Elastic MapReduce

EMR on TKE Operation Guide

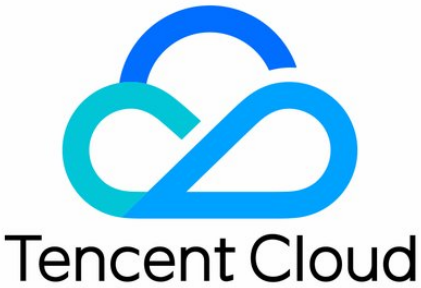

Service Notice

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

EMR on TKE Operation Guide Configuring Cluster Administrative rights Role Authorization

Last updated: 2023-12-25 10:54:36

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 TKE, COS, etc.) to establish clusters and preserve logs, among other functions.

> **Note**
> Upon initial activation of the Elastic MapReduce service, it is imperative that the primary account completes the role authorization process. Failure to do so will result in both the subsidiary and primary accounts 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.

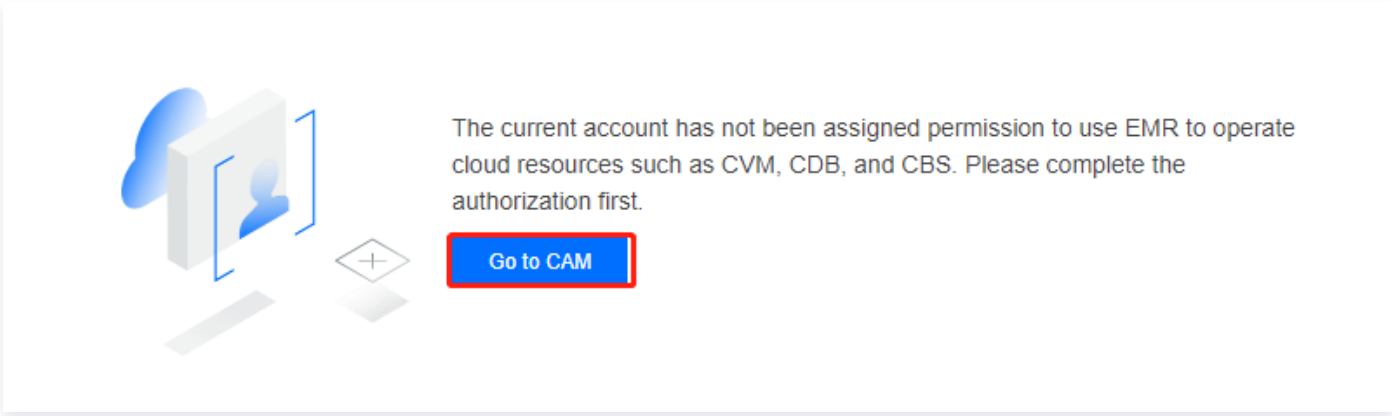

2. Click on **Grant** to assign the default role, EMR_QCSRole, to the service account of Elastic MapReduce.

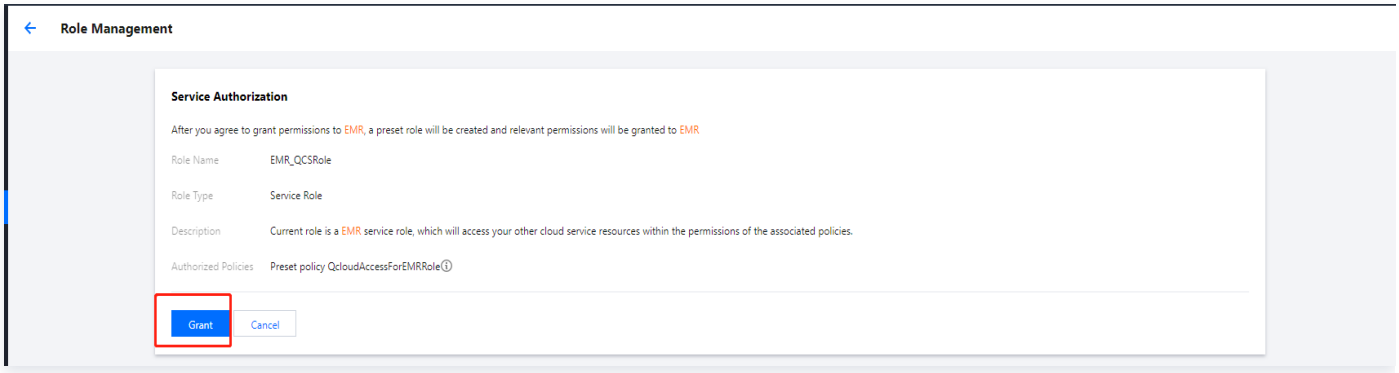

3. Upon completion of the authorization, users are required to refresh the console or purchase page of Elastic MapReduce to resume normal operations. For more detailed policy information related to EMR_QCSRole, please log in to the Access Management Console. For information on the permissions included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

Special Instructions for EMR on TKE Cluster Related Service Role Authorization

