Collection Instances |

| | | | |
|---|----------------------|-------|--|
| Garbage Collection Instances | | | |
| JVM Garbage Collection Duration | GcTimeMillis | ms | JVM Garbage Collection Duration |
| JVM Memory | MemNonHeapUsedM | MB | Size of Non-Heap Memory Used by the Process |
| | MemNonHeapCommittedM | MB | Size of Non-Heap Memory Committed by the Process |
| | MemNonHeapMaxM | MB | Heap Memory Size Utilized by the Process |
| | MemHeapUsedM | MB | Size of Heap Memory Committed by the Process |
| | MemHeapCommittedM | MB | Maximum Heap Memory Size Utilized by the Process |
| | MemHeapMaxM | MB | Size of Non-Heap Memory Used by the Process |
| Retrieve Domain Operations Count | Ops | count | Retrieve Domain Operations Count |
| Retrieve Bulk Domains Operations Count | Ops | count | Retrieve Bulk Domains Operations Count |
| Average Time for Bulk Retrieval of Domains | Time | ms | Average Time for Bulk Retrieval of Domains |
| Average Time for Retrieving Domain | Time | ms | Average Time for Retrieving Domain |
| Bulk Retrieval of Entities Operations Count | Ops | count | Bulk Retrieval of Entities Operations Count |
| Average Time for Bulk Retrieval of Entities | Time | ms | Average Time for Bulk Retrieval of Entities |
| Retrieve Entity Operations Count | Ops | count | Retrieve Entity Operations Count |
| Average Time for Retrieving Entity | Time | ms | Average Time for Retrieving Entity |
| Bulk Retrieval of Events Operations Count | Ops | count | Bulk Retrieval of Events Operations Count |
| Average Time for Bulk Retrieval of Events | Time | ms | Average Time for Bulk Retrieval of Events |
| Bulk Update of Entities Operations Count | Ops | count | Bulk Update of Entities Operations Count |
| Average Time for Bulk Update of Entities | Time | ms | Average Time for Bulk Update of Entities |
| Domain Update Operations Count | Ops | count | Domain Update Operations Count |
| Average Time for Domain Update | Time | ms | Average Time for Domain Update |

ZooKeeper Monitoring Metrics

Last updated: 2023-12-21 16:50:04

Zookeeper

| | | |
|--|--|--|
| | | |
|--|--|--|

| Title | Description | Metric Unit | Description |
|------------------------------|-------------------------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| JVM Memory | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| | MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| CPU utilization | ProcessCpuLoad | % | CPU utilization |
| File Descriptor Count | zk_max_file_descriptor_count | Connections | Maximum File Descriptor Quantity |
| | zk_open_file_descriptor_count | Connections | Number of Open File Descriptors |
| Accumulated CPU Usage Time | ProcessCpuTime | ms | Accumulated CPU Usage Time |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| Number of Working Threads | DaemonThreadCount | Connections | Daemon Thread Count |
| | ThreadCount | Connections | Total Thread Count |
| Number of connections | zk_num_alive_connections | Connections | Number of current connections |
| Latency | zk_avg_latency | ms | Average Latency of zk Processing |
| | zk_max_latency | ms | Maximum Delay in zk Processing |
| | zk_min_latency | ms | Minimum Delay in zk Processing |
| Number of ZONDEs | zk_watch_count | Connections | Number of zk Watches |

| | | | |
|---------------------------------|--------------------------|--|----------------------------------|
| | | ns | |
| | zk_znode_count | Connections | Quantity of zk's znodes |
| | zk_ephemerals_count | Connections | Number of Temporary Nodes in zk |
| Data Size | zk_approximate_data_size | Byte | Data Storage Volume of zk |
| Node status | zk_server_state | 1: Primary 0: Backup 2: Standalone | zk Node Type |
| Received and Sent Packet Volume | zk_packets_received | count/s | zk Data Packet Reception Rate |
| | zk_packets_sent | count/s | zk Data Packet Transmission Rate |
| Number of Queued Requests | zk_outstanding_requests | Connections | Number of Queued Requests |

HBase Monitoring Metrics

Last updated: 2023-12-21 16:50:31

HBASE – Overview

| Title | Description | Metric Unit | Description |
|---|-------------------------|-------------|---|
| Number of Clusters in RIT Region | ritCount | Connections | Number of Regions in Transition |
| | ritCountOverThreshold | Connections | Number of Regions where the Region in Transition Time Exceeds the Threshold |
| Cluster RIT Duration | ritOldestAge | ms | Oldest Age of Region in Transition |
| Average Number of Regions per RS | averageLoad | Connections | Average Number of Regions per RegionServer |
| Number of RS in the Cluster | numRegionServers | Connections | Number of currently active RegionServers |
| | numDeadRegionServers | Connections | Number of Currently Dead RegionServers |
| Data Read/Write Volume of HMaster | receivedBytes | bytes/s | Cluster Data Reception Volume |
| | sentBytes | bytes/s | Cluster Data Transmission Volume |
| Total Request Volume of the Cluster Interface | clusterRequests | count/s | Total Number of Cluster Requests |
| Cluster Assignment Manager Operations | Assign_num_ops | Threads | Number of Assign Region Instances |
| | BulkAssign_num_ops | Threads | Number of Bulk Assign Region Instances |
| Cluster Load Balancing Instances | BalancerCluster_num_ops | Threads | Cluster Load Balancing Instances |

HBASE-HMaster

| Title | Description | Metric Unit | Description |
|------------------------------|-------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| Number of JVM Logs | LogFatal | Connections | Number of Fatal Logs |
| | LogError | Connections | Number of Error Logs |

| | | | |
|------------------------------|--------------------------------|-------------|---|
| | LogWarn | Connections | Number of Warn Logs |
| | LogInfo | Connections | Number of Info Logs |
| JVM Memory | MemNonHeapUsedM | MB | Size of Non-Heap Memory Used by the Process |
| | MemNonHeapCommittedM | MB | Size of Non-Heap Memory Committed by the Process |
| | MemHeapUsedM | MB | Heap Memory Size Utilized by the Process |
| | MemHeapCommittedM | MB | Committed Heap Memory Size of the Process |
| | MemHeapMaxM | MB | Maximum Heap Memory Size Utilized by the Process |
| | MemMaxM | MB | Maximum Memory Size Utilized by the Process |
| Heap Memory Utilization Rate | MemHeapUsedRate | % | Percentage of Configured HeapMemory Currently Utilized by the JVM |
| Number of JVM Threads | ThreadsNew | Connections | Number of Threads in NEW State |
| | ThreadsRunnable | Connections | Number of Threads in RUNNABLE State |
| | ThreadsBlocked | Connections | Number of Threads in BLOCKED State |
| | ThreadsWaiting | Connections | Number of Threads in WAITING State |
| | ThreadsTimedWaiting | Connections | Number of Threads in TIMED WAITING State |
| | ThreadsTerminated | Connections | Current Number of Threads in TERMINATED State |
| RPC Connection Count | numOpenConnections | Connections | RPC Connection Count |
| RPC Exception Count | FailedSanityCheckException | Threads | FailedSanityCheckException Exception Instances |
| | NotServingRegionException | Threads | NotServingRegionException Exception Instances |
| | OutOfOrderScannerNextException | Threads | OutOfOrderScannerNextException Exception Instances |
| | RegionMovedException | Threads | RegionMovedException Exception Instances |
| | RegionTooBusyException | Threads | RegionTooBusyException Exception Instances |
| | UnknownScannerException | Threads | UnknownScannerException Exception Instances |
| RPC Queue Request Count | numCallsInPriorityQueue | Connections | General Queue RPC Request Count |
| | numCallsInReplicationQueue | Connections | Replication Queue RPC Request Count |
| Process Initiation Time | masterActiveTime | s | Master Process Active Duration |
| | masterStartTime | s | Master Process Initiation Time |

HBASE-RegionServer

| | | | |
|--|--|--|--|
| | | | |
|--|--|--|--|

| Title | Description | Metric Unit | Description |
|------------------------------|----------------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| Number of JVM Logs | LogFatal | Connections | Number of Fatal Logs |
| | LogError | Connections | Number of Error Logs |
| | LogWarn | Connections | Number of Warn Logs |
| | LogInfo | Connections | Number of Info Logs |
| JVM Memory | MemNonHeapUsedM | MB | Size of Non-Heap Memory Used by the Process |
| | MemNonHeapCommittedM | MB | Size of Non-Heap Memory Committed by the Process |
| | MemHeapUsedM | MB | Heap Memory Size Utilized by the Process |
| | MemHeapCommittedM | MB | Committed Heap Memory Size of the Process |
| | MemHeapMaxM | MB | Maximum Heap Memory Size Utilized by the Process |
| | MemMaxM | MB | Maximum Memory Size Utilized by the Process |
| Heap Memory Utilization Rate | MemHeapUsedRate | % | Percentage of Configured HeapMemory Currently Utilized by the JVM |
| Number of JVM Threads | ThreadsNew | Connections | Number of Threads in NEW State |
| | ThreadsRunnable | Connections | Number of Threads in RUNNABLE State |
| | ThreadsBlocked | Connections | Number of Threads in BLOCKED State |
| | ThreadsWaiting | Connections | Number of Threads in WAITING State |
| | ThreadsTimedWaiting | Connections | Number of Threads in TIMED WAITING State |
| | ThreadsTerminated | Connections | Current Number of Threads in TERMINATED State |

| | | | |
|-------------------------------|-----------------------------------|-------------|--|
| Number of Regions | regionCount | Connections | Number of Regions |
| Region Localization | percentFilesLocal | % | Proportion of Region's HFile Located on Local HDFS Data Node |
| Region Replica Localization | percentFilesLocalSecondaryRegions | % | Proportion of Region Replica's HFile Located on Local HDFS Data Node |
| Number of RPC Authentications | authenticationFailures | Threads | RPC Authentication Failure Instances |
| | authenticationSuccesses | Threads | RPC Authentication Success Instances |
| RPC Connection Count | numOpenConnections | Connections | RPC Connection Count |
| RPC Exception Count | FailedSanityCheckException | Threads | FailedSanityCheckException Exception Instances |
| | NotServingRegionException | Threads | NotServingRegionException Exception Instances |
| | OutOfOrderScannerNextException | Threads | OutOfOrderScannerNextException Exception Instances |
| | RegionMovedException | Threads | RegionMovedException Exception Instances |
| | RegionTooBusyException | Threads | RegionTooBusyException Exception Instances |
| | UnknownScannerException | Threads | UnknownScannerException Exception Instances |
| RPC Handle Count | numActiveHandler | Connections | RPC Handle Count |
| | numActiveWriteHandler | Connections | RPC Read Handle Count |
| | numActiveReadHandler | Connections | RPC Write Handle Count |
| | numActiveScanHandler | Connections | RPC Scan Handle Count |
| RPC Queue Request Count | numCallsInPriorityQueue | Connections | Priority Queue RPC Request Count |
| | numCallsInReplicationQueue | Connections | Replication Queue RPC Request Count |
| | numCallsInPriorityQueue | Connections | General Queue RPC Request Count |
| | numCallsInWriteQueue | Connections | Write Invocation Queue RPC Request Count |
| | numCallsInReadQueue | Connections | Read Invocation Queue RPC Request Count |
| | numCallsInScanQueue | Connections | Scan Invocation Queue RPC Request Count |
| WAL File Quantity | hlogFileCount | Connections | WAL File Quantity |
| WAL File Size | hlogFileSize | Byte | WAL File Size |
| Memstore Size | memStoreSize | MB | Memstore Size |
| Store Quantity | storeCount | Connections | Store Quantity |
| Storefile Quantity | storeFileCount | Connection | Storefile Quantity |

| | | | |
|---|-----------------------------|-------------|--|
| | | s | |
| Storefile Size | storeFileSize | MB | Storefile Size |
| Disk Write Speed | flushedCellsSize | bytes/s | Disk Write Speed |
| Average Latency | Append_mean | ms | Average Latency of Append Requests |
| | Replay_mean | ms | Average Latency of Replay Requests |
| | Get_mean | ms | Average Latency of Get Requests |
| | updatesBlockedTime | ms | Updates Blocked Time Prior to Memstore Flush |
| Number of Disk Writes by RS | FlushTime_num_ops | Threads | Memstore Flush Disk Write Instances |
| Number of Requests in the Operation Queue | splitQueueLength | Connections | Split Queue Length |
| | compactionQueueLength | Connections | Compaction Queue Length |
| | flushQueueLength | Connections | Region Flush Queue Length |
| Replay Operation Instances | Replay_num_ops | Threads | Replay Operation Instances |
| Number of Slow Operations | slowAppendCount | Threads | Number of Append Requests Exceeding 1s in Duration |
| | slowDeleteCount | Threads | Number of Delete Requests Exceeding 1s in Duration |
| | slowGetCount | Threads | Number of Get Requests Exceeding 1s in Duration |
| | slowIncrementCount | Threads | Number of Increment Requests Surpassing 1s in Duration |
| | slowPutCount | Threads | Number of Put Requests Exceeding 1s in Duration |
| Split Requests | splitRequestCount | Threads | Split Request Number |
| | splitSuccessCount | Threads | Split Successful Instances |
| Quantity of Cache Blocks | blockCacheCount | Connections | Number of Blocks in Block Cache |
| | blockCacheHitCount | Connections | Number of Block Cache Hit Requests |
| | blockCacheMissCount | Connections | Number of Block Cache Miss Requests |
| Read Cache Hit Rate | blockCacheExpressHitPercent | % | Read Cache Hit Rate |
| Memory Occupancy Size of Cache Blocks | blockCacheSize | Byte | Memory Occupancy Size of Cache Blocks |
| Index Size | staticBloomSize | Byte | Uncompressed Static Bloom Filters Size |
| | staticIndexSize | Byte | Uncompressed Static Index Size |
| | storeFileIndexSize | Byte | Index Size in StoreFile on Disk |
| Read/Write Traffic Volume | receivedBytes | bytes/s | Read/Write Traffic Volume |
| | sentBytes | bytes/s | Data Received |
| | | | |

| | | | |
|-----------------------------|-------------------------------|-------------|---|
| Read/Write Request Volume | Total | count/s | Total Request Volume, this value will be less than the sum of read and write requests when a Scan request is present. |
| | Read | count/s | Read Requests |
| | Write | count/s | Volume of Write Requests |
| | Append_num_ops | count/s | Append Request Volume |
| | Mutate_num_ops | count/s | Mutate Request Volume |
| | Delete_num_ops | count/s | Delete Request Volume |
| | Increment_num_ops | count/s | Increment Request Volume |
| | Get_num_ops | count/s | Get Request Volume |
| | Put_num_ops | count/s | Put Request Volume |
| | ScanTime_num_ops | count/s | Scan Request Volume |
| | ScanSize_num_ops | count/s | Scan Request Volume |
| Number of Mutations | mutationsWithoutWALCount | Connections | Number of Mutations |
| Mutation Size | mutationsWithoutWALSize | Byte | Mutation Size |
| Process Initiation Time | regionServerStartTime | s | Process Initiation Time |
| Synchronized Log | source.sizeOfLogQueue | Connections | Synchronized Log Length |
| Synchronization Duration | source.ageOfLastShippedOp | ms | Synchronization Duration |
| Request Volume | ReadRequestCount | count/s | Read Requests per Second |
| | WriteRequestCount | count/s | Write Requests per Second |
| Request Volume | Read | count/s | Read Requests per Second |
| | Write | count/s | Write Requests per Second |
| Store Size | memstoreSize | Byte | Memstore Size |
| | storeFileSize | Byte | StoreFile Size |
| Table-Level Request Latency | getTime_99th_percentile | ms | 99% Request Processing Latency |
| | scanTime_99th_percentile | ms | 99% Request Processing Latency |
| | putTime_99th_percentile | ms | 99% Request Processing Latency |
| | incrementTime_99th_percentile | ms | 99% Request Processing Latency |
| | appendTime_99th_percentile | ms | 99% Request Processing Latency |
| | deleteTime_99th_percentile | ms | 99% Request Processing Latency |
| Request Processing Latency | 99th_percentile | ms | 99% Request Processing Latency |
| | 99.9th_percentile | ms | 99.9% Request Processing Latency |
| Request Queuing Latency | 99th_percentile | ms | 99% Request Queuing Latency |
| | 99.9th_percentile | ms | 99.9% Request Queuing Latency |
| SCAN SIZE | max | bytes | Maximum ScanSize |

| | | | |
|---|----------------------|---------|---|
| | mean | bytes | Average ScanSize |
| | min | bytes | Minimum ScanSize |
| SCAN TIME | max | s | Maximum ScanTime |
| | mean | s | Average ScanTime |
| | min | s | Minimum ScanTime |
| Bulkload Latency | 99th_percentile | ms | Bulkload Latency |
| | 999th_percentile | ms | |
| Append Latency | 99th_percentile | ms | Append Latency |
| | 999th_percentile | ms | |
| Delete Latency | 99th_percentile | ms | Delete Latency |
| | 999th_percentile | ms | |
| MultiGet Latency | 99th_percentile | ms | MultiGet Latency |
| | 999th_percentile | ms | |
| Get Latency | 99th_percentile | ms | Get Latency |
| | 999th_percentile | ms | |
| PutBatch Latency | 99th_percentile | ms | PutBatch Latency |
| | 999th_percentile | ms | |
| Put Latency | 99th_percentile | ms | Put Latency |
| | 999th_percentile | ms | |
| Increment Latency | 99th_percentile | ms | Increment Latency |
| | 999th_percentile | ms | |
| Compacted Count Rate | MinorCompactedCells | count/s | Average Minor Compacted Cells per Second |
| | MajorCompactedCells | count/s | Average Major Compacted Cells per Second |
| Compacted Size Rate | Minor CompactedCells | bytes/s | Average Minor Compacted Cells Size per Second |
| | MajorCompactedCells | bytes/s | Average Major Compacted Cells Size per Second |
| Region Localization | percentFilesLocal | % | Proportion of Region HFile Located in Local HDFS Data Node |
| Average Latency | updatesBlockedTime | ms | Updates Blocked Time Prior to Memstore Flush |
| pauseThresholdExceeded | info | count | Count of INFO Level Pause Alerts |
| | warn | count | Count of WARN Level Pause Alerts |
| Number of Pauses Not Caused by Garbage Collection | ops | count | Number of Pauses Not Caused by Garbage Collection |
| Maximum Duration of Pauses Not Caused by Garbage Collection | max | ms | Maximum Duration of Pauses Not Caused by Garbage Collection |
| Number of Pauses Caused by Garbage Collection | ops | count | Number of Pauses Caused by Garbage Collection |

| | | | |
|---|-----|----|---|
| Maximum Duration of Pauses Caused by Garbage Collection | max | ms | Maximum Duration of Pauses Caused by Garbage Collection |
|---|-----|----|---|

Hive Monitoring Metrics

Last updated: 2023-12-21 16:50:39

HIVE-HiveMetaStore

| Title | Description | Metric Unit | Description |
|------------------------------|-------------------------|----------------|--|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| JVM Memory | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of HeapMemory Committed by JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| | MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| Heap Memory Utilization Rate | MemHeapUsedRate | % | Percentage of Currently Utilized HeapMemory in JVM Relative to the Configured HeapMemory Quantity in JVM |
| Number of File Descriptors | OpenFileDescriptorCount | Connections | Quantity of Open File Descriptors |
| | MaxFileDescriptorCount | Connections | Maximum Number of File Descriptors |
| CPU Utilization Rate | ProcessCpuLoad | % | Process CPU Utilization |
| | SystemCpuLoad | % | System CPU Utilization |
| Proportion of CPU Usage Time | CPURate | seconds/second | Proportion of CPU Usage Time |
| Number of Working Threads | DaemonThreadCount | Connections | Daemon Thread Count |
| | ThreadCount | Connections | Total Thread Count |
| Cumulative CPU Usage Time | ProcessCpuTime | ms | Cumulative CPU Usage Time |
| | | | |

| | | | |
|--|--|-------------|--|
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| GC Additional Sleep Duration | ExtraSleepTime | ms/s | GC Additional Sleep Duration |
| Alter Table Request Time | HIVE.HMS.API_ALTER_TABLE | ms | Average Time for Alter Table Request |
| Alter Table with Environmental Context Request Time | HIVE.HMS.API_ALTER_TABLE_WITH_ENV_CONTEXT | ms | Average Time for Alter Table Request with Environmental Context |
| Create Table Request Time | HIVE.HMS.API_CREATE_TABLE | ms | Average Time for Create Table Request |
| Create Table with Environmental Context Request Time | HIVE.HMS.API_CREATE_TABLE_WITH_ENV_CONTEXT | ms | Average Time for Create Table Request with Environmental Context |
| Drop Table Request Time | HIVE.HMS.API_DROP_TABLE | ms | Average Drop Table Request Time |
| Drop Table with Environmental Context Request Time | HIVE.HMS.API_DROP_TABLE_WITH_ENV_CONTEXT | ms | Average Drop Table with Environmental Context Request Time |
| Get Table Request Time | HIVE.HMS.API_GET_TABLE | ms | Average Get Table Request Time |
| Get Tables Request Time | HIVE.HMS.API_GET_TABLES | ms | Average Get Tables Request Time |
| Get Multiple Tables Request Time | HIVE.HMS.API_GET_MULTITABLE | ms | Average Get Multiple Tables Request Time |
| Get Table Request Time | HIVE.HMS.API_GET_TABLE_REQ | ms | Average Get Table Request Time |
| Get Database Request Time | HIVE.HMS.API_GET_DATABASE | ms | Average Get Database Request Time |
| Get Databases Request Time | HIVE.HMS.API_GET_DATABASES | ms | Average Get Databases Request Time |
| Get All Databases Request Time | HIVE.HMS.API_GET_ALL_DATABASES | ms | Average Get All Databases Request Time |
| Get All Functions Request Time | HIVE.HMS.API_GET_ALL_FUNCTIONS | ms | Average Get All Functions Request Time |
| Current Active Create Table Request Count | HIVE.HMS.ACTIVE_CALLS_API_CREATE_TABLE | Connections | Current Active Create Table Request Count |
| Current Active Drop Table Request Count | HIVE.HMS.ACTIVE_CALLS_API_DROP_TABLE | Connections | Current Active Drop Table Request Count |
| Current Active Alter Table Request Count | HIVE.HMS.ACTIVE_CALLS_API_ALTER_TABLE | Connections | Current Active Alter Table Request Count |

HIVE-HiveServer2

| Title | Description | Metric Unit | Description |
|------------------------------|-------------|-------------|--|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |

| | | | |
|--|-------------------------|----------------|--|
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| JVM Memory | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| | MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| Heap Memory Utilization Ratio | MemHeapUsedRate | % | Percentage of Currently Utilized HeapMemory in JVM Relative to the Configured HeapMemory Quantity in JVM |
| CPU Utilization Rate | ProcessCpuLoad | % | CPU Utilization Rate |
| Proportion of CPU Usage Time | CPUUsedRate | seconds/second | Proportion of CPU Usage Time |
| Number of File Descriptors | MaxFileDescriptorCount | Connections | Maximum Number of File Descriptors |
| | OpenFileDescriptorCount | Connections | Number of Open File Descriptors |
| Cumulative CPU Usage Time | ProcessCpuTime | ms | Cumulative CPU Usage Time |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| Number of Working Threads | DaemonThreadCount | Connections | Number of Daemon Threads |
| | ThreadCount | Connections | Total Number of Threads |
| Driver Execution Latency | 99th_percentile | ms | 99% Latency of Driver Execution |
| | Avg | ms | Average Latency of Driver Execution |
| Number of Open Connections | NumOpenConnections | Connections | Number of Open Connections |
| Current Size of HS2 Asynchronous Thread Pool | PoolSize | Connections | Current Size of HS2 Asynchronous Thread Pool |
| Current Size of HS2 Asynchronous Operation Queue | QueueSize | Connections | Current Size of HS2 Asynchronous Operation Queue |

| | | | |
|---|-------------------------------|-------------|---|
| Number of Hive Operations | Closed | Connections | Number of Closed Operations |
| | Finished | Connections | Number of Completed Operations |
| | Canceled | Connections | Number of Cancelled Operations |
| | Error | Per/ | Number of Erroneous Operations |
| GC Additional Sleep Duration | ExtraSleepTime | ms/s | GC Additional Sleep Duration |
| API Request Count | HIVE.H2.ACTIVE.CALLS.API | Count | Current serializePlan Request Count |
| | | Count | Current semanticAnalyze Request Count |
| | | Count | Current runtask Request Count |
| | | Count | Current releaseLocks Request Count |
| | | Count | Current getSplits Count |
| Duration of SQL Tasks in Pending Status | HIVE.H2.SQL.OPERATION.PENDING | ms | Average Duration of SQL Tasks in Pending Status |
| Duration of SQL Tasks in Running Status | HIVE.H2.SQL.OPERATION.RUNNING | ms | Average Duration of SQL Tasks in Running Status |
| Current Active User Count | HIVE.H2.SQL.OPERATION | Count | Number of Currently Active Users |
| Execution Time of the Query | HIVE.H2.EXECUTING.QUERIES | ms | Average Execution Time of Queries |
| Submission Time of the Query | HIVE.H2.SUBMITTED.QUERIES | ms | Submission Time of the Query |
| 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 | OneMinuteRate of Failed Queries |
| Number of Working Threads | HIVE.H2.THREAD.COUNT | Connections | Number of JVM Blocked Threads |
| | | Connections | Number of JVM Terminate Threads |
| | | Connections | Number of JVM Deadlock Threads |
| | | Connections | Number of New JVM Threads |
| | | Connections | Number of Runnable JVM Threads |
| | | Connections | Number of JVM Timed-Waiting Threads |
| | | Connections | Number of JVM Waiting Threads |
| Number of Sessions | HIVE.H2.OPEN.SESIONS | Connections | Number of Open Sessions |
| Number of Currently Active Sessions | HIVE.H2.ACTIVE.SESIONS | Connections | Number of Active Sessions |

HIVE-HiveWebHcat

| Title | Description | Metric Unit | Description |
|------------------------------|-------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |

Spark Monitoring Metrics

Last updated: 2023-12-21 16:50:50

SPARK-HistoryServer

| Title | Description | Metric Unit | Description |
|------------------------------|-------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |

Presto Monitoring Metrics

Last updated: 2023-12-21 16:50:58

PRESTO – Overview

| Title | Description | Metric Unit | Description |
|-------------------------------------|-------------------------|-------------|---|
| Node Quantity | Active | Connections | Number of Active Nodes |
| | Total | Connections | Total Node Count |
| | Failed | Connections | Number of Failed Nodes |
| Query | RunningQueries | Connections | Total Number of Running Queries |
| | QueuedQueries | Connections | Total Number of Queries in Waiting Status |
| Query Frequency | FailedQueries | Per Minute | Total Number of Failed Queries |
| | AbandonedQueries | Per Minute | Total Number of Abandoned Queries |
| | CanceledQueries | Per Minute | Total Number of Canceled Queries |
| | CompletedQueries | Per Minute | Total Number of Completed Queries |
| | StartedQueries | Per Minute | Total Number of Initiated Queries |
| Data Input/Output Volume Per Minute | InputDataSizeOneMinute | GB/min | Input Data Rate |
| | OutputDataSizeOneMinute | GB/min | Output Data Rate |

PRESTO-Worker

| Title | Description | Metric Unit | Description |
|------------------------------|-----------------|-------------|--|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| Heap Memory Utilization Rate | MemHeapUsedRate | % | Percentage of Currently Utilized HeapMemory in JVM Relative to the Configured HeapMemory Quantity in JVM |
| JVM Memory | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by |

| | | | | | |
|--------------------------|-------------------------------|-------------|--|--|--|
| | | | JVM | | |
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM | | |
| | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM | | |
| | MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM | | |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM | | |
| | MemHeapInitM | MB | Initial Quantity of HeapMem in JVM | | |
| | MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM | | |
| Data Input-Output Rate | InputDataSize.OneMinute.Rate | GB/min | Input Data Rate | | |
| | OutputDataSize.OneMinute.Rate | GB/min | Output Data Rate | | |
| Processes | PeakThreadCount | Connections | Peak Thread Count | | |
| | ThreadCount | Connections | Number of Threads | | |
| | DaemonThreadCount | Connections | Background Thread Count | | |
| Process Runtime Duration | Uptime | s | Process Runtime Duration | | |
| Process Initiation Time | StartTime | s | Process Initiation Time | | |
| File Descriptor Count | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity | | |
| | OpenFileDescriptorCount | Connections | Quantity of Open File Descriptors | | |

PRESTO-Coordinator

| Title | Description | Metric Unit | Description |
|------------------------------|-------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |

| | | | |
|------------------------------|-------------------------|-------------|--|
| JVM Memory | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| | MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| Heap Memory Utilization Rate | MemHeapUsedRate | % | Percentage of Currently Utilized HeapMemory in JVM Relative to the Configured HeapMemory Quantity in JVM |
| Processes | PeakThreadCount | Connections | Peak Thread Count |
| | ThreadCount | Connections | Number of Threads |
| | DaemonThreadCount | Connections | Background Thread Count |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| Process Initiation Time | StartTime | s | Process Initiation Time |
| File Descriptor Count | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| | OpenFileDescriptorCount | Connections | Quantity of Open File Descriptors |

Trino Monitoring Metrics

Last updated: 2023-12-21 16:51:04

TRINO – Overview

| Title | Description | Metric Unit | Description |
|-------------------------------------|-------------------------|-------------|---|
| Node Quantity | Active | Connections | Number of Active Nodes |
| | Total | Connections | Total Node Count |
| | Failed | Connections | Number of Failed Nodes |
| Query | RunningQueries | Connections | Total Number of Running Queries |
| | QueuedQueries | Connections | Total Number of Queries in Waiting Status |
| Query Frequency | FailedQueries | Per Minute | Total Number of Failed Queries |
| | AbandonedQueries | Per Minute | Total Number of Abandoned Queries |
| | CanceledQueries | Per Minute | Total Number of Canceled Queries |
| | CompletedQueries | Per Minute | Total Number of Completed Queries |
| | StartedQueries | Per Minute | Total Number of Initiated Queries |
| Data Input/Output Volume Per Minute | InputDataSizeOneMinute | GB/min | Input Data Rate |
| | OutputDataSizeOneMinute | GB/min | Output Data Rate |

TRINO-WORKER

| Title | Description | Metric Unit | Description |
|------------------------------|-----------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| JVM Memory | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |

| | | | |
|------------------------------|-------------------------------|-------------|--|
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| | MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| Heap Memory Utilization Rate | MemHeapUsedRate | % | Percentage of Currently Utilized HeapMemory in JVM Relative to the Configured HeapMemory Quantity in JVM |
| Data Input-Output Rate | InputDataSize.OneMinute.Rate | GB/min | Input Data Rate |
| | OutputDataSize.OneMinute.Rate | GB/min | Output Data Rate |
| Number of Working Threads | PeakThreadCount | Connections | Peak Thread Count |
| | ThreadCount | Connections | Number of Threads |
| | DaemonThreadCount | Connections | Background Thread Count |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| File Descriptor Count | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| | OpenFileDescriptorCount | Connections | Quantity of Open File Descriptors |

TRINO-Coordinator

| Title | Description | Metric Unit | Description |
|------------------------------|-----------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| JVM Memory | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | | | |

| | | | |
|------------------------------|-------------------------|-------------|--|
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| | MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| Heap Memory Utilization Rate | MemHeapUsedRate | % | Percentage of Currently Utilized HeapMemory in JVM Relative to the Configured HeapMemory Quantity in JVM |
| Number of Working Threads | PeakThreadCount | Connections | Peak Thread Count |
| | ThreadCount | Connections | Number of Threads |
| | DaemonThreadCount | Connections | Background Thread Count |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| Process Initiation Time | StartTime | s | Process Initiation Time |
| File Descriptor Count | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| | OpenFileDescriptorCount | Connections | Quantity of Open File Descriptors |

ClickHouse Monitoring Metrics

Last updated: 2023-12-21 16:51:12

CLICKHOUSE-Metrics

| Title | Description | Metric Unit | Description |
|--|----------------------------|-------------|---|
| Number of Network Connections | tcp | Connections | Number of connections to the TCP server |
| | http | Connections | Number of connections to the HTTP server |
| Number of ZK event subscriptions | watches | Connections | Number of ZK event subscriptions |
| Number of temporary nodes stored in ZooKeeper | ephemeralNode | Connections | Number of temporary nodes stored in ZooKeeper |
| Number of active tasks in the BackgroundPool | backgroundPoolTask | Connections | Number of active tasks in the BackgroundProcessingPool |
| | backgroundSchedulePoolTask | Connections | Number of active tasks in the BackgroundSchedulePool |
| Number of threads waiting for locks in the Context | contextLockWait | Connections | Number of threads waiting for locks in the Context |
| Number of suppressed Insert queries | delayedInserts | Connections | Number of suppressed Insert queries |
| Number of requests in the data source of the cache type dictionary | dictCacheRequests | Connections | Number of requests in the data source of the cache type dictionary |
| Number of pending asynchronous insertions to distributed tables | distributedSend | Connections | Number of pending asynchronous insertions to distributed tables |
| Threads | global | Connections | Number of threads in the global thread pool |
| | globalActive | Connections | Number of active threads in the global thread pool |
| | local | Connections | Number of threads in the local thread pool |
| | localActive | Connections | Number of active threads in the local thread pool |
| Number of replicas participating in the leader election | leaderElection | Connections | Number of replicas participating in the leader election |
| Quantity of replicated tables | leaderReplica | Connections | Quantity of Replicated tables that are leaders |
| | readonlyReplica | Connections | Quantity of Replicated tables in read-only state |
| Total amount of allocated memory | memoryTracking | GB | Total amount of memory allocated in the currently executing queries |
| | backgroundProcessingPool | GB | Total amount of memory allocated in the background processing pool |
| | backgroundSchedulePool | GB | Total amount of memory allocated in the background scheduling pool |

| | | | |
|---|-------------------------|-------------|---|
| | memoryTrackingForMerges | GB | Total amount of memory allocated for background merges |
| Number of merges currently being executed in the background | merge | Connections | Number of merges currently being executed in the background |
| Number of open files | read | Connections | Number of open readable files |
| | write | Connections | Number of open writable files |
| Number of table alterations | partMutation | Connections | Number of table alterations |
| Number of threads processing queries | queryThread | Connections | Number of threads processing queries |
| Number of halted or pending queries | queryPreempted | Connections | Number of halted or pending queries |
| Number of system calls | read | Connections | Number of read system calls |
| | write | Connections | Number of write system calls |
| Number of data blocks | fetch | Connections | Quantity of data blocks collected from replicas |
| | send | Connections | Quantity of blocks sent to replicas |
| | check | Connections | Quantity of data blocks undergoing consistency checks |
| Modifications to the server | revision | Connections | Modifications to the server |
| Version No. | version | 1 | Version No. |
| Number of threads awaiting possession of read-write locks | waitingRead | Connections | Number of threads waiting for a read-write lock to read a table |
| | waitingWrite | Connections | Number of threads waiting for a read-write lock to write a table |
| | activeRead | Connections | Number of threads holding a read lock in a table's read-write lock |
| | activeWrite | Connections | Number of threads holding a write lock in a table's read-write lock |
| Number of rows in buffer tables | storageBufferRows | Connections | Number of rows in buffer tables |
| Number of bytes in buffer tables | storageBufferBytes | MB | Number of bytes in buffer tables |

CLICKHOUSE-Events

| Title | Description | Metric Unit | Description |
|-------------------------|-------------|-------------|---|
| Queries | total | Connections | Total number of potential queries to be executed |
| | select | Connections | Number of potential Select queries to be executed |
| | insert | Connections | Number of potential Insert queries to be executed |
| Number of rows inserted | insertedRow | Connections | Number of rows inserted into all tables |

| | | | |
|---|---------------|------------------------|--|
| | s | | |
| Number of bytes inserted | insertedBytes | Connections | Number of bytes inserted into all tables |
| Total time spent waiting for system calls | read | Microsecond (μ s) | Total time spent waiting for read system calls |
| | write | Microsecond (μ s) | Total time spent waiting for write system calls |
| Number of open files | fileOpen | Connections | Number of open files |
| Number of read/write operations from file descriptors | read | Connections | Number of reads from the file descriptor |
| | write | Connections | Number of writes to the file descriptor |
| Number of bytes read/written from file descriptors | read | MB | Number of bytes read from the file descriptor |
| | write | MB | Number of bytes written to the file descriptor |
| Total time expended by the processing thread | realtime | Microsecond (μ s) | Total time expended by the processing thread |
| | user | Microsecond (μ s) | Total time expended by the processing thread executing CPU instructions within the user space |
| | system | Microsecond (μ s) | Total time expended by the processing thread executing CPU instructions within the operating system kernel space |
| Number of compiled regular expressions | regexCreated | Connections | Number of compiled regular expressions |

CLICKHOUSE-Asynchronous_metrics

| Title | Description | Metric Unit | Description |
|---|--------------------------|------------------------|--|
| Cache size of marks in StorageMergeTree | markCacheBytes | MB | Cache size of marks in StorageMergeTree |
| Number of cache files for marks in StorageMergeTree | markCacheFiles | Connections | Number of cache files for marks in StorageMergeTree |
| Maximum value of active data blocks | maxPartCountForPartition | Connections | Number of the largest active data blocks in partitions |
| Quantity of databases | databaseCount | Connections | Quantity of databases |
| Number of data tables | tableCount | Connections | Number of data tables |
| Maximum latency of replica | absolute | Microsecond (μ s) | Maximum latency of the replica queue |
| | relative | Microsecond (μ s) | Maximum value of the absolute latency difference from other blocks |
| Size of the pending operations queue | replicasMaxQueueSize | Connections | Size of the pending operations queue |

Druid Monitoring Metrics

Last updated: 2023-12-21 16:51:29

Druid-broker

| Description | Unit | Description |
|----------------------------|-------------|---|
| YGC | Threads | Young Garbage Collection Frequency |
| FGC | Threads | Full Garbage Collection Frequency |
| FGCT | s | Full Garbage Collection Time Consumption |
| GCT | s | Garbage Collection Time Expenditure |
| YGCT | s | Young Garbage Collection Time Consumption |
| S0 | % | Survivor 0 Region Memory Usage Proportion |
| E | % | Eden Region Memory Usage Proportion |
| CCS | % | Compressed Class Space Region Memory Usage Proportion |
| S1 | % | Survivor 1 Region Memory Usage Proportion |
| O | % | Old Region Memory Usage Proportion |
| M | % | Metaspace Region Memory Usage Proportion |
| MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| ThreadsNew | Connections | Current Quantity of Threads in NEW State |
| ThreadsRunnable | Connections | Current Quantity of Threads in RUNNABLE State |
| ThreadsBlocked | Connections | Current Quantity of Threads in BLOCKED State |
| ThreadsWaiting | Connections | Current Quantity of Threads in WAITING State |
| ThreadsTimedWaiting | Connections | Current Quantity of Threads in TIMED_WAITING State |
| ThreadsTerminated | Connections | Current Number of Threads in TERMINATED State |
| LogFatal | Threads | Number of FATAL Level Logs |
| LogError | Threads | Number of ERROR Level Logs |
| LogWarn | Threads | Number of WARN Level Logs |
| LogInfo | Threads | Number of INFO Level Logs |
| jetty.numOpenConnections | Connections | Number of Active Jetty Connections |
| segment.scan.pending | Connections | Number of segments in queue waiting to be scanned |
| broker.query.count | Threads | number of total queries |
| broker.query.success.count | Threads | number of queries successfully processed |

| | | |
|--------------------------------|-------------|--|
| broker.query.failed.count | Threads | number of failed queries |
| broker.query.interrupted.count | Threads | number of queries interrupted due to cancellation or timeout |
| normal.count | Connections | Number of Queries with a Delay of < 1s |
| abnormal.count | Connections | Number of Queries with a Delay of >= 1s |