When creating or utilizing an EMR on TKE cluster, it is necessary to directly write or compute data on COS (Object Storage). To ensure data security, EMR needs to be granted read and write access to COS resources via temporary keys. This requires granting the relevant EMR service role `EMR_QCSLinkedRoleInApplicationDataAccess` and binding the preset policy `QcloudAccessForEMRLinkedRoleInApplicationDataAccess`.

1. When conducting an EMR on TKE cluster list, it is necessary to verify whether the EMR service related role `EMR_QCSLinkedRoleInApplicationDataAccess` is bound.

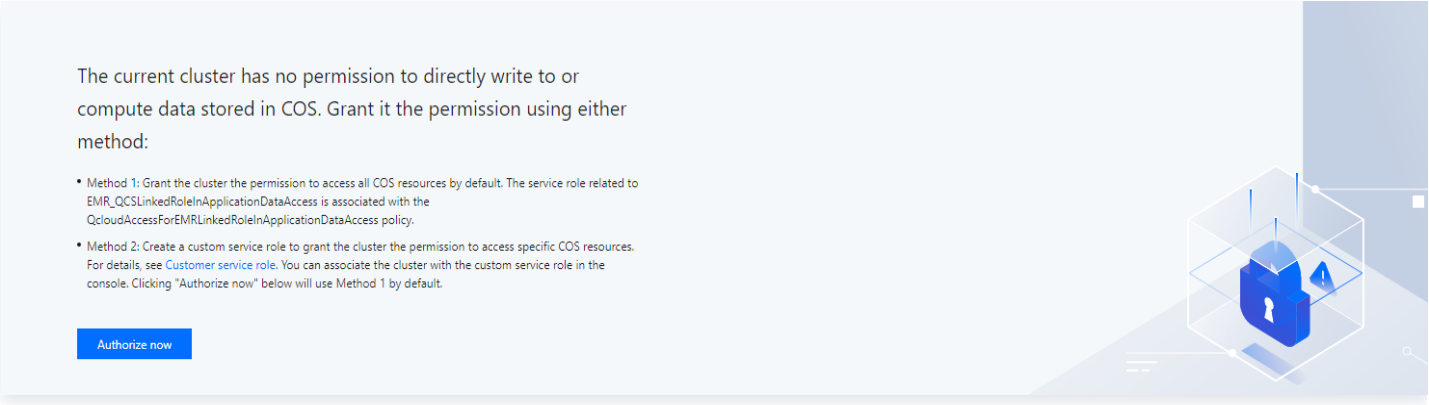

2. If the EMR service related role `EMR_QCSLinkedRoleInApplicationDataAccess` does not exist, authorization binding is required.

Note

> If there is a need to specify cluster access to corresponding COS resource permissions in a more refined manner, please refer to the Custom Service Role settings for more details.

Authentication Instructions for EMR on TKE Cluster

- The permission settings for sub-accounts and collaborators are consistent with the EMR on CVM version. For more details, please refer to Collaborator/Sub-account Permissions.
- The settings for tag authentication and interface authentication are consistent with the EMR on CVM version. For more details, please refer to the Authentication Granularity Scheme.