Druid- coordinator

| Description | Unit | Description |
|----------------------------|-------------|--|
| segment.assigned.count | Connections | Number of Segments Loaded into the Druid Cluster |
| segment.moved.count | Connections | Number of Segments Relocated within the Druid Cluster |
| segment.dropped.count | Connections | Number of Segments Expired and Deleted within the Druid Cluster |
| segment.deleted.count | Connections | Number of Segments Deleted within the Druid Cluster due to Rule Configuration |
| segment.unneeded.count | Connections | Number of Segments Deleted within the Druid Cluster due to Being Set as No Longer in Use |
| 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 | Connections | Number of segments that failed to load. |
| segment.loadQueue.count | Connections | Number of segments to load. |
| segment.dropQueue.count | Connections | Number of segments to drop. |
| segment.overshadowed.count | Connections | Number of overshadowed segments. |
| tier.historical.count | Connections | Number of available historical nodes in each tier. |
| tier.replication.factor | Connections | 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 | Connections | Number of Compact Tasks |
| YGC | Threads | Young Garbage Collection Frequency |
| FGC | Threads | Full Garbage Collection Frequency |
| FGCT | s | Full Garbage Collection Time Consumption |
| GCT | s | Garbage Collection Time Expenditure |

| | | |
|-------------------------------|-------------|---|
| YGCT | s | Young Garbage Collection Time Consumption |
| S0 | % | Survivor 0 Region Memory Usage Proportion |
| E | % | Eden Region Memory Usage Proportion |
| CCS | % | Compressed Class Space Region Memory Usage Proportion |
| S1 | % | Survivor 1 Region Memory Usage Proportion |
| O | % | Old Region Memory Usage Proportion |
| M | % | Metaspace Region Memory Usage Proportion |
| MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| ThreadsNew | Connections | Current Quantity of Threads in NEW State |
| ThreadsRunnable | Connections | Current Quantity of Threads in RUNNABLE State |
| ThreadsBlocked | Connections | Current Quantity of Threads in BLOCKED State |
| ThreadsWaiting | Connections | Current Quantity of Threads in WAITING State |
| ThreadsTimedWaiting | Connections | Current Quantity of Threads in TIMED_WAITING State |
| ThreadsTerminated | Connections | Current Number of Threads in TERMINATED State |
| LogFatal | Threads | Number of FATAL Level Logs |
| LogError | Threads | Number of ERROR Level Logs |
| LogWarn | Threads | Number of WARN Level Logs |
| LogInfo | Threads | Number of INFO Level 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 | Connections | 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 | Connections | 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 | Connections | Number of segments (including replicas) left to load until segments that should be loaded in the cluster are available for queries. |
| jetty.numOpenConnections | Connections | Number of Active Jetty Connections |

Druid- historical

| Description | Unit | Description |
|--------------------------------|-------------|--|
| YGC | Threads | Young Garbage Collection Frequency |
| FGC | Threads | Full Garbage Collection Frequency |
| FGCT | s | Full Garbage Collection Time Consumption |
| GCT | s | Garbage Collection Time Expenditure |
| YGCT | s | Young Garbage Collection Time Consumption |
| S0 | % | Survivor 0 Region Memory Usage Proportion |
| E | % | Eden Region Memory Usage Proportion |
| CCS | % | Compressed Class Space Region Memory Usage Proportion |
| S1 | % | Survivor 1 Region Memory Usage Proportion |
| O | % | Old Region Memory Usage Proportion |
| M | % | Metaspace Region Memory Usage Proportion |
| MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| ThreadsNew | Connections | Current Quantity of Threads in NEW State |
| ThreadsRunnable | Connections | Current Quantity of Threads in RUNNABLE State |
| ThreadsBlocked | Connections | Current Quantity of Threads in BLOCKED State |
| ThreadsWaiting | Connections | Current Quantity of Threads in WAITING State |
| ThreadsTimedWaiting | Connections | Current Quantity of Threads in TIMED_WAITING State |
| ThreadsTerminated | Connections | Current Number of Threads in TERMINATED State |
| LogFatal | Threads | Number of FATAL Level Logs |
| LogError | Threads | Number of ERROR Level Logs |
| LogWarn | Threads | Number of WARN Level Logs |
| LogInfo | Threads | Number of INFO Level Logs |
| jetty.numOpenConnections | Connections | Number of Active Jetty Connections |
| segment.scan.pending | Connections | 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 | Threads | Total Number of Historical Queries |
| historical.query.success.count | Threads | Number of Successful Historical Queries |

| | | |
|------------------------------------|-------------|--|
| historical.query.failed.count | Threads | Number of Failed Historical Queries |
| historical.query.interrupted.count | Threads | Number of Historical Queries Interrupted |
| normal.count | Connections | Number of Queries with a Delay of < 1s |
| abnormal.count | Connections | Number of Queries with a Delay of >= 1s |

Druid- middleManager

| Description | Unit | Description |
|--------------------------|-------------|---|
| YGC | Threads | Young Garbage Collection Frequency |
| FGC | Threads | Full Garbage Collection Frequency |
| FGCT | s | Full Garbage Collection Time Consumption |
| GCT | s | Garbage Collection Time Expenditure |
| YGCT | s | Young Garbage Collection Time Consumption |
| S0 | % | Survivor 0 Region Memory Usage Proportion |
| E | % | Eden Region Memory Usage Proportion |
| CCS | % | Compressed Class Space Region Memory Usage Proportion |
| S1 | % | Survivor 1 Region Memory Usage Proportion |
| O | % | Old Region Memory Usage Proportion |
| M | % | Metaspace Region Memory Usage Proportion |
| MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| ThreadsNew | Connections | Current Quantity of Threads in NEW State |
| ThreadsRunnable | Connections | Current Quantity of Threads in RUNNABLE State |
| ThreadsBlocked | Connections | Current Quantity of Threads in BLOCKED State |
| ThreadsWaiting | Connections | Current Quantity of Threads in WAITING State |
| ThreadsTimedWaiting | Connections | Current Quantity of Threads in TIMED_WAITING State |
| ThreadsTerminated | Connections | Current Number of Threads in TERMINATED State |
| LogFatal | Threads | Number of FATAL Level Logs |
| LogError | Threads | Number of ERROR Level Logs |
| LogWarn | Threads | Number of WARN Level Logs |
| LogInfo | Threads | Number of INFO Level Logs |
| jetty.numOpenConnections | Connections | Number of Active Jetty Connections |

Druid- overlord

| Description | Unit | Description |
|--------------------------|-------------|---|
| YGC | Threads | Young Garbage Collection Frequency |
| FGC | Threads | Full Garbage Collection Frequency |
| FGCT | s | Full Garbage Collection Time Consumption |
| GCT | s | Garbage Collection Time Expenditure |
| YGCT | s | Young Garbage Collection Time Consumption |
| S0 | % | Survivor 0 Region Memory Usage Proportion |
| E | % | Eden Region Memory Usage Proportion |
| CCS | % | Compressed Class Space Region Memory Usage Proportion |
| S1 | % | Survivor 1 Region Memory Usage Proportion |
| O | % | Old Region Memory Usage Proportion |
| M | % | Metaspace Region Memory Usage Proportion |
| MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| ThreadsNew | Connections | Current Quantity of Threads in NEW State |
| ThreadsRunnable | Connections | Current Quantity of Threads in RUNNABLE State |
| ThreadsBlocked | Connections | Current Quantity of Threads in BLOCKED State |
| ThreadsWaiting | Connections | Current Quantity of Threads in WAITING State |
| ThreadsTimedWaiting | Connections | Current Quantity of Threads in TIMED_WAITING State |
| ThreadsTerminated | Connections | Current Number of Threads in TERMINATED State |
| LogFatal | Threads | Number of FATAL Level Logs |
| LogError | Threads | Number of ERROR Level Logs |
| LogWarn | Threads | Number of WARN Level Logs |
| LogInfo | Threads | Number of INFO Level Logs |
| jetty.numOpenConnections | Connections | Number of Active Jetty Connections |

Druid- router

| Description | Unit | Description |
|-------------|---------|--|
| YGC | Threads | Young Garbage Collection Frequency |
| FGC | Threads | Full Garbage Collection Frequency |
| FGCT | s | Full Garbage Collection Time Consumption |
| GCT | s | Garbage Collection Time Expenditure |

| | | |
|--------------------------|-------------|---|
| YGCT | s | Young Garbage Collection Time Consumption |
| S0 | % | Survivor 0 Region Memory Usage Proportion |
| E | % | Eden Region Memory Usage Proportion |
| CCS | % | Compressed Class Space Region Memory Usage Proportion |
| S1 | % | Survivor 1 Region Memory Usage Proportion |
| O | % | Old Region Memory Usage Proportion |
| M | % | Metaspace Region Memory Usage Proportion |
| MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| ThreadsNew | Connections | Current Quantity of Threads in NEW State |
| ThreadsRunnable | Connections | Current Quantity of Threads in RUNNABLE State |
| ThreadsBlocked | Connections | Current Quantity of Threads in BLOCKED State |
| ThreadsWaiting | Connections | Current Quantity of Threads in WAITING State |
| ThreadsTimedWaiting | Connections | Current Quantity of Threads in TIMED_WAITING State |
| ThreadsTerminated | Connections | Current Number of Threads in TERMINATED State |
| LogFatal | Threads | Number of FATAL Level Logs |
| LogError | Threads | Number of ERROR Level Logs |
| LogWarn | Threads | Number of WARN Level Logs |
| LogInfo | Threads | Number of INFO Level Logs |
| jetty.numOpenConnections | Connections | Number of Active Jetty Connections |

Kudu Monitoring Metrics

Last updated: 2023-12-21 16:51:37

KUDU-Overview

| Title | Description | Metric Unit | Description |
|-------------------------------|------------------------|-------------|--|
| Tablet Count | TabletRunning | Connections | The total number of tablets currently running in all tablet servers. |
| Tablet Replica Difference | ClusterReplicaSkew | Connections | The discrepancy between the number of replicas on the tablet server bearing the maximum replicas and the tablet server bearing the minimum replicas. |
| TServer Thread Count | ThreadsRunning | Connections | The total number of threads currently running in all tablet servers. |
| Master Thread Count | ThreadsRunning | Connections | The total number of threads currently running in all masters. |
| TServer Log Count | ErrorMessages | Connections | The total number of ERROR level log messages issued in all processes. |
| Master Log Count | ErrorMessages | Connections | The total number of ERROR level log messages issued in all processes. |
| | WarningMessages | Connections | The total number of WARNING level log messages issued in all processes. |
| Excessive Write Request Count | OversizedWriteRequests | Connections | The number of excessive write requests to the system catalog tablet rejected by the master after startup. |

KUDU-Server

| Title | Description | Metric Unit | Description |
|------------------------------|----------------------|-------------|---|
| Block Cache Hits | BlockCacheHit | Threads | The number of anticipated and successful block lookups. This value is used in lieu of cache_hits when determining cache efficiency. |
| | BlockCacheMiss | Threads | The number of anticipated but unsuccessful block lookups. This value is used instead of cache_misses when determining cache efficiency. |
| Block Cache Utilization Rate | BlockCacheUsage | bytes | Memory Occupied by Block Cache |
| File Cache Hits | FileCacheHit | Threads | The number of anticipated and successful file descriptor lookups. This value is used instead of cache_hits when assessing cache efficiency. |
| | FileCacheMiss | Threads | The number of anticipated but unsuccessful file descriptor lookups. This value is used instead of cache_misses when assessing cache efficiency. |
| File Cache Utilization Rate | FileCacheUsage | Connections | Number of Entries in File Cache |
| Scanner | ActiveScanners | Connections | Number of Active Scanners Currently in Operation |
| | ExpiredScanners | Connections | Number of Scanners that have Expired due to Inactivity since the Service was Initiated |
| Number of | BlockUnderManagement | Connections | Number of Data Blocks Currently Managed |

| | | | |
|---|-------------------------------|------------------------|--|
| Blocks Managed by the Block Manager | BlockOpenReading | Connections | Number of Data Blocks Currently Open for Reading |
| | BlockOpenWriting | Connections | Number of Data Blocks Currently Open for Writing |
| Byte Count Managed by the Block Manager | BytesUnderManagement | bytes | Byte Count of Data Blocks Currently Managed |
| Number of Containers Managed by the Block Manager | ContainersUnderManagement | Connections | Number of Log Block Containers |
| | FullContainersUnderManagement | Connections | Number of Complete Log Block Containers |
| Number of Tablet Leaders | NumRaftLeaders | Connections | Quantity of Tablet Replicas for Raft Leaders |
| Number of Tablet Sessions | OpenClientSessions | Connections | Current Number of Open Tablet Replication Client Sessions on this Server |
| | OpenSourceSessions | Connections | Current Number of Open Tablet Replication Source Sessions on this Server |
| Tablet Count | TabletBootstrapping | Connections | Current Number of Tablets Undergoing Bootstrap |
| | TabletFailed | Connections | Number of Failed Tablets |
| | TabletInitialized | Connections | Current Number of Initialized Tablets |
| | TabletNotInitialized | Connections | Current Number of Uninitialized Tablets |
| | TabletRunning | Connections | Current Number of Running Tablets/Current Number of Active Threads |
| | TabletShutdown | Connections | Current Number of Closed Tablets |
| | TabletStopped | Connections | Current Number of Halted Tablets |
| CPU Time | CpuStime | Millisecond (ms) | Total System CPU Time of the Process |
| | CpuUtime | Millisecond (ms) | Total User CPU Time of the Process |
| Data path | DataDirsFailed | Connections | Number of Data Directories Currently in a Fault State on the Disk |
| | DataDirsFull | Connections | Number of Data Directories Currently Filled on the Disk |
| Thread | ThreadsRunning | Connections | Number of Threads Currently Running |
| Context | InvoluntarySwitches | Threads | Involuntary Context Switches |
| | VoluntarySwitches | Threads | Voluntary Context Switches |
| Spinlock | SpinlockContentionTime | Microsecond (μ s) | The amount of time consumed by contention on the internal spinlock since the server was initiated. |
| Log message | ErrorMessages | Connections | The number of ERROR level log messages issued by the application. |
| | WarningMessages | Connections | The number of WARNING level log messages issued by the application. |
| Number of | TotalCount | Connections | Total |

| | | | |
|---------------------------|-------------------------|------------------------|---|
| operations in the queue. | Min | Connections | Minimum number of waiting tasks in the queue. |
| | Max | Connections | Maximum number of waiting tasks in the queue. |
| | Mean | Connections | Average number of waiting tasks in the queue. |
| | Percentile_99_9 | Connections | 99.9th percentile of waiting tasks in the queue. |
| Operation execution time. | TotalCount | Microsecond (μ s) | Total number of operations. |
| | Min | Microsecond (μ s) | Minimum execution time. |
| | Max | Microsecond (μ s) | Maximum execution time. |
| | Mean | Microsecond (μ s) | Average execution time. |
| | Percentile_99_9 | Microsecond (μ s) | 99.9th percentile of execution time. |
| Queue waiting time. | TotalCount | Microsecond (μ s) | Total number of operations. |
| | Min | Microsecond (μ s) | Minimum waiting time. |
| | Max | Microsecond (μ s) | Maximum waiting time. |
| | Mean | Microsecond (μ s) | Average waiting time. |
| | Percentile_99_9 | Microsecond (μ s) | 99.9th percentile of waiting time. |
| Allocated bytes. | AllocatedBytes | bytes | Bytes used by the application. This typically does not match the memory usage reported by the operating system, as it does not include TCMalloc overhead or memory fragmentation. |
| Mixed clock error. | HybridClockError | Microsecond (μ s) | Maximum server clock error; returns $2^{64}-1$ when the underlying clock cannot be read. |
| Mixed clock timestamp. | HybridClockTimestamp | Microsecond (μ s) | Mixed clock timestamp; returns $2^{64}-1$ when the underlying clock cannot be read. |
| TCMalloc memory. | HeapSize | bytes | TCMalloc reserved system memory bytes. |
| | CurrentThreadCacheBytes | bytes | Measurement of memory currently in use by TCMalloc (for small objects). |
| | TotalThreadCacheBytes | bytes | Memory limit of TCMalloc for small objects. |
| TCMalloc PageHeap | FreeBytes | bytes | The number of bytes in mappable pages available in the page heap. |
| | UnMappedBytes | bytes | The number of bytes in unmapped free pages within the page heap. |
| RPC Requests. | ConnectionsAccepted | Connections | The number of incoming TCP connections to the RPC server. |
| | QueueOverflow | Connections | The number of RPCs discarded due to a full service queue. |
| | TimesOutInQueue | Connections | The number of RPCs that timed out and were therefore not processed while waiting in the service queue. |
| RPC FetchData | TotalCount | Microsecond (μ s) | Total number of operations. |

| | | | |
|------------------|-----------------|------------------|---------------------------------------|
| | Min | Microsecond (μs) | Minimum Processing Time. |
| | Max | Microsecond (μs) | Maximum Processing Time. |
| | Mean | Microsecond (μs) | Average Processing Time |
| | Percentile_99_9 | Microsecond (μs) | 99.9th Percentile of Processing Time. |
| RPC AlterSchema | TotalCount | Microsecond (μs) | Total number of operations. |
| | Min | Microsecond (μs) | Minimum Processing Time. |
| | Max | Microsecond (μs) | Maximum Processing Time. |
| | Mean | Microsecond (μs) | Average Processing Time |
| | Percentile_99_9 | Microsecond (μs) | 99.9th Percentile of Processing Time. |
| RPC CreateTablet | TotalCount | Microsecond (μs) | Total number of operations. |
| | Min | Microsecond (μs) | Minimum Processing Time. |
| | Max | Microsecond (μs) | Maximum Processing Time. |
| | Mean | Microsecond (μs) | Average Processing Time |
| | Percentile_99_9 | Microsecond (μs) | 99.9th Percentile of Processing Time. |
| RPC DeleteTablet | TotalCount | Microsecond (μs) | Total number of operations. |
| | Min | Microsecond (μs) | Minimum Processing Time. |
| | Max | Microsecond (μs) | Maximum Processing Time. |
| | Mean | Microsecond (μs) | Average Processing Time |
| | Percentile_99_9 | Microsecond (μs) | 99.9th Percentile of Processing Time. |
| RPC Quiesce | TotalCount | Microsecond (μs) | Total number of operations. |
| | Min | Microsecond (μs) | Minimum Processing Time. |
| | Max | Microsecond (μs) | Maximum Processing Time. |
| | Mean | Microsecond (μs) | Average Processing Time |