Creating a Cluster

Last updated: 2023-12-25 10:55:37

Scenario

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

Instructions

Log into the EMR console, click on 'Create Cluster' in the EMR on TKE cluster list page, and complete the necessary configurations on the purchase page. When the cluster status in the cluster list displays as 'Running', it signifies that the cluster has been successfully created.

Configuration items	Configuration Item Description
Cluster name	Distinguish between different EMR clusters by setting unique cluster names. The system generates these names randomly, but they can be modified.
Regions	The cluster is deployed in a physical data center, with each region representing a separate physical data center. The internal networks of cloud servers in different regions do not communicate with each other. Note 1. Once a cluster is created, the region cannot be changed. Please choose carefully. 2. Only container clusters within the same region can provide resources for the cluster.
Container type	The service role is deployed using resources provided by the underlying container, supporting both TKE standard clusters and TKE Serverless clusters. Note: Select a specific container cluster to provide Pod resources for the deployment of this EMR on TKE cluster.
Cluster network	This is used for purchasing the database, ensuring network consistency with the container cluster network.
AZ & Subnet	This is used for purchasing the database, ensuring network consistency with the container cluster network.
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 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. 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.

	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.
Billing Mode	Only the pay-as-you-go billing mode is supported. Pay-as-you-go: Payment is made according to the duration of use. Real-name authentication is required for the account, and a fee equivalent to 2 hours of usage needs to be frozen at the time of activation (coupons cannot be used as freezing vouchers). The frozen resource fee will be refunded upon termination.
Instance Edition	The components and their versions bundled with different product editions vary.
Deploy task	Optional components can be customized and deployed according to individual needs, with at least one component service required for deployment. Supported deployment forms include: default deployment and associated deployment. Associated Deployment: It refers to the source of service for the current cluster (other EMR clusters or DLC on the cloud). Default Deployment: This involves deploying the current service to the selected container resources, with the ability to customize the specifications and quantity of deployment resources (deployed on a service role dimension). This method is highly recommended.
Associating	Select a cluster from EMR on CVM or EMR on TKE that contains the current service. 1. When selecting an EMR on CVM cluster, it is imperative that the current service is deployed within that cluster. 2. When selecting an EMR on TKE cluster, it is essential to configure the current service within the cluster to be capable of providing external services. 3. MetaStore (HiveMetaStore) supports the selection of DLC service sources.
Default	Data Volume Directory: Provides a data storage directory for the container. Currently, it only supports using the host path (HostPath mounting method) or creating a new PVCPOD. The service role dimension deployment information is optional. External Access Settings: Configure whether the current role is enabled to provide external services. It supports optional selections for public network LB access and internal network LB access. Node Scheduling Strategy: Establish the deployment scheduling strategy for the current role. According to the scheduling rules, Pods can be scheduled to nodes that meet the expected Label criteria.
Hive Metadata Database	If the Hive component is selected, the Hive metadata database offers two storage methods: The first type is the default cluster, where the Hive metadata is stored in the MetaDB, which is independently purchased for the cluster. The second type involves linking to an external Hive metadata database. 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. Please refer to Hive Metadata Management for more information. 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.

COS bucket	This is used for storing information such as logs and jar packages.
Tag	When creating, you can add tags to the cluster or node resources to facilitate the management of these resources. You can bind up to 5 tags, and the tag keys must be unique.
Set a password	Establish a WebUI password. The current password is solely for the initial setup of the WebUI access password.
Terms of Service	Agree to the "Elastic MapReduce Service Level Agreement" and "Refund Agreement".

Cluster Management
Adjusting Pod quantity

Last updated: 2023-12-25 10:55:54

Scenario

In instances where the resources utilized by the services within the cluster are either insufficient or excessive, the cluster's load capacity can be altered by adjusting the number of PODs deployed for the service role.