| | | | |
|--|-------------------------|------------------|--|
| | Percentile_99_9 | Microsecond (μs) | 99.9th Percentile of Processing Time. |
| RPC Scan | TotalCount | Microsecond (μs) | Total number of operations. |
| | Min | Microsecond (μs) | Minimum Processing Time. |
| | Max | Microsecond (μs) | Maximum Processing Time. |
| | Mean | Microsecond (μs) | Average Processing Time |
| | Percentile_99_9 | Microsecond (μs) | 99.9th Percentile of Processing Time. |
| RPC ScannerKeepAlive | TotalCount | Microsecond (μs) | Total number of operations. |
| | Min | Microsecond (μs) | Minimum Processing Time. |
| | Max | Microsecond (μs) | Maximum Processing Time. |
| | Mean | Microsecond (μs) | Average Processing Time |
| | Percentile_99_9 | Microsecond (μs) | 99.9th Percentile of Processing Time. |
| RPC Write | TotalCount | Microsecond (μs) | Total number of operations. |
| | Min | Microsecond (μs) | Minimum Processing Time. |
| | Max | Microsecond (μs) | Maximum Processing Time. |
| | Mean | Microsecond (μs) | Average Processing Time |
| | Percentile_99_9 | Microsecond (μs) | 99.9th Percentile of Processing Time. |
| Number of Write Refusals Due to Queue Overload | QueueOverloadRejections | count | Number of Write Requests Rejected Due to Queue Overload. |
| Scanner Speed. | ScannedFromDiskRate | bytes/s | Scanner Data Volume Per Second. |
| | ScannerReturnedRate | bytes/s | Scanner Data Volume Per Second. |
| Total Scanner Volume. | ScannedFromDisk | bytes | Total Data Volume Scanned from Hard Disk. |
| | ScannerReturned | bytes | Total Data Volume Returned by Scanner. |
| Total Volume of Row Operations. | RowsInserted | count | Quantity of Rows Inserted by Node. |
| | RowsDeleted | count | Quantity of Rows Deleted by Node. |
| | RowsUpserted | count | Quantity of Rows Upserted by Node. |
| | RowsUpdated | count | Quantity of Rows Updated by Node. |

| | | | |
|------------------------------------|---------------------------------------|------------------|---|
| Rate of Row Operations. | RowsInsertedRate | count/s | Quantity of Rows Inserted by Node per Second. |
| | RowsDeletedRate | count/s | Quantity of Rows Deleted by Node per Second. |
| | RowsUpsertedRate | count/s | Quantity of Rows Upserted by Node per Second. |
| | RowsUpdatedRate | count/s | Quantity of Rows Updated by Node per Second. |
| ExpScanner | ExpiredScanners | keys/second | Average Number of Scanners Expired per Second Due to Inactivity Since Service Launch Within the Metric Collection Period. |
| Mem Rowset | Total | bytes | Size of Tablet's Memrowset Utilized by Node. |
| Memory Refreshment | DeltaMemStore | count | Quantity of DeltaMemStore Refreshments. |
| | MemRowSet | count | Quantity of MemRowSet Refreshments. |
| Disk Rowsets Statistics | Total | count | Total Quantity of Node Tablet Disk Rowsets. |
| | Avg | count | Average Quantity of Node Tablet Disk Rowsets. |
| | Max | count | Maximum Quantity of Node Tablet Disk Rowsets. |
| Size of Tablet Data. | OnDisk | bytes | Size of Tablet Data on Node. |
| Average Height of Disk Rowsets. | Total | count | Average Height of Total Disk Rowsets on Node Tablet. |
| | Avg | count | Average Height of Disk Rowsets on Node Tablet. |
| | Max | count | Maximum Average Height of Disk Rowsets on Node Tablet. |
| Statistics of Compactions Running. | RowSet | count | Total Size of RowSet Mergers on Node Tablet. |
| | Major Delta | count | Total Size of Major Delta Mergers on Node Tablet. |
| | Minor Delta | count | Total Size of Minor Delta Mergers on Node Tablet. |
| Tablet Cache Refresh. | Bytes Flushed | bytes/s | Average Per-Second Tablet Cache Refresh Data Volume on Node within Metric Collection Period. |
| RPC Request Rejection. | leader | keys/second | Average Per-Second RPC Requests Rejected Due to Memory Pressure on Leader within Metric Collection Period. |
| | follower | keys/second | Average Per-Second RPC Requests Rejected Due to Memory Pressure on Follower within Metric Collection Period. |
| Access Queue Duration. | TotalCount,Percentile_99,Min,Max,Mean | Microsecond (μs) | 99th Percentile of Processing Time for Access RPC Requests in the Work Queue. |
| Scanner Duration. | TotalCount,Percentile_99,Min,Max,Mean | Microsecond (μs) | 99th Percentile of Scan Duration. |

KUDU-Master

| Title | Description | Metric Unit | Description |
|------------------------------|-----------------|-------------|--|
| Block Cache Hits | BlockCacheHit | Threads | Number of Times a Block is Expected and Found. Use this Value Instead of Cache Hits when Determining Cache Efficiency. |
| | BlockCacheMiss | Threads | Number of Expected Lookups for Blocks Not Generated. Use this Value to Determine Cache Efficiency, Rather than Cache Misses. |
| Block Cache Utilization Rate | BlockCacheUsage | bytes | Memory Occupied by Block Cache |
| File Cache | FileCacheHit | Threads | The number of anticipated and successful file descriptor lookups. |

| | | | |
|---|-------------------------------|------------------------|---|
| Hits | | | This value is used instead of cache_hits when assessing cache efficiency. |
| | FileCacheMiss | Threads | The number of anticipated but unsuccessful file descriptor lookups. This value is used instead of cache_misses when assessing cache efficiency. |
| File Cache Utilization Rate | FileCacheUsage | Connections | Number of Entries in File Cache |
| Number of Blocks Managed by the Block Manager | BlockUnderManagement | Connections | Number of Data Blocks Currently Managed |
| | BlockOpenReading | Connections | Number of Data Blocks Currently Open for Reading |
| | BlockOpenWriting | Connections | Number of Data Blocks Currently Open for Writing |
| Byte Count Managed by the Block Manager | BytesUnderManagement | bytes | Byte Count of Data Blocks Currently Managed |
| Number of Containers Managed by the Block Manager | ContainersUnderManagement | Connections | Number of Log Block Containers |
| | FullContainersUnderManagement | Connections | Number of Complete Log Block Containers |
| CPU Time | CpuStime | Millisecond (ms) | Total System CPU Time of the Process |
| | CpuUtime | Millisecond (ms) | Total User CPU Time of the Process |
| Thread | ThreadsRunning | Connections | Number of Threads Currently Running |
| Data path | DataDirsFailed | Connections | Number of Data Directories Currently in a Fault State on the Disk |
| | DataDirsFull | Connections | Number of Data Directories Currently Filled on the Disk |
| Allocated bytes. | AllocatedBytes | bytes | Bytes used by the application. This typically does not match the memory usage reported by the operating system, as it does not include TCMalloc overhead or memory fragmentation. |
| Log message | ErrorMessages | Connections | The number of ERROR level log messages issued by the application. |
| | WarningMessages | Connections | The number of WARNING level log messages issued by the application. |
| Context | InvoluntarySwitches | Threads | Involuntary Context Switches |
| | VoluntarySwitches | Threads | Voluntary Context Switches |
| Number of operations in the queue. | TotalCount | Connections | Total |
| | Min | Connections | Minimum number of waiting tasks in the queue. |
| | Max | Connections | Maximum number of waiting tasks in the queue. |
| | Mean | Connections | Average number of waiting tasks in the queue. |
| | Percentile_99_9 | Connections | 99.9th percentile of waiting tasks in the queue. |
| Queue waiting time. | TotalCount | Microsecond (μ s) | Total number of operations. |
| | Min | Microsecond (μ s) | Minimum waiting time. |

| | | | |
|------------------------------------|-------------------------|------------------|--|
| | Max | Microsecond (μs) | Maximum waiting time. |
| | Mean | Microsecond (μs) | Average waiting time. |
| | Percentile_99_9 | Microsecond (μs) | 99.9th percentile of waiting time. |
| Operation execution time. | TotalCount | Microsecond (μs) | Total number of operations. |
| | Min | Microsecond (μs) | Minimum execution time. |
| | Max | Microsecond (μs) | Maximum execution time. |
| | Mean | Microsecond (μs) | Average execution time. |
| | Percentile_99_9 | Microsecond (μs) | 99.9th percentile of execution time. |
| Spinlock | SpinlockContentionTime | Microsecond (μs) | The amount of time consumed by contention on the internal spinlock since the server was initiated. |
| Excessive Number of Read Requests. | OversizedWriteRequests | Connections | Number of Excessive Write Requests to the System Catalog Tablet Rejected After Startup. |
| Mixed clock error. | HybridClockError | Microsecond (μs) | Maximum server clock error; returns $2^{64}-1$ when the underlying clock cannot be read. |
| Mixed clock timestamp. | HybridClockTimestamp | Microsecond (μs) | Mixed clock timestamp; returns $2^{64}-1$ when the underlying clock cannot be read. |
| Tablet Replica Discrepancy. | ClusterReplicaSkew | Connections | The discrepancy between the number of replicas on the tablet server bearing the maximum replicas and the tablet server bearing the minimum replicas. |
| Number of Tablet Leaders | NumRaftLeaders | Connections | Quantity of Tablet Replicas for Raft Leaders |
| Number of Tablet Sessions | OpenSourceSessions | Connections | Current Number of Open Tablet Replication Source Sessions on this Server |
| TCMalloc memory. | HeapSize | bytes | TCMalloc reserved system memory bytes. |
| | CurrentThreadCacheBytes | bytes | Measurement of memory currently in use by TCMalloc (for small objects). |
| | TotalThreadCacheBytes | bytes | Memory limit of TCMalloc for small objects. |
| TCMalloc PageHeap | FreeBytes | bytes | The number of bytes in mappable pages available in the page heap. |
| | UnMappedBytes | bytes | The number of bytes in unmapped free pages within the page heap. |
| RPC Requests. | ConnectionsAccepted | Connections | The number of incoming TCP connections to the RPC server. |
| | QueueOverflow | Connections | The number of RPCs discarded due to a full service queue. |
| | TimesOutInQueue | Connections | The number of RPCs that timed out and were therefore not processed while waiting in the service queue. |

| | | | |
|-----------------------|-----------------|------------------|---------------------------------------|
| RPC RunLeaderElection | TotalCount | Microsecond (μs) | Total number of operations. |
| | Min | Microsecond (μs) | Minimum Processing Time. |
| | Max | Microsecond (μs) | Maximum Processing Time. |
| | Mean | Microsecond (μs) | Average Processing Time |
| | Percentile_99_9 | Microsecond (μs) | 99.9th Percentile of Processing Time. |
| RPC ConnectToMaster | TotalCount | Microsecond (μs) | Total number of operations. |
| | Min | Microsecond (μs) | Minimum Processing Time. |
| | Max | Microsecond (μs) | Maximum Processing Time. |
| | Mean | Microsecond (μs) | Average Processing Time |
| | Percentile_99_9 | Microsecond (μs) | 99.9th Percentile of Processing Time. |
| RPC Ping | TotalCount | Microsecond (μs) | Total number of operations. |
| | Min | Microsecond (μs) | Minimum Processing Time. |
| | Max | Microsecond (μs) | Maximum Processing Time. |
| | Mean | Microsecond (μs) | Average Processing Time |
| | Percentile_99_9 | Microsecond (μs) | 99.9th Percentile of Processing Time. |
| RPC TSHearbeat | TotalCount | Microsecond (μs) | Total number of operations. |
| | Min | Microsecond (μs) | Minimum Processing Time. |
| | Max | Microsecond (μs) | Maximum Processing Time. |
| | Mean | Microsecond (μs) | Average Processing Time |
| | Percentile_99_9 | Microsecond (μs) | 99.9th Percentile of Processing Time. |
| RPC FetchData | TotalCount | Microsecond (μs) | Total number of operations. |
| | Min | Microsecond (μs) | Minimum Processing Time. |
| | Max | Microsecond (μs) | Maximum Processing Time. |
| | Mean | Microsecond | Average Processing Time |

| | | | |
|--|-----------------|---------------------------|---------------------------------------|
| | | (μ s) | |
| | Percentile_99_9 | Microsecond (μ s) | 99.9th Percentile of Processing Time. |

Alluxio Monitoring Metrics

Last updated: 2023-12-21 16:51:44

ALLUXIO-Cluster

| Title | Description | Metric Unit | Description |
|--------------------------------------|-------------------------------|-------------|---|
| Total Data Read/Write Volume | BytesReadAlluxio | Bytes | Total bytes read from Alluxio storage as reported by all workers. |
| | BytesReadUfsAll | Bytes | Total bytes read from all Alluxio UFSes by all workers. |
| | BytesWrittenAlluxio | Bytes | Total bytes written to Alluxio storage across all workers. |
| | BytesWrittenUfsAll | Bytes | Total bytes written to all Alluxio UFSes by all workers. |
| Data Read/Write Throughput | BytesReadAlluxioThroughput | Bytes | Throughput of data read from Alluxio storage by all workers. |
| | BytesReadUfsThroughput | Bytes | Throughput of data read from all Alluxio UFSes by all workers. |
| | BytesWrittenAlluxioThroughput | Bytes | Throughput of data written to Alluxio storage by all workers. |
| | BytesWrittenUfsThroughput | Bytes | Throughput of data written to all Alluxio UFSes by all workers. |
| Capacity at the worker's tier level. | CapacityFree | Bytes | Total available bytes across all layers of all workers. |
| | CapacityTotal | Bytes | Total capacity across all layers of all workers. |
| | CapacityUsed | Bytes | Total number of bytes used across all layers of all workers. |
| Total number of workers. | Workers | Connections | Total number of active workers in the cluster. |

ALLUXIO-Master

| Title | Description | Metric Unit | Description |
|----------------------------|--------------------|-------------|---|
| CompleteFile operation. | CompleteFileOps | Connections | Total Number of CompleteFile Operations |
| | FilesCompleted | Connections | Total Number of Successful CompleteFile Operations |
| CreateDirectory operation. | CreateDirectoryOps | Connections | Total Number of CreateDirectory Operations |
| | DirectoriesCreated | Connections | Total Number of Successful CreateDirectory Operations |
| CreateFile Operation | CreateFileOps | Connections | Total Number of CreateFile Operations |
| | FilesCreated | Connections | Total Number of Successful CreateFile Operations |
| Delete Operation | DeletePathOps | Connections | Total Number of Delete Operations |
| | PathsDeleted | Connections | Total Number of Successful Delete Operations |

| | | | |
|----------------------------------|---------------------|-------------|--|
| FreeFile Operation | FreeFileOps | Connections | Total Number of FreeFile Operations |
| | FilesFreed | Connections | Total Number of Successful FreeFile Operations |
| GetFileBlockInfo Operation | GetFileBlockInfoOps | Connections | Total Number of GetFileBlockInfo Operations |
| | FileBlockInfosGot | Connections | Total Number of Successful GetFileBlockInfo Operations |
| GetFileInfo Operation | GetFileInfoOps | Connections | Total Number of GetFileInfo Operations |
| | FileInfosGot | Connections | Total Number of Successful GetFileInfo Operations |
| GetNewBlock Operation | GetNewBlockOps | Connections | Total Number of GetNewBlock Operations |
| | NewBlocksGot | Connections | Total Number of Successful GetNewBlock Operations |
| Mount Operation | MountOps | Connections | Total Number of Mount Operations |
| | PathsMounted | Connections | Total Number of Successful Mount Operations |
| Unmount Operation | UnmountOps | Connections | Total Number of Unmount Operations |
| | PathsUnmounted | Connections | Total Number of Successful Unmount Operations |
| Rename Operation | RenamePathOps | Connections | Total Number of Rename Operations |
| | PathsRenamed | Connections | Total Number of Successful Rename Operations |
| SetAcl Operation | SetAclOps | Connections | Total Number of SetAcl Operations |
| SetAttribute Operation | SetAttributeOps | Connections | Total Number of SetAttribute Operations |
| Total Number of File Operations | FilesPersisted | Connections | Total Number of Successfully Saved Files |
| | FilesPinned | Connections | Current Total Number of Pinned Files |
| Total Number of File Directories | TotalPaths | Connections | Total Number of Files and Directories in Alluxio Namespace |
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |

| | | | |
|------------|----------------------|----|---|
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| JVM Memory | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| | MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |

ALLUXIO-Worker

| Title | Description | Metric Unit | Description |
|---------------------------------------|-----------------------------|-------------|--|
| Asynchronous Cache Request | AsyncCacheDuplicateRequests | Connections | Total Number of Duplicate Async Cache Requests Received by Worker |
| | AsyncCacheRequests | Connections | Total Number of Async Cache Requests Received by Worker |
| Quantity of Asynchronous Cache Blocks | AsyncCacheFailedBlocks | Connections | Total Number of Blocks Failed in Worker Async Cache |
| | AsyncCacheRemoteBlocks | Connections | Total Number of Blocks Requiring Asynchronous Caching from Remote Source |
| | AsyncCacheSucceededBlocks | Connections | Total Number of Blocks Successfully Cached by Worker Async |
| | AsyncCacheUfsBlocks | Connections | Total Number of Data Blocks Requiring Async Cache from Local Source |
| Blocks | BlocksAccessed | Connections | Total Number of Times Any Block within this Worker has been Accessed |
| | BlocksCached | Connections | Total Number of Blocks Used for Caching Data in the Worker |
| | BlocksCancelled | Connections | Total Number of Temporary Blocks Aborted in the Worker |
| | BlocksDeleted | Connections | Total Number of Blocks in this Worker Listed as Deleted by External Requests |
| | BlocksEvicted | Connections | Total Number of Blocks Evicted in the Worker |
| | BlocksLost | Connections | Total Number of Data Blocks Lost in the Worker |
| | BlocksPromoted | Connections | Total Number of Times Any Data Block in the Worker Has Been Moved to a New Layer |
| Capacity at the worker's tier level. | CapacityFree | Bytes | Total Available Bytes Across All Layers of the Worker |

| | | | |
|------------------------------|----------------------|---------|--|
| | CapacityTotal | Bytes | Total Capacity Across All Layers of the Worker |
| | CapacityUsed | Bytes | Total Number of Bytes Used Across All Layers of the Worker |
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| JVM Memory | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| | MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |

PrestoSQL Monitoring Metrics

Last updated: 2023-12-21 16:51:52

Note

PrestoSQL metrics are currently only supported for PrestoSQL322 and versions PrestoSQL350 and above.

PrestoSQL – Overview

| Title | Description | Metric Unit | Description |
|--|-------------------------------|-------------|--|
| Node Quantity | Active | Connections | Number of Active Nodes |
| | Total | Connections | Total Node Count |
| | Failed | Connections | Number of Failed Nodes |
| Query | RunningQueries | Connections | Total Number of Running Queries |
| | QueuedQueries | Connections | Total Number of Queries in Waiting Status |
| Query Frequency | FailedQueries | Per Minute | Total Number of Failed Queries |
| | AbandonedQueries | Per Minute | Total Number of Abandoned Queries |
| | CanceledQueries | Per Minute | Total Number of Canceled Queries |
| | CompletedQueries | Per Minute | Total Number of Completed Queries |
| | StartedQueries | Per Minute | Total Number of Initiated Queries |
| Data Input/Output Volume Per Minute | InputDataSizeOneMinute | GB/min | Input Data Rate |
| | OutputDataSizeOneMinute | GB/min | Output Data Rate |
| Number of User-Induced Failed Queries Per Minute | UserErrorFailures | count/min | Number of User-Induced Failed Queries Per Minute |
| Number of Internal Error-Induced Failed Queries Per Minute | InternalFailures | count/min | Number of Internal Error-Induced Failed Queries Per Minute |
| Number of Resource Insufficiency-Induced Failed Queries Per Minute | InsufficientResourcesFailures | count/min | Number of Resource Insufficiency-Induced Failed Queries Per Minute |
| Number of External Error-Induced Failed Queries Per Minute | ExternalFailures | count/min | Number of External Error-Induced Failed Queries Per Minute |
| Number of Records Processed Per Minute by Queries | ConsumedInputRows | count/min | Number of Records Processed Per Minute by Queries |
| Volume of Bytes Processed Per Minute by Queries | ConsumedInput | bytes/min | Volume of Bytes Processed Per Minute by Queries |
| CPU Time Utilized for Query Processing Per Minute | ConsumedCpuTimeSecs | count/min | CPU Time Utilized for Query Processing Per Minute |
| Number of Splits | BlockedSplits | count | Number of Blocked Splits |
| | RunningSplits | count | Number of Active Splits |

| | | | |
|---|--------------------------|-----------|---|
| | TotalSplits | count | Total Number of Splits for Task Executors |
| | WaitingSplits | count | Number of Splits Awaiting Execution on TaskExecutor |
| All Tasks Registered with TaskExecutor | Tasks | count | All Tasks Registered with TaskExecutor |
| Number of Tasks Failing per Minute | FailedTasks | count/min | Number of Tasks Failing per Minute |
| Number of Data Rows Input and Output per Minute by Task | InputPositionsOneMinute | count/min | Number of Data Rows Input by Task per Minute |
| | OutputPositionsOneMinute | count/min | Number of Data Rows Output by Task per Minute |

PrestoSQL-Worker

| Title | Description | Metric Unit | Description |
|------------------------------|-------------------------------|-------------|--|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| Heap Memory Utilization Rate | MemHeapUsedRate | % | Percentage of Currently Utilized HeapMemory in JVM Relative to the Configured HeapMemory Quantity in JVM |
| JVM Memory | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| | MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| Data Input-Output Rate | InputDataSize.OneMinute.Rate | GB/min | Input Data Rate |
| | OutputDataSize.OneMinute.Rate | GB/min | Output Data Rate |

| | | | |
|--------------------------|-------------------------|-------------|-----------------------------------|
| Processes | PeakThreadCount | Connections | Peak Thread Count |
| | ThreadCount | Connections | Total Thread Count |
| | DaemonThreadCount | Connections | Daemon Thread Quantity |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| Process Initiation Time | StartTime | s | Process Initiation Time |
| File Descriptor Count | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| | OpenFileDescriptorCount | Connections | Quantity of Open File Descriptors |

PRESTOSQL-Coordinator

| Title | Description | Metric Unit | Description |
|------------------------------|----------------------|-------------|--|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| Heap Memory Utilization Rate | MemHeapUsedRate | % | Percentage of Currently Utilized HeapMemory in JVM Relative to the Configured HeapMemory Quantity in JVM |
| JVM Memory | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| | MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| Processes | PeakThreadCount | Connections | Peak Thread Count |
| | ThreadCount | Connections | Number of Threads |
| | DaemonThreadCount | Connections | Background Thread Count |

| | | | |
|--------------------------|-------------------------|-------------|-----------------------------------|
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| Process Initiation Time | StartTime | s | Process Initiation Time |
| File Descriptor Count | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| | OpenFileDescriptorCount | Connections | Quantity of Open File Descriptors |

Impala Monitoring Metrics

Last updated: 2023-12-21 16:51:59

Note

Impala metrics are currently only supported for Impala 3.4.0 and higher versions.

Impala-CATALOG

| Title | Description | Metric Unit | Description |
|--|--------------------------------|-------------|--|
| Resident Memory Set | RSS | bytes | Resident Memory Set |
| JVM Memory | MemHeapInitM | MB | Peak Quantity of Initial JVM Heap Memory |
| | MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemNonHeapInitM | MB | Initial Quantity of NonHeapMemory in JVM |
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| Heap Memory Utilization Rate | MemHeapUsedRate | % | Percentage of Currently Utilized HeapMemory in JVM Relative to the Configured HeapMemory Quantity in JVM |
| Heartbeat Interval from Daemon Process to StateStore | Last | s | Most Recent Heartbeat Interval from Daemon Process to StateStore |
| | Max | s | Maximum Heartbeat Interval from Daemon Process to StateStore |
| | Mean | s | Average Heartbeat Interval from Daemon Process to StateStore |
| | Min | s | Minimum Heartbeat Interval from Daemon Process to StateStore |
| | Stddev | s | Standard Deviation of Heartbeat Intervals from Daemon Process to StateStore |
| TCMalloc memory. | Used | bytes | Bytes Utilized by the Program |
| | PageheapFreeBytes | bytes | Bytes of Free Mapped Pages in the Page Heap |
| | PageheapUnmappedBytes | bytes | Bytes of Unmapped Free Pages in the Page Heap |
| | PhysicalBytesReserved | bytes | Amount of Physical Memory Used by the Computing Process |
| | TotalBytesReserved | bytes | Bytes of System Memory Reserved by TCMalloc |
| Number of active connections | Thrift_Server_Connections_Used | Connections | Number of active connections |

| | | | |
|--------------------------|-------------------------|-------------|----------------------------------|
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| File Descriptor Count | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| | OpenFileDescriptorCount | Connections | Number of Open File Descriptors |
| Threads | ThreadCount | Connections | Total Thread Count |
| | DaemonThreadCount | Connections | Daemon Thread Count |
| CPU utilization | SystemCpuLoad | Connections | System CPU Utilization |

Impala-STATESTORE

| Title | Description | Metric Unit | Description |
|-----------------------------|-----------------------|-------------|---|
| Resident Memory Set | RSS | bytes | Resident Memory Set |
| TCMalloc memory. | Used | bytes | Bytes Utilized by the Program |
| | PageheapFreeBytes | bytes | Bytes of Free Mapped Pages in the Page Heap |
| | PageheapUnmappedBytes | bytes | Bytes of Unmapped Free Pages in the Page Heap |
| | PhysicalBytesReserved | bytes | Amount of Physical Memory Used by the Computing Process |
| | TotalBytesReserved | bytes | Bytes of System Memory Reserved by TCMalloc |
| Number of connections | Used | Connections | Number of active connections |
| Running Threads | Count | Connections | Running Threads |
| StateStore Subscriber Count | Count | Connections | StateStore Subscriber Count |

Impala-DAEMON

| Title | Description | Metric Unit | Description |
|------------------------------|----------------------|-------------|--|
| JVM Memory | MemHeapInitM | MB | Peak Quantity of Initial JVM Heap Memory |
| | MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemNonHeapInitM | MB | Initial Quantity of NonHeapMemory in JVM |
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| Heap Memory Utilization Rate | MemHeapUsedRate | % | Percentage of Currently Utilized HeapMemory in JVM Relative to the Configured HeapMemory Quantity in JVM |

| | | | |
|--|-------------------------|-------------|---|
| TCMalloc memory. | Used | bytes | Bytes Utilized by the Program |
| | PageheapFreeBytes | bytes | Bytes of Free Mapped Pages in the Page Heap |
| | PageheapUnmappedBytes | bytes | Bytes of Unmapped Free Pages in the Page Heap |
| | PhysicalBytesReserved | bytes | Amount of Physical Memory Used by the Computing Process |
| | TotalBytesReserved | bytes | Bytes of System Memory Reserved by TCMalloc |
| Threads | ThreadCount | Connections | Total Thread Count |
| | DaemonThreadCount | Connections | Daemon Thread Count |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| File Descriptor Count | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| | OpenFileDescriptorCount | Connections | Number of Open File Descriptors |
| CPU utilization | SystemCpuLoad | Connections | System CPU Utilization |
| Beeswax API Client Connections | Use | Connections | Active Beeswax API Connections |
| | Conn_In_Use | Connections | Active Beeswax API Connections with this Impala Daemon |
| | TotalConns | Connections | Total Active Beeswax API Connections with this Impala Daemon |
| | ConnSetupQueueSize | Connections | Beeswax API Connections Received and Awaiting Establishment by this Impala Daemon |
| HS2 API Client Connections | Use | Connections | Active HS2 API Connections |
| | Conn_In_Use | Connections | Active HS2 API Connections |
| | TotalConns | Connections | Total HS2 API Connections Established Throughout the Lifespan of this Impala Daemon |
| | ConnSetupQueueSize | Connections | HS2 API Connections Received and Awaiting Establishment by this Impala Daemon |
| Thread Manager | RunningThreads | Connections | Running Threads |
| | TotalCreatedThreads | Connections | Number of Threads Created Throughout the Lifespan |
| Memory Manager Constraints | Limit | Bytes | Excess Memory Quantity Beyond Its Limit (Default Value -1) |
| Excess Memory Quantity Beyond Its Limit (Default Value -1) | OverLimit | Bytes | Number of Threads Created Throughout the Lifespan |
| HS2 API Client Waiting Time for Connection Establishment | P20 | us | HS2 API Client Waiting Time for Connection Establishment |
| | P50 | | |
| | P70 | | |
| | P90 | | |
| | P95 | | |
| | P99.9 | | |
| Beeswax API Client Waiting | P20 | us | Beeswax API Client Waiting Time for Service |

| | | | |
|---|---------------------|-------------|---|
| Time for Service Thread Establishment | P50 | | Thread Establishment |
| | P70 | | |
| | P90 | | |
| | P95 | | |
| | P99.9 | | |
| Number of Beeswax API Connections That Have Timed Out | TimeOutCnncRequests | Connections | Number of Beeswax API Connections That Have Timed Out |
| Time Spent on Parsing Request Pool Requests (Milliseconds) | Total | ms | Time Spent on Parsing Request Pool Requests (Milliseconds) |
| Number of Cache Misses in External Data Source Cache Class | Misses | Connections | Number of Cache Misses in External Data Source Cache Class |
| Number of Connection Requests to the Configured Impala Backend Server That Have Timed Out | ConnSetupQueueSize | Connections | Number of Connection Requests to the Configured Impala Backend Server That Have Timed Out |
| Number of Connection Requests to the Configured Impala Backend That Have Timed Out | TimeOutCnncRequests | Connections | Number of Connection Requests to the Configured Impala Backend That Have Timed Out |
| Total Number of Impala Backend Client Connections Established by This Impala Daemon | TotalConnections | Connections | Total Number of Impala Backend Client Connections Established by This Impala Daemon |
| Time Spent by Impala Backend's Client Waiting for Connection Establishment | P20 | us | Time Spent by Impala Backend's Client Waiting for Connection Establishment |
| | P50 | | |
| | P70 | | |
| | P90 | | |
| | P95 | | |
| | P99.9 | | |
| | Count | | |
| | Sum | | |
| Time Spent by Impala Backend's Client Waiting for Service Threads | P20 | us | Time Spent by Impala Backend's Client Waiting for Service Threads |
| | P50 | | |
| | P70 | | |
| | P90 | | |
| | P95 | | |
| | P99.9 | | |
| | Count | | |
| | Sum | | |
| Time Spent by HS2 API Client | P20 | us | Time Spent by HS2 API Client Waiting for |

| | | | |
|--|-------------------|-------------|--|
| Waiting for Service Thread Establishment | P50 | | Service Thread Establishment |
| | P70 | | |
| | P90 | | |
| | P95 | | |
| | P99.9 | | |
| | Count | | |
| | Sum | | |
| Time Spent by HS2 HTTP API Client Waiting for Service Thread | P20 | us | Time Spent by HS2 HTTP API Client Waiting for Service Thread |
| | P50 | | |
| | P70 | | |
| | P90 | | |
| | P95 | | |
| | P99.9 | | |
| | Count | | |
| Sum | | | |
| DataStreamService: Number of Service Queue Overflows Rejected | RpcsQueueOverflow | Connections | DataStreamService: Number of Service Queue Overflows Rejected |
| ControlStreamService: Number of Service Queue Overflows Rejected | RpcsQueueOverflow | Connections | ControlStreamService: Number of Service Queue Overflows Rejected |
| DataStreamService: Number of Bytes Utilized | PeakUsageBytes | Bytes | Peak Number of Bytes Utilized by Memtracker DataStreamService |
| | CurrentUsageBytes | Bytes | Current Number of Bytes Utilized by Memtracker DataStreamService |
| ControlService: Number of Bytes Utilized | PeakUsageBytes | Bytes | Peak Number of Bytes Utilized by Memtracker ControlService |
| | CurrentUsageBytes | Bytes | Current Number of Bytes Utilized by Memtracker ControlService |
| Resident Set Size (RSS) of This Process | RSS | Bytes | Resident Set Size (RSS) of This Process |
| Total Number of Backends Registered in StateStore | Total | Connections | Total Number of Backends Registered in StateStore |
| Distribution of Query Latency | P20 | us | Distribution of Query Latency |
| | P50 | | |
| | P70 | | |
| | P90 | | |
| | P95 | | |
| | P99.9 | | |
| | Count | | |
| Sum | | | |

| | | | |
|---|-----------------------------|-------------|---|
| Number of Open HDFS Files for Writing | NumFilesOpenForInsert | Connections | Number of Open HDFS Files |
| Scan Ranges Read During the Process Lifetime | ScanRangesTotal | Connections | Scan Ranges Read During the Process Lifetime |
| Number of Open Beeswax Sessions | NumOpenBeeswaxSessions | Connections | Number of Open Beeswax Sessions |
| Total Number of Query Fragments Processed During the Process Lifetime | NumFragments | Connections | Total Number of Query Fragments Processed During the Process Lifetime |
| Total Number of Scan Ranges Read During the Process Lifetime Without Volum Metadata | ScanRangesNumMissingVolumId | Connections | Total Number of Scan Ranges Read During the Process Lifetime Without Volum Metadata |
| Number of Hedged Reads Attempts | HedgedReadOps | Connections | Number of Hedged Reads Attempts |
| Total Number of Queries Processed During the Process Lifetime | NumQueries | Connections | Total Number of Queries Processed During the Process Lifetime |
| Total Number of Rows Supporting Cache HS2 FETCH_FIRST | ResultSetCacheTotalNumRows | Connections | Total Number of Rows for HS2 FETCH_FIRST Supported by Cache |
| Total Number of Queries Registered on this Impala Server | NumQueriesRegistered | Connections | Total Number of Queries Registered on this Impala Server |
| Total Number of BE Queries | NumQueriesExecuted | Connections | Total Number of BE Queries |
| Number of Sessions Terminated Due to Inactivity | NumSessionsExpired | Connections | Number of Sessions Terminated Due to Inactivity |
| Number of Queries Terminated Due to Inactivity | NumQueriesExpired | Connections | Number of Queries Terminated Due to Inactivity |
| Number of Opened HS2 Sessions | NumOpenHS2Sessions | Connections | Number of Opened HS2 Sessions |
| Number of Tables in the Catalog | NumTables | Connections | Quantity of Catalog Tables |
| Number of Databases in the Catalog | NumDatabases | Connections | Quantity of Catalog Databases |
| Bytes Written to Disk by the IO Manager | BytesWritten | Connections | Bytes Written to Disk by the IO Manager |
| Number of Files Opened by the IO Manager | NumOpenFiles | Connections | Number of Files Opened by the IO Manager |
| Number of HDFS File Handles Utilized | NumFileHandlesOutstanding | Bytes | Number of HDFS File Handles Utilized |
| Number of Local Bytes Read | LocalBytesRead | Bytes | Number of Local Bytes Read by the IO Manager |

Ranger Monitoring Metrics

Last updated: 2023-12-21 16:52:05

Ranger-Admin

| Title | Description | Metric Unit | Description |
|------------------------------|-------------------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| JVM Memory | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of HeapMemory Committed by JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| | MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| CPU Utilization Rate | ProcessCpuLoad | % | CPU Utilization Rate |
| Number of File Handles | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| | OpenFileDescriptorCount | Connections | Number of Open File Descriptors |
| Accumulated CPU Usage Time | ProcessCpuTime | ms | Accumulated CPU Usage Time |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| Threads | ThreadCount | Connections | Number of Threads |
| | PeakThreadCount | Connections | Peak Thread Count |
| | DaemonThreadCount | Connections | Background Thread Count |

Ranger-UserSync

| Title | Description | Metric Unit | Description |
|------------------------------|-------------------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| JVM Memory | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of HeapMemory Committed by JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| | MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| CPU Utilization Rate | ProcessCpuLoad | % | CPU Utilization Rate |
| Number of File Handles | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| | OpenFileDescriptorCount | Connections | Number of Open File Descriptors |
| Accumulated CPU Usage Time | ProcessCpuTime | ms | Accumulated CPU Usage Time |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| Threads | ThreadCount | Connections | Number of Threads |
| | PeakThreadCount | Connections | Peak Thread Count |
| | DaemonThreadCount | Connections | Background Thread Count |

COSRanger Monitoring Metrics

Last updated: 2023-12-21 16:52:14

Note

Metrics such as authentication failure statistics, unauthenticated statistics, and authenticated statistics are currently without data. The feature is in the process of being launched. Your patience is appreciated.

COSranger-CosRangerServer

| Title | Description | Metric Unit | Description |
|------------------------------|-------------------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| JVM Memory | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of HeapMemory Committed by JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| | MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| CPU Utilization Rate | ProcessCpuLoad | % | CPU Utilization Rate |
| Number of File Handles | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| | OpenFileDescriptorCount | Connections | Number of Open File Descriptors |

| | | | |
|--------------------------------------|----------------------|---------------------|---|
| CPU Usage Duration | ProcessCpuTime | ms | Accumulated CPU Usage Time |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| Number of Working Threads | ThreadCount | Connections | Number of Threads |
| | PeakThreadCount | Connections | Peak Thread Count |
| | DaemonThreadCount | Connections | Background Thread Count |
| - | Leader | - | Is it the Primary Node for COSRanger? |
| Check Statistics | PermissionAllowCount | Count (Occurrences) | Total Number of Successful Authorizations |
| | AuthDenyCnt | Count (Occurrences) | Total Number of Authentication Failures |
| | PermissionDenyCount | Count (Occurrences) | Total Number of Authorization Failures |
| Successful Authentication Statistics | Qps | Count (Occurrences) | Queries Per Second |
| | Total_5m | Count (Occurrences) | Total Request Count for Five Minutes |
| | Total_1m | Count (Occurrences) | Total Request Count for One Minute |
| | Qps_5m | Count (Occurrences) | Five-Minute Average Request Count |
| | Qps_1m | Count (Occurrences) | One-Minute Average Request Count |
| Authentication Failure Statistics | Qps | Count (Occurrences) | Queries Per Second |
| | Total_5m | Count (Occurrences) | Total Request Count for Five Minutes |
| | Total_1m | Count (Occurrences) | Total Request Count for One Minute |
| | Qps_5m | Count (Occurrences) | Five-Minute Average Request Count |
| | Qps_1m | Count (Occurrences) | One-Minute Average Request Count |
| Unauthorized Access Statistics | Qps | Count (Occurrences) | Queries Per Second |
| | Total_5m | Count (Occurrences) | Total Request Count for Five Minutes |
| | Total_1m | Count (Occurrences) | Total Request Count for One Minute |
| | Qps_5m | Count (Occurrences) | Five-Minute Average Request Count |
| | Qps_1m | Count (Occurrences) | One-Minute Average Request Count |
| Authorized Access Statistics | Qps | Count (Occurrences) | Queries Per Second |
| | | | |