> **Note**
> The current expansion of POD specifications defaults to the POD specifications selected when creating a new cluster. If there is a need to adjust the service resource configuration, it should be set in Change Configuration.

Instructions

1. Log in to the EMR on TKE Console, and in the cluster list, click on the **Cluster ID/Name** to access the cluster details page.
2. From the cluster details page, select **Cluster Services > Specific Service** and proceed to the role management page.
3. On the role management page, select the service role that needs to be scaled up/down and click on **Adjust POD Quantity**.
4. Adjust the POD quantity according to actual needs, and click **Confirm** to initiate the adjustment.

Changing Configurations

Last updated: 2023-12-25 10:56:08

Scenario

When the service role's CPU or memory size is insufficient, it can be enhanced by modifying the configuration. This document introduces the relevant operations for changing the instance configuration through the EMR on TKE Console.

> **Note**
> - Modifying the configuration may cause running jobs to fail. Please ensure to carry out this operation without affecting the business operations.
> - During the configuration modification process, resources may be preempted, causing the POD rebuild to fail and the service to become unavailable. It is advisable to perform this operation during periods of low business activity.

Preparations

1. Adjustments to the POD resource configuration will be automatically deducted on an item-by-item basis. Please ensure that your account balance is sufficient.
2. After the configuration of pay-as-you-go nodes is changed, the cost will be frozen again according to the unit price of the latest configuration. For more details, please refer to EMR on TKE Billing Instructions > Purchase Instructions.

Instructions

1. Log in to the EMR on TKE Console, and in the cluster list, click on the **Cluster ID/Name** to access the cluster details page.
2. On the cluster details page, select **Cluster Services > A Particular Service** and proceed to the role management page.
3. On the role management page, select the service role that needs to be reconfigured and click on **Change Configuration**.
4. On the configuration adjustment page, confirm the relevant reconfiguration information, select the target CPU and memory configuration, and carefully read the important tips and check the box to agree to the change information.
5. After confirming that the charges are correct, click on **Confirm** to adjust the configuration.

Task Center

Last updated: 2023-12-25 10:56:19

The detailed information pertaining to the cluster tasks in EMR on TKE aligns seamlessly with the version of EMR on CVM. For further elucidation, kindly refer to the Task Center.

Reboot

Last updated: 2023-12-25 10:56:31

Scenario

Reconstruction refers to the annihilation of all POD resources deployed at the current role dimension, and the repurchase of POD resources consistent with the original configuration specifications. This is primarily applicable when there is an "increase or decrease" in the configuration file, utilizing the reconstruction function to activate its configuration.

> ⚠ **Note:**
> - Restart: This is solely applicable when there is a change in the configuration parameters, enabling them to take effect.
> - Reconstruction: There are instances of reconstruction failure. Such failure may occur when the load of the service role cannot be initiated, leading to service unavailability. Please proceed with caution and ensure operations are conducted without impacting the business.

Instructions

1. Log in to the EMR on TKE Console, and in the cluster list, click on the **Cluster ID/Name** to access the cluster details page.
2. On the cluster details page, select **Cluster Services > Service (for example, select Hive)**, and proceed to the role management page.
3. On the role management page, select the service role that needs to be reconstructed and click on **Reconstruct**.
4. On the reconstruction configuration page, confirm the relevant configuration information, select the reconstruction method (supports: rolling reconstruction and rapid reconstruction), and carefully read the important prompts and check the box to agree to the reconstruction information.
5. Click on **Confirm** to reconstruct the POD.
6. In the POD list, click on the POD status to view the status details, where you can see up to 10 status data entries for the current POD in the last hour.

Managing Service Deployment Instructions

Last updated: 2023-12-25 10:56:46

Deployment instructions are utilized to examine the types of clusters supported for deployment by EMR on TKE, service role information, default Pod specification data, and data volume directories. This facilitates users in referencing and adjusting service deployment resource configurations.