| | | | |
|--|----------|---------------------|--------------------------------------|
| | Total_5m | Count (Occurrences) | Total Request Count for Five Minutes |
| | Total_1m | Count (Occurrences) | Total Request Count for One Minute |
| | Qps_5m | Count (Occurrences) | Five-Minute Average Request Count |
| | Qps_1m | Count (Occurrences) | One-Minute Average Request Count |
| AccessStat_DELETE Operation Statistics | Qps | Count (Occurrences) | Queries Per Second |
| | Total_5m | Count (Occurrences) | Total Request Count for Five Minutes |
| | Total_1m | Count (Occurrences) | Total Request Count for One Minute |
| | Qps_5m | Count (Occurrences) | Five-Minute Average Request Count |
| | Qps_1m | Count (Occurrences) | One-Minute Average Request Count |
| AccessStat_LIST Operation Statistics | Qps | Count (Occurrences) | Queries Per Second |
| | Total_5m | Count (Occurrences) | Total Request Count for Five Minutes |
| | Total_1m | Count (Occurrences) | Total Request Count for One Minute |
| | Qps_5m | Count (Occurrences) | Five-Minute Average Request Count |
| | Qps_1m | Count (Occurrences) | One-Minute Average Request Count |
| AccessStat_READ Operation Statistics | Qps | Count (Occurrences) | Queries Per Second |
| | Total_5m | Count (Occurrences) | Total Request Count for Five Minutes |
| | Total_1m | Count (Occurrences) | Total Request Count for One Minute |
| | Qps_5m | Count (Occurrences) | Five-Minute Average Request Count |
| | Qps_1m | Count (Occurrences) | One-Minute Average Request Count |
| AccessStat_WRITE Operation Statistics | Qps | Count (Occurrences) | Queries Per Second |
| | Total_5m | Count (Occurrences) | Total Request Count for Five Minutes |
| | Total_1m | Count (Occurrences) | Total Request Count for One Minute |
| | Qps_5m | Count (Occurrences) | Five-Minute Average Request Count |
| | Qps_1m | Count | One-Minute Average Request Count |

| | | (Occurrences) | |
|---|----------|---------------------|--------------------------------------|
| Rpc_getRangerAuthPolicy Invocation Statistics | Qps | Count (Occurrences) | Queries Per Second |
| | Total_5m | Count (Occurrences) | Total Request Count for Five Minutes |
| | Total_1m | Count (Occurrences) | Total Request Count for One Minute |
| | Qps_5m | Count (Occurrences) | Five-Minute Average Request Count |
| | Qps_1m | Count (Occurrences) | One-Minute Average Request Count |
| Rpc_checkPermission Invocation Statistics | Qps | Count (Occurrences) | Queries Per Second |
| | Total_5m | Count (Occurrences) | Total Request Count for Five Minutes |
| | Total_1m | Count (Occurrences) | Total Request Count for One Minute |
| | Qps_5m | Count (Occurrences) | Five-Minute Average Request Count |
| | Qps_1m | Count (Occurrences) | One-Minute Average Request Count |
| Rpc_getDelegationToken Invocation Statistics | Qps | Count (Occurrences) | Queries Per Second |
| | Total_5m | Count (Occurrences) | Total Request Count for Five Minutes |
| | Total_1m | Count (Occurrences) | Total Request Count for One Minute |
| | Qps_5m | Count (Occurrences) | Five-Minute Average Request Count |
| | Qps_1m | Count (Occurrences) | One-Minute Average Request Count |
| Rpc_renewDelegationToken Invocation Statistics | Qps | Count (Occurrences) | Queries Per Second |
| | Total_5m | Count (Occurrences) | Total Request Count for Five Minutes |
| | Total_1m | Count (Occurrences) | Total Request Count for One Minute |
| | Qps_5m | Count (Occurrences) | Five-Minute Average Request Count |
| | Qps_1m | Count (Occurrences) | One-Minute Average Request Count |
| Rpc_cancelDelegationToken Invocation Statistics | Qps | Count (Occurrences) | Queries Per Second |
| | Total_5m | Count (Occurrences) | Total Request Count for Five Minutes |
| | Total_1m | Count (Occurrences) | Total Request Count for One Minute |

| | | | |
|---|-------------|------------------------|--|
| | Qps_5m | Count (Occurrences) | Five-Minute Average Request Count |
| | Qps_1m | Count (Occurrences) | One-Minute Average Request Count |
| Rpc_getSTS Invocation Statistics | Qps | Count (Occurrences) | Queries Per Second |
| | Total_5m | Count (Occurrences) | Total Request Count for Five Minutes |
| | Total_1m | Count (Occurrences) | Total Request Count for One Minute |
| | Qps_5m | Count (Occurrences) | Five-Minute Average Request Count |
| | Qps_1m | Count (Occurrences) | One-Minute Average Request Count |
| CosRpc_getSTS Invocation Duration | Cost_Avg | μ s (Microseconds) | Average Time Consumption within the Current Second |
| | Cost_Avg_1m | μ s (Microseconds) | One-Minute Average Time Consumption |
| | Cost_Avg_5m | μ s (Microseconds) | Five-Minute Average Time Consumption |
| | Cost_Max | μ s (Microseconds) | Maximum Time Consumption within the Current Second |
| | Cost_Max_1m | μ s (Microseconds) | Maximum Time Consumption within One Minute |
| | Cost_Max_5m | μ s (Microseconds) | Maximum Time Consumption within Five Minutes |
| | Cost_Min | μ s (Microseconds) | Minimum Time Consumption within the Current Second |
| | Cost_Min_1m | μ s (Microseconds) | Minimum Time Consumption within One Minute |
| | Cost_Min_5m | μ s (Microseconds) | Minimum Time Consumption within Five Minutes |
| Time Consumption for cosRpc_renewDelegationToken Invocation | Cost_Avg | μ s (Microseconds) | Average Time Consumption within the Current Second |
| | Cost_Avg_1m | μ s (Microseconds) | One-Minute Average Time Consumption |
| | Cost_Avg_5m | μ s (Microseconds) | Five-Minute Average Time Consumption |
| | Cost_Max | μ s (Microseconds) | Maximum Time Consumption within the Current Second |
| | Cost_Max_1m | μ s (Microseconds) | Maximum Time Consumption within One Minute |
| | Cost_Max_5m | μ s (Microseconds) | Maximum Time Consumption within Five Minutes |
| | Cost_Min | μ s (Microseconds) | Minimum Time Consumption within the Current Second |
| | | | |

| | | | |
|--|-------------|----------------------|--|
| | Cost_Min_1m | μs (Microseconds) | Minimum Time Consumption within One Minute |
| | Cost_Min_5m | μs (Microseconds) | Minimum Time Consumption within Five Minutes |
| Time Consumption for cosRpc_cancelDelegationToken Invocation | Cost_Avg | μs (Microseconds) | Average Time Consumption within the Current Second |
| | Cost_Avg_1m | μs (Microseconds) | One-Minute Average Time Consumption |
| | Cost_Avg_5m | μs (Microseconds) | Five-Minute Average Time Consumption |
| | Cost_Max | μs (Microseconds) | Maximum Time Consumption within the Current Second |
| | Cost_Max_1m | μs (Microseconds) | Maximum Time Consumption within One Minute |
| | Cost_Max_5m | μs (Microseconds) | Maximum Time Consumption within Five Minutes |
| | Cost_Min | μs (Microseconds) | Minimum Time Consumption within the Current Second |
| | Cost_Min_1m | μs (Microseconds) | Minimum Time Consumption within One Minute |
| | Cost_Min_5m | μs (Microseconds) | Minimum Time Consumption within Five Minutes |
| Time Consumption for cosRpc_getDelegationToken Invocation | Cost_Avg | μs (Microseconds) | Average Time Consumption within the Current Second |
| | Cost_Avg_1m | μs (Microseconds) | One-Minute Average Time Consumption |
| | Cost_Avg_5m | μs (Microseconds) | Five-Minute Average Time Consumption |
| | Cost_Max | μs (Microseconds) | Maximum Time Consumption within the Current Second |
| | Cost_Max_1m | μs (Microseconds) | Maximum Time Consumption within One Minute |
| | Cost_Max_5m | μs (Microseconds) | Maximum Time Consumption within Five Minutes |
| | Cost_Min | μs (Microseconds) | Minimum Time Consumption within the Current Second |
| | Cost_Min_1m | μs (Microseconds) | Minimum Time Consumption within One Minute |
| | Cost_Min_5m | μs (Microseconds) | Minimum Time Consumption within Five Minutes |
| Time Consumption for cosRpc_checkPermission Invocation | Cost_Avg | μs (Microseconds) | Average Time Consumption within the Current Second |
| | Cost_Avg_1m | μs (Microseconds) | One-Minute Average Time Consumption |
| | Cost_Avg_5m | μs (Microseconds) | Five-Minute Average Time Consumption |
| | Cost_Max | μs | Maximum Time Consumption within the |

| | | | |
|--|-------------|----------------------|--|
| | | (Microseconds) | Current Second |
| | Cost_Max_1m | μs (Microseconds) | Maximum Time Consumption within One Minute |
| | Cost_Max_5m | μs (Microseconds) | Maximum Time Consumption within Five Minutes |
| | Cost_Min | μs (Microseconds) | Minimum Time Consumption within the Current Second |
| | Cost_Min_1m | μs (Microseconds) | Minimum Time Consumption within One Minute |
| | Cost_Min_5m | μs (Microseconds) | Minimum Time Consumption within Five Minutes |
| Time Consumption for cosRpc_getRangerAuthPolicy Invocation | Cost_Avg | μs (Microseconds) | Average Time Consumption within the Current Second |
| | Cost_Avg_1m | μs (Microseconds) | One-Minute Average Time Consumption |
| | Cost_Avg_5m | μs (Microseconds) | Five-Minute Average Time Consumption |
| | Cost_Max | μs (Microseconds) | Maximum Time Consumption within the Current Second |
| | Cost_Max_1m | μs (Microseconds) | Maximum Time Consumption within One Minute |
| | Cost_Max_5m | μs (Microseconds) | Maximum Time Consumption within Five Minutes |
| | Cost_Min | μs (Microseconds) | Minimum Time Consumption within the Current Second |
| | Cost_Min_1m | μs (Microseconds) | Minimum Time Consumption within One Minute |
| | Cost_Min_5m | μs (Microseconds) | Minimum Time Consumption within Five Minutes |

Doris Monitoring Metrics

Last updated: 2023-12-21 16:52:21

Doris-FE

| Title | Description | Metric Unit | Description |
|--------------------------------|------------------------|-------------|--|
| Node Information | FeNodeNum | count | Total Number of FE Nodes |
| | BeAliveNum | count | Number of Active BE Nodes |
| | BkDeadNum | count | Number of Deceased Broker Nodes |
| Number of CONNECTIONS | Num | count | Number of JVM Connections for FE Nodes |
| Number of JVM Threads | Total | count | Total number of threads in the JVM of FE nodes, inclusive of both daemon and non-daemon threads. |
| | Peak | count | Peak Value of JVM Threads for FE Nodes |
| Garbage Collection Frequency | YoungGC | count | Number of Young GC Events in JVM of FE Nodes |
| | OldGC | count | Number of Old GC Events in JVM of FE Nodes |
| Garbage Collection Duration | YoungGC | s | Duration of Young GC Events in JVM of FE Nodes |
| | OldGC | s | Duration of Old GC Events in JVM of FE Nodes |
| Query Latency in FE Nodes | Quantile75 | ms | 75th Percentile of FE Query Latency |
| | Quantile95 | ms | 95th Percentile of FE Query Latency |
| | Quantile99 | ms | 99th Percentile of FE Query Latency |
| | Quantile999 | ms | 99.9th Percentile of FE Query Latency |
| Peak Score of TABLET_COMPACTON | MAX | score | Maximum Compaction Score during FE Tablet Compaction |
| Quantity of SCHEDULED_TABLET | ScheduledTablet | count | Number of Scheduled Tablets in FE |
| Request Response | QPS | count | Queries Per Second |
| | RPS | count | Number of Requests Processable Per Second |
| Query Failure Rate | ErrRate | % | Query Error Rate |
| Cache Query | SqlModelHitQuery | count | Number of Queries Hitting the Cache in SQL Mode |
| | PartitionModelHitQuery | count | Number of Queries Hit Through Partition |
| | SqlModelQuery | count | Number of Queries Identified in SQL Cache Mode |
| | PartitionModelQuery | count | Number of Queries Identified in Partition Cache Mode |
| | CachePartitionHit | count | Number of Partitions Hit in Query via Cache |
| | CachePartitionScan | count | Total Number of Partitions Scanned in Query |
| Number of Rows in ROUTINE_LOAD | TotalRows | count | FE Routine Load Row Count |
| | ErrorRows | count | FE Routine Load Number of Erroneous Rows |
| TRANSACTION Status Statistics | Reject | count | FE Number of Rejected Transactions |

| | | | |
|------------------------------|---------------------|-------|--|
| | Begin | count | FE Number of Initiated Transactions |
| | Success | count | FE Number of Successful Transactions |
| | Failed | count | FE Number of Failed Transactions |
| IMAGE Quantity | Write | count | FE Quantity of Image Writes |
| | Push | count | FE Quantity of Image Pushes |
| ALTER Task Statistics | RollupRunning | count | Quantity of Alter Jobs in Progress, of ROLLUP Type |
| | SchemaChangeRunning | count | Quantity of Alter Jobs in Progress, of SCHEMA_CHANGE Type |
| UNKNOWN_LOAD Task Statistics | UNKNOWN | count | Quantity of Load Jobs of UNKNOWN Type, with UNKNOWN Status |
| | PENDING | count | Quantity of Load Jobs of UNKNOWN Type, with PENDING Status |
| | ETL | count | Quantity of Load Jobs of UNKNOWN Type, in ETL Status |
| | LOADING | count | Quantity of Load Jobs of UNKNOWN Type, in LOADING Status |
| | COMMITTED | count | Quantity of Load Jobs of UNKNOWN Type, in COMMITTED Status |
| | FINISHED | count | Quantity of Load Jobs of UNKNOWN Type, in FINISHED Status |
| | CANCELLED | count | Quantity of Load Jobs of UNKNOWN Type, in CANCELLED Status |
| SPARK_LOAD Task Statistics | UNKNOWN | count | Quantity of Load Jobs of SPARK Type, in UNKNOWN Status |
| | PENDING | count | Quantity of Load Jobs of SPARK Type, in PENDING Status |
| | ETL | count | Quantity of Load Jobs of SPARK Type, in ETL Status |
| | LOADING | count | Quantity of Load Jobs of SPARK Type, in LOADING Status |
| | COMMITTED | count | Quantity of Load Jobs of SPARK Type, in COMMITTED Status |
| | FINISHED | count | Quantity of Load Jobs of SPARK Type, in FINISHED Status |
| | CANCELLED | count | Quantity of Load Jobs of SPARK Type, in CANCELLED Status |
| DELETE_LOAD Task Statistics | UNKNOWN | count | Quantity of Load Jobs of DELETE Type, in UNKNOWN Status |
| | PENDING | count | Quantity of Load Jobs of DELETE Type, in PENDING Status |
| | ETL | count | Quantity of Load Jobs of DELETE Type, in ETL Status |
| | LOADING | count | Quantity of Load Jobs of DELETE Type, in LOADING Status |
| | COMMITTED | count | Quantity of Load Jobs of DELETE Type, in COMMITTED Status |
| | FINISHED | count | Quantity of Load Jobs of DELETE Type, in FINISHED Status |
| | CANCELLED | count | Quantity of Load Jobs of DELETE Type, in CANCELLED Status |
| INSERT_LOAD Task | UNKNOWN | count | Quantity of Load Jobs of INSERT Type, in UNKNOWN |

| Statistics | | | Status |
|-----------------------------|-----------|-------|---|
| Statistics | PENDING | count | Quantity of Load Jobs of INSERT Type, in PENDING Status |
| | ETL | count | Quantity of Load Jobs of INSERT Type, in ETL Status |
| | LOADING | count | Quantity of Load Jobs of INSERT Type, in LOADING Status |
| | COMMITTED | count | Quantity of Load Jobs of INSERT Type, in COMMITTED Status |
| | FINISHED | count | Quantity of Load Jobs of INSERT Type, in FINISHED Status |
| | CANCELLED | count | Quantity of Load Jobs of INSERT Type, in CANCELLED Status |
| BROKER_LOAD Task Statistics | UNKNOWN | count | Quantity of Load Jobs of BROKER Type, in UNKNOWN Status |
| | PENDING | count | Quantity of Load Jobs of BROKER Type, in PENDING Status |
| | ETL | count | Quantity of Load Jobs of BROKER Type, in ETL Status |
| | LOADING | count | Quantity of Load Jobs of BROKER Type, in LOADING Status |
| | COMMITTED | count | Quantity of Load Jobs of BROKER Type, in COMMITTED Status |
| | FINISHED | count | Quantity of Load Jobs of BROKER Type, in FINISHED Status |
| MINI_LOAD Task Statistics | UNKNOWN | count | Quantity of Load Jobs of MINI Type, in UNKNOWN Status |
| | PENDING | count | Quantity of Load Jobs of MINI Type, in PENDING Status |
| | ETL | count | Quantity of Load Jobs of MINI Type, in ETL Status |
| | LOADING | count | Quantity of Load Jobs of MINI Type, in LOADING Status |
| | COMMITTED | count | Quantity of Load Jobs of MINI Type, in COMMITTED Status |
| | FINISHED | count | Quantity of Load Jobs of MINI Type, in FINISHED Status |
| HADOOP_LOAD Task Statistics | CANCELLED | count | Quantity of Load Jobs of MINI Type, in CANCELLED Status |
| | UNKNOWN | count | Quantity of Load Jobs of HADOOP Type, in UNKNOWN Status |
| | PENDING | count | Quantity of Load Jobs of HADOOP Type, in PENDING Status |
| | ETL | count | Quantity of Load Jobs of HADOOP Type, in ETL Status |
| | LOADING | count | Quantity of Load Jobs of HADOOP Type, in LOADING Status |
| | COMMITTED | count | Quantity of Load Jobs of HADOOP Type, in COMMITTED Status |
| HADOOP_LOAD Task Statistics | FINISHED | count | Quantity of Load Jobs of HADOOP Type, in FINISHED Status |
| | CANCELLED | count | Quantity of Load Jobs of HADOOP Type, in CANCELLED Status |

Doris-BE

| Title | Description | Metric Unit | Description |
|--|----------------------|------------------------|---|
| Utilization of THRIFT | Broker | count | Number of Brokers Utilizing Thrift |
| | Backend | count | Number of BEs Utilizing Thrift |
| | Extdatasource | count | Number of Extdatasources Utilizing Thrift |
| | Frontend | count | Number of FEs Utilizing Thrift |
| STREAMING_LOAD Statistics | RequestsTotal | count | Streaming Load Request Quantity |
| | CurrentProcessing | count | Current Process Count for Streaming Load |
| | PipeCount | count | Streaming Load Pipe Quantity |
| STREAMING_LOAD Duration | Duration | ms | Streaming Load Duration |
| STREAMING_LOAD Data Volume | LoadTotal | bytes | Streaming Load Imported Data Size |
| FRAGMENT Statistics | PlanFragment | count | Plan Fragment Quantity |
| | Endpoint | count | Quantity of DataStreams |
| | RequestsTotal | count | Request Count of Fragment |
| FRAGMENT Request Time | Duration | μ s (Microseconds) | Request Time of Fragment |
| BE Memory | Total | bytes | Size of BE Memory Pool |
| | Allocated | bytes | Size of Allocated BE Memory |
| Maximum Score of TABLET_COMPACTON | CumulativeMax | score | Maximum Base Compaction Score within Tablet |
| | BaseMax | score | Maximum Compaction Score of Tablet Base |
| Volume of COMPACTION Data | Cumulative | bytes | Data Volume of Cumulative Compaction |
| | Base | bytes | Data Volume of Base Compaction |
| Volume of COMPACTION_DELTAS Data | Cumulative | bytes | Data Volume of Cumulative Compaction Deltas |
| | Base | bytes | Data Volume of Base Compaction Deltas |
| Quantity of MemPool Utilized by COMPACTION | CurrentConsumption | count | Aggregate MemPool Utilized by Compaction (All Compaction Threads) |
| Number of Process File Handles | Used | count | Number of File Handles Utilized by BE Process |
| | SoftLimit | count | Soft Limit Quantity of File Handles for BE Process |
| | HardLimit | count | Hard Limit Quantity of File Handles for BE Process |
| Number of Running Threads in Process | NUM | count | Number of Threads Running in the BE Process |
| ENGINE REQUESTS Statistics | FailedBaseCompaction | count | Number of Failed Engine Requests of Type Base_Compaction |
| | FailedCultCompt | count | Number of Failed Engine Requests of Type Cumulative_Compaction |
| | TotalBaseCompaction | count | Total Number of Engine Requests of Type Base_Compaction |
| | TotalCultCompt | count | Total Number of Engine Requests of Type |