Deployment Information Description

Cluster Type	Component	Role	Default Pod Specifications			Data volume directory
			CPU Limit (Cores)	Memory Limit (GB)	Pod Quantity Range	
TKE	hive	HiveMetaStore	1-100	2-100	1-100	Create New PVC, Host Path
		HiveServer2	1-100	2-100	1-100	Create New PVC, Host Path
	hue	Hue	2-100	4-100	1-100	Create New PVC, Host Path
	kyuubi	KyuubiServer	1-100	2-100	2-100	Create New PVC, Host Path
	ranger	Ranger	1-100	2-100	1-100	Create New PVC, Host Path
		RangerUsersync	1-100	2-100	1-100	Create New PVC, Host Path
	rss	Coordinator	1-100	2-100	2-100	Create New PVC, Host Path
		ShuffleServer	2-100	4-100	3-100	Create New PVC
	spark	SparkJobHistoryServer	1-16	1-64	1-2	Create New PVC, Host Path
	trino	TrinoCoordinator	1-100	2-100	1-2	Create New PVC, Host Path
		TrinoWorker	1-100	2-100	2-100	Create New PVC, Host Path
	zookeeper	QuorumPeerMain	1-100	2-100	3-100	Create New PVC

	knox	Knox	1–100	2–100	1–100	Create New PVC, Host Path
TKE Serverless	hive	HiveMetaStore	1–100	2–100	1–100	Create New PVC
		HiveServer2	1–100	2–100	1–100	
	hue	Hue	2–100	4–100	1–100	
	kyuubi	KyuubiServer	1–100	2–100	2–100	
	ranger	Ranger	1–100	2–100	1–100	
		RangerUsersync	1–100	2–100	1–100	
	rss	Coordinator	1–100	2–100	2–100	
		ShuffleServer	2–100	4–100	3–100	
	spark	SparkJobHistoryServer	1–16	1–64	1–2	
	trino	TrinoCoordinator	1–100	2–100	1–2	
		TrinoWorker	1–100	2–100	2–100	
	zookeeper	QuorumPeerMain	1–100	2–100	3–100	
	knox	Knox	1–100	2–100	1–100	

Note

The maximum specification of the service role Pod in the cluster will be dynamically adjusted in conjunction with the available resources of the selected container cluster. The actual maximum may be less than the maximum given in the table.

User Management

Last updated: 2023-12-25 10:56:57

Under the version of EMR on TKE, the deployment of the OpenLdap service supports user management. For operations related to user management and viewing, please refer to the User Management section of the EMR on CVM version.

Adding Components

Last updated: 2023-12-25 10:57:10

Upon the creation of the cluster, additional components supported by the currently selected product version can be incorporated through the addition of new components. This document provides a guide on how to add new components via the console.

> **Note**
> The addition of new components is exclusively supported for the components of the current EMR product version.

Instructions

1. Log in to the EMR on TKE Console, and in the cluster list, click on the **Cluster ID/Name** to access the cluster details page.
2. Select **Cluster Information > Basic Configuration**, then choose **Add New Component** in the service architecture diagram to incorporate components that have not been installed in the cluster.
3. In the absence of a metadata database in the cluster, the installation of Hue and Ranger components necessitates the purchase of a cloud database instance storage unit for metadata storage.
4. The Hive component metadata storage offers two methods of storage: The first is the cluster default, where metadata is stored in MetaDB. The second is association, either with EMR-MetaDB or a self-built MySQL database. In this case, the metadata will be stored in the associated database and will not be destroyed with the cluster.
5. The decision to purchase MetaDB is consistent with that made when purchasing a new cluster, as is the choice of Hive metadata database.
6. After selecting the components, simply click on **Create**.

Restarting Service

Last updated: 2023-12-25 10:57:32

Scenario

Upon upgrading or modifying the configuration of service roles within a cluster, it is necessary to restart the corresponding services to activate the changes. This document provides a guide on how to restart services via the console.

Instructions

1. Log in to the EMR on TKE Console, and in the cluster list, click on the **Cluster ID/Name** to access the cluster details page.
2. If you need to restart the entire component service, you can select **Cluster Service > Specific Service**, and on the right side of the navigation bar on this page, choose **Operation > Restart Service**. When you select the service role as ALL, the entire service will be restarted.