Cumulative_Compaction

Doris-BK

| Title | Description | Metric Unit | Description |
|------------------------------|-------------------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| JVM Memory | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of HeapMemory Committed by JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of HeapMem in JVM |
| | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| | MemNonHeapInitM | MB | Initial Quantity of NonHeapMem in JVM |
| CPU utilization | ProcessCpuLoad | % | CPU utilization |
| Number of File Handles | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| | OpenFileDescriptorCount | Connections | Number of Open File Descriptors |
| CPU Usage Duration | ProcessCpuTime | ms | Accumulated CPU Usage Time |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| Number of Working Threads | ThreadCount | Connections | Number of Threads |
| | PeakThreadCount | Connections | Peak Thread Count |
| | DaemonThreadCount | Connections | Background Thread Count |

Kylin Monitoring Metrics

Last updated: 2023-12-21 16:52:28

Kylin- Kylin

| Title | Description | Metric Unit | Description |
|------------------------------|-------------------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| JVM Memory | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of HeapMemory Committed by JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of JVM HeapMemory |
| | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of NonHeapMemory Committed by JVM |
| | MemNonHeapInitM | MB | Initial Quantity of JVM NonHeapMemory |
| File Descriptor Count | OpenFileDescriptorCount | Connections | Quantity of Open File Descriptors |
| | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| CPU utilization | ProcessCpuLoad | % | Process CPU Utilization |
| | SystemCpuLoad | % | System CPU Utilization |
| Number of Working Threads | DaemonThreadCount | Connections | Daemon Thread Count |
| | ThreadCount | Connections | Total Thread Count |
| Accumulated CPU Usage Time | ProcessCpuTime | ms | Accumulated CPU Usage Time |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |

Zeppelin Monitoring Metrics

Last updated: 2023-12-21 16:52:36

Zeppelin-Zeppelin

| Title | Description | Metric Unit | Description |
|------------------------------|-------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |

Oozie Monitoring Metrics

Last updated: 2023-12-21 16:52:43

Oozie-Oozie

| Title | Description | Metric Unit | Description |
|------------------------------|-------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |

Storm Monitoring Metrics

Last updated: 2023-12-21 16:52:58

Storm-Nimbus

| Title | Description | Metric Unit | Description |
|------------------------------|-------------------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| JVM Memory | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of HeapMemory Committed by JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of JVM HeapMemory |
| | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of NonHeapMemory Committed by JVM |
| | MemNonHeapInitM | MB | Initial Quantity of JVM NonHeapMemory |
| File Descriptor Count | OpenFileDescriptorCount | Connections | Quantity of Open File Descriptors |
| | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| CPU utilization | ProcessCpuLoad | % | Process CPU Utilization |
| Number of Working Threads | DaemonThreadCount | Connections | Daemon Thread Count |
| | PeakThreadCount | Connections | Peak Thread Count |
| | ThreadCount | Connections | Total Thread Count |
| Accumulated CPU Usage Time | ProcessCpuTime | ms | Accumulated CPU Usage Time |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |

Storm-Supervisor

| Title | Description | Metric Unit | Description |
|------------------------------|-------------------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| JVM Memory | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of HeapMemory Committed by JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of JVM HeapMemory |
| | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of NonHeapMemory Committed by JVM |
| | MemNonHeapInitM | MB | Initial Quantity of JVM NonHeapMemory |
| File Descriptor Count | OpenFileDescriptorCount | Connections | Quantity of Open File Descriptors |
| | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| CPU utilization | ProcessCpuLoad | % | Process CPU Utilization |
| Number of Working Threads | DaemonThreadCount | Connections | Daemon Thread Count |
| | PeakThreadCount | Connections | Peak Thread Count |
| | ThreadCount | Connections | Total Thread Count |
| Accumulated CPU Usage Time | ProcessCpuTime | ms | Accumulated CPU Usage Time |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |

Livy Monitoring Metrics

Last updated: 2023-12-21 16:53:04

Livy-LivyServer

| Title | Description | Metric Unit | Description |
|------------------------------|-------------------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| JVM Memory | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of HeapMemory Committed by JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of JVM HeapMemory |
| | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of NonHeapMemory Committed by JVM |
| | MemNonHeapInitM | MB | Initial Quantity of JVM NonHeapMemory |
| File Descriptor Count | OpenFileDescriptorCount | Connections | Quantity of Open File Descriptors |
| | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| CPU utilization | ProcessCpuLoad | % | Process CPU Utilization |
| Number of Working Threads | DaemonThreadCount | Connections | Daemon Thread Count |
| | PeakThreadCount | Connections | Peak Thread Count |
| | ThreadCount | Connections | Total Thread Count |
| Accumulated CPU Usage Time | ProcessCpuTime | ms | Accumulated CPU Usage Time |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |

Kyuubi Monitoring Metrics

Last updated: 2023-12-21 16:53:15

Kyuubi-KyuubiServer

| Title | Description | Metric Unit | Description |
|------------------------------|-------------------------|-------------|---|
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | FGCT | s | Full Garbage Collection Time Consumption |
| | GCT | s | Garbage Collection Time Expenditure |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| JVM Memory | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemHeapCommittedM | MB | Quantity of HeapMemory Committed by JVM |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapInitM | MB | Initial Quantity of JVM HeapMemory |
| | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of NonHeapMemory Committed by JVM |
| | MemNonHeapInitM | MB | Initial Quantity of JVM NonHeapMemory |
| File Descriptor Count | OpenFileDescriptorCount | Connections | Quantity of Open File Descriptors |
| | MaxFileDescriptorCount | Connections | Maximum File Descriptor Quantity |
| CPU utilization | ProcessCpuLoad | % | Process CPU Utilization |
| Number of Working Threads | DaemonThreadCount | Connections | Daemon Thread Count |
| | PeakThreadCount | Connections | Peak Thread Count |
| | ThreadCount | Connections | Total Thread Count |
| Accumulated CPU Usage Time | ProcessCpuTime | ms | Accumulated CPU Usage Time |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| Frequency of Events | Ratio | count/s | Frequency of Events Entering the Running |

| | | | |
|--|-------|---------|---|
| Entering the Running Operational State | | | Operational State |
| Frequency of Events Entering the Pending Operational State | Ratio | count/s | Frequency of Events Entering the Pending Operational State |
| Frequency of Events Entering the Initialized Operational State | Ratio | count/s | Frequency of Events Entering the Initialized Operational State |
| Frequency of Events Entering the Finished Operational State | Ratio | count/s | Frequency of Events Entering the Finished Operational State |
| Frequency of Events Entering the Closed Operational State | Ratio | count/s | Frequency of Events Entering the Closed Operational State |
| Frequency of FetchResultRows Method Invocation | Ratio | count/s | Frequency of FetchResultRows Method Invocation |
| Frequency of FetchLogRows Method Invocation | Ratio | count/s | Frequency of FetchLogRows Method Invocation |
| Execution Time of GetResultSetMetadata Method Invocation | P50 | ms | Median Execution Time of GetResultSetMetadata Method Invocation |
| | P75 | | 75th Percentile of Execution Time for GetResultSetMetadata Method Invocation |
| | P95 | | 95th Percentile of Execution Time for GetResultSetMetadata Method Invocation |
| | P99 | | 99th Percentile of Execution Time for GetResultSetMetadata Method Invocation |
| | P999 | | 999th Percentile of Execution Time for GetResultSetMetadata Method Invocation |
| Execution Time for CloseOperation Method Invocation | P50 | ms | Median Execution Time for CloseOperation Method Invocation |
| | P75 | | 75th Percentile of Execution Time for CloseOperation Method Invocation |
| | P95 | | 95th Percentile of Execution Time for CloseOperation Method Invocation |
| | P99 | | 99th Percentile of Execution Time for CloseOperation Method Invocation |
| | P999 | | 999th Percentile of Execution Time for CloseOperation Method Invocation |
| Execution Time for GetOperationStatus Method Invocation | P50 | ms | Median Execution Time for GetOperationStatus Method Invocation |
| | P75 | | 75th Percentile of Execution Time for GetOperationStatus Method Invocation |
| | P95 | | 95th Percentile of Execution Time for GetOperationStatus Method Invocation |
| | P99 | | 99th Percentile of Execution Time for GetOperationStatus Method Invocation |
| | P999 | | 999th Percentile of Execution Time for |

| | | | |
|---|------|----|--|
| | | | GetOperationStatus Method Invocation |
| Execution Time for GetInfo Method Invocation | P50 | ms | 50th Percentile of Execution Time for GetInfo Method Invocation |
| | P75 | | 75th Percentile of Execution Time for GetInfo Method Invocation |
| | P95 | | 95th Percentile of Execution Time for GetInfo Method Invocation |
| | P99 | | 99th Percentile of Execution Time for GetInfo Method Invocation |
| | P999 | | 999th Percentile of Execution Time for GetInfo Method Invocation |
| Execution Time for FetchResults Method Invocation | P50 | ms | Median Execution Time for FetchResults Method Invocation |
| | P75 | | 75th Percentile of Execution Time for FetchResults Method Invocation |
| | P95 | | 95th Percentile of Execution Time for FetchResults Method Invocation |
| | P99 | | 99th Percentile of Execution Time for FetchResults Method Invocation |
| | P999 | | 999th Percentile of Execution Time for FetchResults Method Invocation |
| Execution Time for ExecuteStatement Method Invocation | P50 | ms | Median Execution Time for ExecuteStatement Method Invocation |
| | P75 | | 75th Percentile of Execution Time for ExecuteStatement Method Invocation |
| | P95 | | 95th Percentile of Execution Time for ExecuteStatement Method Invocation |
| | P99 | | 99th Percentile of Execution Time for ExecuteStatement Method Invocation |
| | P999 | | 99.9th Percentile of Execution Time for ExecuteStatement Method Invocation |
| Execution Time for CloseSession Method Invocation | P50 | ms | Median Execution Time for CloseSession Method Invocation |
| | P75 | | 75th Percentile of Execution Time for CloseSession Method Invocation |
| | P95 | | 95th Percentile of Execution Time for CloseSession Method Invocation |
| | P99 | | 99th Percentile of Execution Time for CloseSession Method Invocation |
| | P999 | | 999th Percentile of Execution Time for CloseSession Method Invocation |
| Execution Time for OpenSession Method Invocation | P50 | ms | Median Execution Time for OpenSession Method Invocation |
| | P75 | | 75th Percentile of Execution Time for OpenSession Method Invocation |
| | P95 | | 95th Percentile of Execution Time for OpenSession Method Invocation |

| | | | |
|--|---------------------|-------------|--|
| | P99 | | 99th Percentile of Execution Time for OpenSession Method Invocation |
| | P999 | | 999th Percentile of Execution Time for OpenSession Method Invocation |
| Current Number of Launch Engine Operations | Count | Connections | Current Number of Launch Engine Operations |
| Current Number of Execute Statement Operations | Count | Connections | Current Number of Execute Statement Operations |
| Current Number of Operations | Count | Connections | Current Number of Operations |
| Number of JVM Threads | ThreadsNew | Connections | Number of Threads in NEW State |
| | ThreadsRunnable | Connections | Number of Threads in RUNNABLE State |
| | ThreadsBlocked | Connections | Number of Threads in BLOCKED State |
| | ThreadsWaiting | Connections | Number of Threads in WAITING State |
| | ThreadsTimedWaiting | Connections | Number of Threads in TIMED WAITING State |
| | ThreadsTerminated | Connections | Number of Threads in Terminated State |
| Current Active Connection Count | Count | count | Current Active Connection Count |
| Cumulative Connection Count | Count | count | Cumulative Connection Count |
| Active Thread Count in Thread Pool | Count | count | Active Thread Count in Thread Pool |
| Available Thread Count in Thread Pool | Count | count | Available Thread Count in Thread Pool |
| Failed Metadata Request Count | Count | count | Failed Metadata Request Count |
| Current Metadata Request Count | Count | count | Current Metadata Request Count |
| Metadata Request Count | Count | count | Metadata Request Count |
| Cumulative Operation Count | Count | count | Cumulative Operation Count |

StarRocks Monitoring Metrics

Last updated: 2023-12-21 16:53:24

StarRocks-BE

| Title | Description | Metric Unit | Description |
|---------------------------------------|-----------------------|-------------|---|
| Compaction Deltas Data Volume | Cumulative | rowsets | Data Volume of Cumulative Compaction Deltas |
| | Base | rowsets | Data Volume of Base Compaction Deltas |
| | Update | rowsets | Data Volume of Update Compaction Deltas |
| Volume of COMPACTION Data | Cumulative | bytes | Data Volume of Cumulative Compaction Bytes |
| | Base | bytes | Data Volume of Base Compaction Bytes |
| | Update | bytes | Data Volume of Update Compaction Bytes |
| Maximum Score of TABLET COMPACTION | CumulativeMax | score | Maximum Base Compaction Score within Tablet |
| | BaseMax | score | Maximum Compaction Score of Tablet Base |
| ENGINE Request Failure Statistics (1) | base_compaction | count | Number of Engine Failure Requests of Type Base_Compaction |
| | clone | count | Number of Engine Failure Requests of Type Clone |
| | create_rollup | count | Number of Engine Failure Requests of Type Create_Rollup |
| | create_tablet | count | Number of Engine Failure Requests of Type Create_Tablet |
| | cumulative_compaction | count | Number of Engine Failure Requests of Type Cumulative_Compaction |
| | delete | count | Number of Engine Failure Requests of Type Delete |
| | finish_task | count | Number of Engine Failure Requests of Type Finish_Task |
| Engine Request Failure Statistics (2) | publish | count | Number of Engine Failure Requests of Type Publish |
| | report_all_tablets | count | Number of Engine Failure Requests of Type Report_All_Tablets |
| | report_disk | count | Number of Engine Failure Requests of Type Report_Disk |
| | report_tablet | count | Number of Engine Failure Requests of Type Report_Tablet |
| | report_task | count | Number of Engine Failure Requests of Type Report_Task |
| | schema_change | count | Number of Engine Failure Requests of Type Schema_Change |
| Engine Request Statistics (1) | base_compaction | count | Number of Engine Failure Requests of Type Base_Compaction |
| | clone | count | Number of Engine Failure Requests of Type Clone |
| | create_rollup | count | Number of Engine Failure Requests of Type Create_Rollup |
| | create_tablet | count | Number of Engine Failure Requests of Type |

| | | | |
|--------------------------------------|-----------------------|------------------------|---|
| | | | Create_Tablet |
| | cumulative_compaction | count | Number of Engine Failure Requests of Type Cumulative_Compaction |
| | delete | count | Number of Engine Failure Requests of Type Delete |
| | drop_tablet | count | Number of Engine Failure Requests of Type Drop_Tablet |
| | finish_task | count | Number of Engine Failure Requests of Type Finish_Task |
| Engine Request Statistics (2) | publish | count | Number of Engine Failure Requests of Type Publish |
| | report_all_tablets | count | Number of Engine Failure Requests of Type Report_All_Tablets |
| | report_disk | count | Number of Engine Failure Requests of Type Report_Disk |
| | report_tablet | count | Number of Engine Failure Requests of Type Report_Tablet |
| | report_task | count | Number of Engine Failure Requests of Type Report_Task |
| | schema_change | count | Number of Engine Failure Requests of Type Schema_Change |
| | storage_migrate | count | Number of Engine Failure Requests of Type Storage_Migrate |
| FRAGMENT Statistics | PlanFragment | count | Plan Fragment Quantity |
| | Endpoint | count | Quantity of DataStreams |
| FRAGMENT Request Time | Duration | Microsecond (μ s) | Request Time of Fragment |
| TXN Request Statistics | begin | count | Quantity of Requests with 'begin' as TXN Type |
| | commit | count | Quantity of Requests with 'commit' as TXN Type |
| | exec | count | Quantity of Requests with 'exec' as TXN Type |
| | rollback | count | Quantity of Requests with 'rollback' as TXN Type |
| STREAMING LOAD Data Volume | LoadTotal | bytes | Data Size Imported via Stream Load |
| STREAMING LOAD Statistics | CurrentProcessing | count | Current Process Count for Streaming Load |
| | PipeCount | count | Number of Pipes in Streaming Load |
| STREAMING LOAD Duration | Duration | ms | Streaming Load Duration |
| BE Memory | Total | bytes | Size of BE Memory Pool |
| | Allocated | bytes | Size of Allocated BE Memory |
| Number of Process File Handles | Used | count | Number of File Handles Utilized by BE Process |
| | SoftLimit | count | Soft Limit Quantity of File Handles for BE Process |
| | HardLimit | count | Hard Limit Quantity of File Handles for BE Process |
| Number of Running Threads in Process | Thread | count | Number of Threads Running in the BE Process |
| | | | |

| | | | |
|--------------------------------|----------|-------|------------------------------------|
| Quantity of THRIFT Utilization | Broker | count | Number of Brokers Utilizing Thrift |
| | Backend | count | Number of BEs Utilizing Thrift |
| | Frontend | count | Number of FEs Utilizing Thrift |
| TABLET Write Statistics | Writer | count | BE TABLET Write Statistics |

StarRocks-FE

| Title | Description | Metric Unit | Description |
|--------------------------------|---------------------|-------------|---|
| ALTER Task Statistics | RollupRunning | count | Number of Alter Jobs in Progress, of Type ROLLUP |
| | SchemaChangeRunning | count | Number of Alter Jobs in Progress, of Type SCHEMA_CHANGE |
| IMAGE Statistics | Write | count | FE Quantity of Image Writes |
| | Push | count | FE Quantity of Image Pushes |
| Quantity of SCHEDULED_TABLET | ScheduledTablet | count | Number of Scheduled Tablets in FE |
| TRANSACTION Status Statistics | Reject | count | FE Number of Rejected Transactions |
| | Begin | count | FE Number of Initiated Transactions |
| | Success | count | FE Number of Successful Transactions |
| | Failed | count | FE Number of Failed Transactions |
| JVM Memory_HEAP | max | bytes | Maximum Heap Memory |
| | committed | bytes | Committed Heap Memory |
| | used | bytes | Utilized Heap Memory |
| JVM Memory_NONHEAP | committed | bytes | Committed Non-Heap Memory |
| | used | bytes | Utilized Non-Heap Memory |
| JVM Memory_OLD | used | bytes | Utilized Old Generation Memory |
| | peak_used | bytes | Maximum Utilized Old Generation Memory |
| | max | bytes | Maximum Old Generation Memory |
| JVM Memory_YOUNG | used | bytes | Utilized Young Memory |
| | peak_used | bytes | Maximum Utilized Young Memory |
| | max | bytes | Maximum Young Memory |
| ROUTINE LOAD QUEUE Dimension | report queue | count | Size of the FE Report Queue |
| Number of Rows in ROUTINE_LOAD | TotalRows | count | Number of Rows in FE Routine Load |
| | ErrorRows | count | FE Routine Load Number of Erroneous Rows |
| Size of ROUTINE LOAD | Receive | bytes | Size of the FE Routine Load |
| Peak Score of TABLET_COMPACTON | MAX | score | Maximum Compaction Score during FE Tablet Compaction |
| EDITLOG Write Delay | Quantile75 | ms | 75th Percentile of FE Editlog Write Delay |
| | Quantile95 | ms | 95th Percentile of FE Editlog Write Delay |

| | | | |
|------------------------------|-------------|---------|---|
| | Quantile98 | ms | 98th Percentile of FE Editlog Write Delay |
| | Quantile99 | ms | 99th Percentile of FE Editlog Write Delay |
| | Quantile999 | ms | 99.9th Percentile of FE Editlog Write Delay |
| Garbage Collection Frequency | YoungGC | count | Number of Young GC Events in JVM of FE Nodes |
| | OldGC | count | FE Node JVM Old GC Event Count |
| Garbage Collection Duration | YoungGC | Second. | Duration of Young GC Events in JVM of FE Nodes |
| | OldGC | Second. | Duration of Old GC Events in JVM of FE Nodes |
| Number of JVM Threads | Total | count | Total Number of Threads in FE Node's JVM |
| | Peak | count | Peak JVM Threads of FE Node |
| BROKER_LOAD Task Statistics | UNKNOWN | count | Quantity of Load Jobs of BROKER Type, in UNKNOWN Status |
| | PENDING | count | Quantity of Load Jobs of BROKER Type, in PENDING Status |
| | ETL | count | Quantity of Load Jobs of BROKER Type, in ETL Status |
| | LOADING | count | Quantity of Load Jobs of BROKER Type, in LOADING Status |
| | COMMITTED | count | Quantity of Load Jobs of BROKER Type, in COMMITTED Status |
| | FINISHED | count | Quantity of Load Jobs of BROKER Type, in FINISHED Status |
| | CANCELLED | count | Quantity of Load Jobs of BROKER Type, in CANCELLED Status |
| DELETE_LOAD Task Statistics | UNKNOWN | count | Quantity of Load Jobs of DELETE Type, in UNKNOWN Status |
| | PENDING | count | Quantity of Load Jobs of DELETE Type, in PENDING Status |
| | ETL | count | Quantity of Load Jobs of DELETE Type, in ETL Status |
| | LOADING | count | Quantity of Load Jobs of DELETE Type, in LOADING Status |
| | COMMITTED | count | Quantity of Load Jobs of DELETE Type, in COMMITTED Status |
| | FINISHED | count | Quantity of Load Jobs of DELETE Type, in FINISHED Status |
| | CANCELLED | count | Quantity of Load Jobs of DELETE Type, in CANCELLED Status |
| HADOOP_LOAD Task Statistics | UNKNOWN | count | Quantity of Load Jobs of HADOOP Type, in UNKNOWN Status |
| | PENDING | count | Quantity of Load Jobs of HADOOP Type, in PENDING Status |
| | ETL | count | Quantity of Load Jobs of HADOOP Type, in ETL Status |
| | LOADING | count | Quantity of Load Jobs of HADOOP Type, in LOADING Status |
| | COMMITTED | count | Quantity of Load Jobs of HADOOP Type, in COMMITTED Status |

| | | | |
|------------------------------|----------------------|-------|--|
| | FINISHED | count | Quantity of Load Jobs of HADOOP Type, in FINISHED Status |
| | CANCELLED | count | Quantity of Load Jobs of HADOOP Type, in CANCELLED Status |
| INSERT_LOAD Task Statistics | UNKNOWN | count | Quantity of Load Jobs of INSERT Type, in UNKNOWN Status |
| | PENDING | count | Quantity of Load Jobs of INSERT Type, in PENDING Status |
| | ETL | count | Quantity of Load Jobs of INSERT Type, in ETL Status |
| | LOADING | count | Quantity of Load Jobs of INSERT Type, in LOADING Status |
| | COMMITTED | count | Quantity of Load Jobs of INSERT Type, in COMMITTED Status |
| | FINISHED | count | Quantity of Load Jobs of INSERT Type, in FINISHED Status |
| | CANCELLED | count | Quantity of Load Jobs of INSERT Type, in CANCELLED Status |
| ROUTINE_LOAD Task Statistics | NEED_SCHEDULE | count | Routine Load Jobs Statistics, State=NEED_SCHEDULE |
| | RUNNING | count | Routine Load Jobs Statistics, State=RUNNING |
| | PAUSED | count | Routine Load Jobs Statistics, State=PAUSED |
| | STOPPED | count | Routine Load Jobs Statistics, State=STOPPED |
| | CANCELLED | count | Routine Load Jobs Statistics, State=CANCELLED |
| SPARK_LOAD Task Statistics | UNKNOWN | count | Quantity of Load Jobs of SPARK Type, in UNKNOWN Status |
| | PENDING | count | Quantity of Load Jobs of SPARK Type, in PENDING Status |
| | ETL | count | Quantity of Load Jobs of SPARK Type, in ETL Status |
| | LOADING | count | Quantity of Load Jobs of SPARK Type, in LOADING Status |
| | COMMITTED | count | Quantity of Load Jobs of SPARK Type, in COMMITTED Status |
| | FINISHED | count | Quantity of Load Jobs of SPARK Type, in FINISHED Status |
| | CANCELLED | count | Quantity of Load Jobs of SPARK Type, in CANCELLED Status |
| FE MASTER | FE Master | count | Indication of FE Master Status; 1 for Master, 0 for Follower |
| Node Information | FeNodeNum | count | Total Number of FE Nodes |
| | BeNodeNum | count | Total Number of BE Nodes |
| | BeAliveNum | count | Number of Active BE Nodes |
| | BeDecommissioned Num | count | Number of Active BE Nodes |
| | BkDeadNum | count | Number of Deceased Broker Nodes |

| | | | |
|---------------------------|-------------|---------|---|
| Request Response | QPS | count/s | Queries Per Second |
| | RPS | count/s | Number of Requests Processable Per Second |
| FE Query Statistics | total | count | Total Number of FE Queries |
| | err | count | Total Number of FE Query Errors |
| | timeout | count | Number of FE Query Timeouts |
| | success | count | Total Number of Successful FE Queries |
| | slow | count | Total Number of Slow FE Queries |
| Query Failure Rate | ErrRate | % | Query Error Rate |
| Query Latency in FE Nodes | Quantile75 | ms | 75th Percentile of FE Query Latency |
| | Quantile95 | ms | 95th Percentile of FE Query Latency |
| | Quantile99 | ms | 99th Percentile of FE Query Latency |
| | Quantile999 | ms | 99.9th Percentile of FE Query Latency |
| Number of CONNECTIONS | Num | count | FE Node Connection Quantity |

StarRocks-Broker

| Title | Description | Metric Unit | Description |
|------------------------------|----------------------|-------------|---|
| CPU utilization | ProcessCpuLoad | % | Process CPU Utilization |
| CPU Usage Duration | ProcessCpuTime | ms | Accumulated CPU Usage Time |
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | GCT | s | Garbage Collection Time Expenditure |
| | FGCT | s | Full Garbage Collection Time Consumption |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| JVM Memory | MemHeapInitM | MB | Initial Quantity of JVM Heap Memory |
| | MemNonHeapInitM | MB | Initial Quantity of JVM NonHeapMemory |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapCommittedM | MB | Quantity of Currently Committed HeapMemory in JVM |
| | MemHeapUsedM | MB | Current Quantity of HeapMemory in Use by JVM |
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |

| | | | |
|---------------------------|-------------------------|-------|---|
| | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| Number of File Handles | OpenFileDescriptorCount | count | Number of Open File Descriptors |
| | MaxFileDescriptorCount | count | Maximum File Descriptor Quantity |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| Number of Working Threads | PeakThreadCount | count | Peak Thread Count |
| | ThreadCount | count | Total Thread Count |
| | DaemonThreadCount | count | Daemon Thread Quantity |

Kafka Monitoring Metrics

Last updated: 2023-12-21 16:53:31

| Title | Description | Metric Unit | Description |
|------------------------------|-------------------------|-------------|---|
| CPU utilization | ProcessCpuLoad | % | Process CPU Utilization |
| CPU Usage Duration | ProcessCpuTime | ms | Accumulated CPU Usage Time |
| Garbage Collection Frequency | YGC | Threads | Young Garbage Collection Frequency |
| | FGC | Threads | Full Garbage Collection Frequency |
| Garbage Collection Duration | GCT | s | Garbage Collection Time Expenditure |
| | FGCT | s | Full Garbage Collection Time Consumption |
| | YGCT | s | Young Garbage Collection Time Consumption |
| Memory Region Occupancy | O | % | Old Region Memory Usage Proportion |
| | M | % | Metaspace Region Memory Usage Proportion |
| | CCS | % | Compressed Class Space Region Memory Usage Proportion |
| | S0 | % | Survivor 0 Region Memory Usage Proportion |
| | S1 | % | Survivor 1 Region Memory Usage Proportion |
| | E | % | Eden Region Memory Usage Proportion |
| JVM Memory | MemHeapInitM | MB | Initial Quantity of JVM HeapMemory |
| | MemNonHeapInitM | MB | Initial Quantity of JVM NonHeapMemory |
| | MemHeapMaxM | MB | Configured Quantity of HeapMemory in JVM |
| | MemHeapCommittedM | MB | Current Committed Quantity of JVM HeapMemory |
| | MemHeapUsedM | MB | Current Utilized Quantity of JVM HeapMemory |
| | MemNonHeapCommittedM | MB | Quantity of Currently Committed NonHeapMemory in JVM |
| | MemNonHeapUsedM | MB | Current Quantity of NonHeapMemory in Use by JVM |
| Number of File Handles | OpenFileDescriptorCount | count | Number of Open File Descriptors |
| | MaxFileDescriptorCount | count | Maximum File Descriptor Quantity |
| Process Runtime Duration | Uptime | s | Process Runtime Duration |
| Number of Working Threads | PeakThreadCount | count | Peak Thread Count |
| | ThreadCount | count | Total Thread Count |
| | DaemonThreadCount | count | Daemon Thread Quantity |
| Broker Production Traffic | OneMinuteRate | bytes/s | One-Minute Broker Message Production Traffic |
| Broker Consumption Traffic | OneMinuteRate | bytes/s | One-Minute Broker Message Consumption Traffic |

| | | | |
|--|--------------------------------|---------|---|
| Rejected Consumption Traffic | OneMinuteRate | bytes/s | One-Minute Topic Request Rejection Rate |
| Fetch Failure Request Count | OneMinuteRate | count/s | One-Minute Fetch Failed Request Count |
| Produce Failure Request Count | OneMinuteRate | count/s | One-Minute Produce Failed Request Count |
| Message Production Quantity | OneMinuteRate | count/s | One-Minute Message Production Rate |
| Traffic Read from Other Brokers | OneMinuteRate | bytes/s | One-Minute Rate of Reading from Other Brokers |
| Traffic Read from Other Brokers | OneMinuteRate | bytes/s | One-Minute Rate of Reading from Other Brokers |
| Number of Fetch Requests | OneMinuteRate | count/s | One-Minute Total Fetch Request Rate |
| Number of Produce Requests | OneMinuteRate | count/s | One-Minute Total Produce Request Rate |
| ControllerBroker | IsControllerBroker | - | The metric value on the Broker where the Controller is located is 1, while the value on other Brokers is 0. |
| LeaderElection Rate | OneMinuteRate | count/s | One-Minute LeaderElection Rate |
| LeaderElection Latency | 99thPercentile | ms | LeaderElection Latency_99thPercentile |
| | 999thPercentile | | LeaderElection Latency_999thPercentile |
| | Mean | | LeaderElection Latency_Mean |
| UncleanLeaderElections Rate | OneMinuteRate | count/s | One-Minute UncleanLeaderElections Rate |
| GlobalPartition Count | GlobalPartitionCount | count | The Global Partition Count Observed by This Controller |
| OfflinePartitions Count | OfflinePartitionCount | count | The Number of Offline Partitions Observed by This Controller |
| GlobalTopic Count | GlobalTopicCount | count | The Quantity of GlobalTopics Observed by This Controller |
| Number of Offline Log Directories | OfflineLogDirectory | count | Quantity of Offline Log Directories |
| LogFlush Rate | OneMinuteRate | calls/s | One-Minute Message Log Refresh Rate |
| LogFlush Latency | 99thPercentile | ms | LogFlush Delay_99th Percentile |
| | 999thPercentile | | LogFlush Delay_999th Percentile |
| | Mean | | LogFlush Delay_Mean |
| Average Idle Rate of Network Processor | NetworkProcessorAvgIdlePercent | % | Average Idle Ratio of Network Thread Pool Threads |
| ISR Expansion Rate | OneMinuteRate | count/s | One Minute ISR Expansion Rate |
| ISR Contraction Rate | OneMinuteRate | count/s | One Minute ISR Contraction Rate |
| Quantity of Replicas | LeaderReplicaCount | count | Number of Leader Replicas |
| | OfflineReplicaCount | count | Number of Offline Replicas |
| Number of Partitions | PartitionCount | count | Number of partitions |

| | | | |
|---|---------------------------|------|--|
| | UnderMinIsrPartitionCount | | Number of Partitions under Minimum In-Sync Replica (ISR) Count |
| | UnderReplicatedPartitions | | Number of UnderReplicatedPartitions |
| FetchConsumer Request Latency | 99thPercentile | ms | FetchConsumer Request Time - 99th Percentile |
| | 999thPercentile | | FetchConsumer Request Time - 999th Percentile |
| | Mean | | Average FetchConsumer Request Time |
| FetchFollower Request Latency | 99thPercentile | ms | FetchFollower Request Time - 99th Percentile |
| | 999thPercentile | | FetchFollower Request Time - 999th Percentile |
| | Mean | | Average FetchFollower Request Time |
| Produce Request Latency | 99thPercentile | ms | Produce Request Time - 99th Percentile |
| | 999thPercentile | | Produce Request Time - 999th Percentile |
| | Mean | | Average Produce Request Time |
| Request Queue Size | RequestQueueSize | size | Request Queue Size |
| Purgatory Size | Fetch | size | Number of Requests Waiting in Fetch Purgatory |
| | Produce | size | Number of Requests Waiting in Producer Purgatory |
| Average Idle Rate of Request Processing | OneMinuteRate | % | One-Minute Idle Rate for Request Processing |
| ZooKeeper Request Latency | 99thPercentile | ms | ZooKeeper Request Latency_99th Percentile |
| | 999thPercentile | | ZooKeeper Request Latency_999th Percentile |
| | Mean | | ZooKeeper Request Latency_Mean |

Alarm Configurations

Last updated: 2023-12-21 16:53:42

Scenario

Elastic MapReduce has been integrated into the Tencent Cloud Observability Platform. Users can configure alarm strategies for Elastic MapReduce node and service monitoring indicators on the Tencent Cloud Observability Platform console.

Note

- Elastic MapReduce has been integrated with the default alarms of the Tencent Cloud Observability Platform, which will automatically create default alarm strategies. For information on the indicators/events supported by the default Elastic MapReduce strategy or alarm rules, please refer to the [Default Strategy Description](#).
- Supports the manual creation of alarm strategies, set as the default alarm strategy. Once set successfully, newly purchased instances will automatically associate with the default strategy, eliminating the need for manual addition.

Instructions

- Log into the Tencent Cloud Observability Platform console and select **Alarm Management** > **Strategy Management** from the left navigation bar.
- On the "Strategy Management" page, click on **Create New Strategy**.
- In the pop-up "Create New Alarm Strategy" window, refer to the table below to configure basic information, alarm rules, and create a new notification template.

| Configuration Specifications | Configuration items | Note | |
|------------------------------|-----------------------|--|--|
| Basic information | Rule Name | Custom Rule Name | |
| | Remarks | Custom Rule Remarks | |
| | Test mode | Supported Cloud Product Monitoring Types | |
| | Policy types | Select the cloud product policy type you wish to monitor. | |
| | Project | After setting the associated project, you can swiftly filter the alarm policies under this project in the alarm policy list. | |
| Configuring Alarm Rules | Alert object | By selecting the Instance ID, the alarm policy will be bound to the chosen instance. By selecting the Instance Group, the alarm policy will be bound to the chosen instance grouping. By selecting all objects, the alarm policy will be bound to all instances for which the current account has permissions. | |
| | Trigger condition | Manual Configuration(Metric alarm) | Alarm Trigger Condition: A semantically meaningful condition composed of metrics, comparison relationships, thresholds, statistical periods, and duration periods. You can expand the details of the trigger condition and set an appropriate alarm threshold based on the trend of metric changes in the chart. |
| | | Manual Configuration (Event Alarm) | In the event of anomalies in cloud product resources or underlying infrastructure services, you can create event alarms to promptly notify you to take action. |
| | | Trigger condition (choosing Select template) | Click the Select Template button and choose a pre-configured template from the dropdown list. For specific configuration details, please refer to Configuring Trigger Condition Templates . |
| Configure alarm | Notification Template | By default, the system binds to a preset notification template (with the primary account administrator as the recipient and SMS and email as the delivery channels). | |

notification
(optional)

Each alarm policy can bind to a maximum of three notification templates. For more information on configuring notification templates, please refer to [Creating a New Notification Template](#).

4. Click **Complete** after filling in the information.

Alarm History

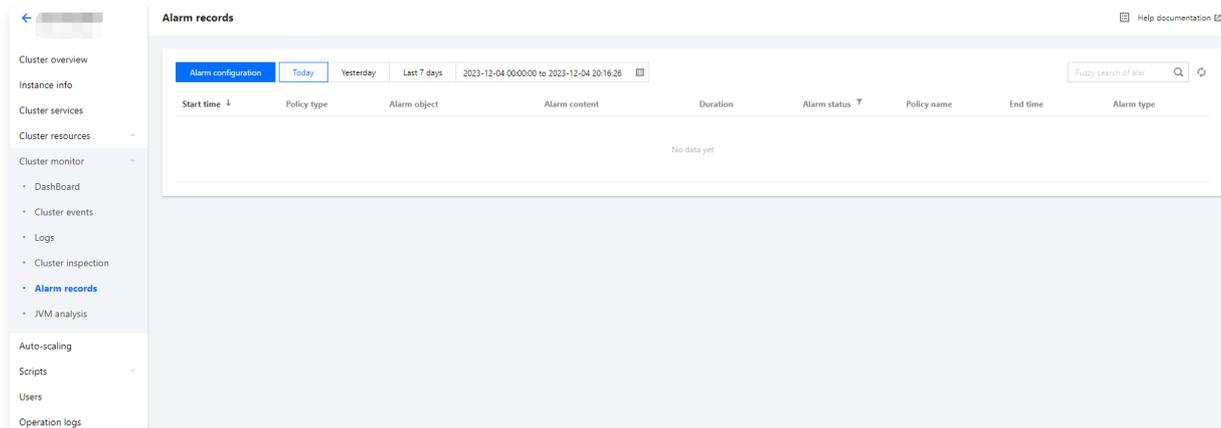
Last updated: 2023-12-21 18:19:24

Feature Overview

The Alarm History feature enables you to retrospect and review the alarm history of the current cluster for the past six months. You can also swiftly navigate to the corresponding alarm policy and subscribe to it via the Alarm History page by policy type.

Instructions

1. Log in to the [EMR Console](#), and in the cluster list, click on the corresponding cluster **ID/Name** to access the cluster details page.
2. In the Cluster Details page, select **Cluster Monitoring > Alarm records**. You can perform a fuzzy search for the corresponding alarms based on the alarm content, or filter by time range and alarm status. Sorting can be done based on the start time.



3. By clicking on **Alarm Configuration**, you can swiftly access the Tencent Cloud Observability Platform's alarm configuration page to set up alarm policies. For specific operations, please refer to [Alarm Configuration](#).