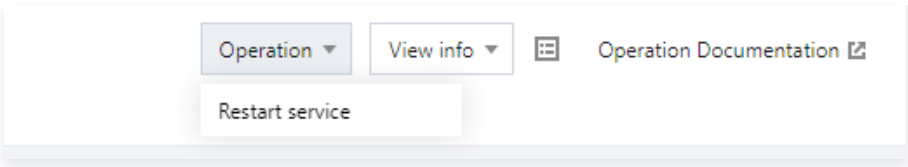

3. If you need to restart a specific role type of the service, you can select **Cluster Service > Specific Service**, check the service role that needs to be restarted on the role management page, and click **Restart Service**.
4. After confirming that everything is correct, click **Confirm** to restart the service.

WebUI Access

Last updated: 2023-12-25 10:57:44

The WebUI portal feature is a native UI access capability provided by EMR components, enabling swift public network access to the native UI through the Knox component.

Preparations

The cluster has installed the Knox component, and the Knox component has enabled public network access.

Instructions

1. Log in to the EMR on TKE Console, and in the cluster list, click on the **Cluster ID/Name** to access the cluster details page.
2. Based on your needs, select **Cluster Services > Specific Service**. On the right side of the navigation bar on this page, select **View Information > View WebUI** to access.

> **Note**
> - The access address requires authentication. The username is 'root', and the default password is the one entered when creating the cluster.
> - If the Knox component was not installed or public access was not enabled for the Knox component during cluster creation, it will not be possible to enter the WebUI interface of the relevant component via the native WebUI access address on the component management page.

Role Management

Last updated: 2023-12-25 10:57:55

Role management offers operational functions at both the role service level and specific POD dimensions, such as reboot operations. POD health status monitoring is supported to facilitate real-time understanding of role process statuses.

Instructions

1. Log in to the EMR on TKE Console, and in the cluster list, click on the **Cluster ID/Name** to access the cluster details page.
2. Within the cluster details page, select **Cluster Services**, and click on the corresponding service to enter the role management page.
3. The role list displays information such as the current service role's POD health status, running/expected POD count, resource configuration, and most recent reboot time. The POD list showcases the health status of POD resources, POD status, POD IP, number of reboots, and the most recent reboot time.

> **Note**
> When the POD status is Running or Succeeded, the POD health status is displayed as normal. If the POD status is Pending, Failed, Unknown, or if the corresponding service's Container is not ready, the POD health status is displayed as abnormal.

Monitoring and Operations
Monitoring Dashboard

Last updated: 2023-12-25 10:59:50

Offers configuration for EMR on TKE monitoring policies and alarm history functions, while also providing the ability to view cluster service monitoring indicators and historical trends.

Monitored metrics

The service monitoring metrics provided by EMR on TKE are as follows:

Sevice	Descriptions Of Monitoring Metrics
Hive	Hive Monitoring Metrics
Spark	Spark Monitoring Metrics
Trino	Trino Monitoring Metrics
Hue	-
Kyuubi	Kyuubi Monitoring Metrics
RSS	-
Ranger	Ranger Monitoring Metrics
Zookeeper	Zookeeper Monitoring Metrics
Knox	-

Logs

Last updated: 2023-12-25 11:00:06

Feature Overview

The Log Search feature offers the capability to gather and search operational logs of components, supporting keyword searches within the core service logs and node system logs of the current cluster. This allows for swift examination of crucial service logs directly from the console.

Instructions

1. Sign in to the EMR Console, and within the EMR on TKE cluster list, click on the corresponding **Cluster ID/Name** to access the cluster details page.
2. On the cluster details page, navigate to **Cluster Information > Service Architecture > Service Name** or select **Service Management > Service Name**, then click on **Logs**. You can filter by the current service's PodName and time range to view log content.
   Alternatively, on the cluster details page, go to **Basic Configuration > Service Architecture**, select **Service Name > Logs**, choose the PodName and time range to conduct a log search.

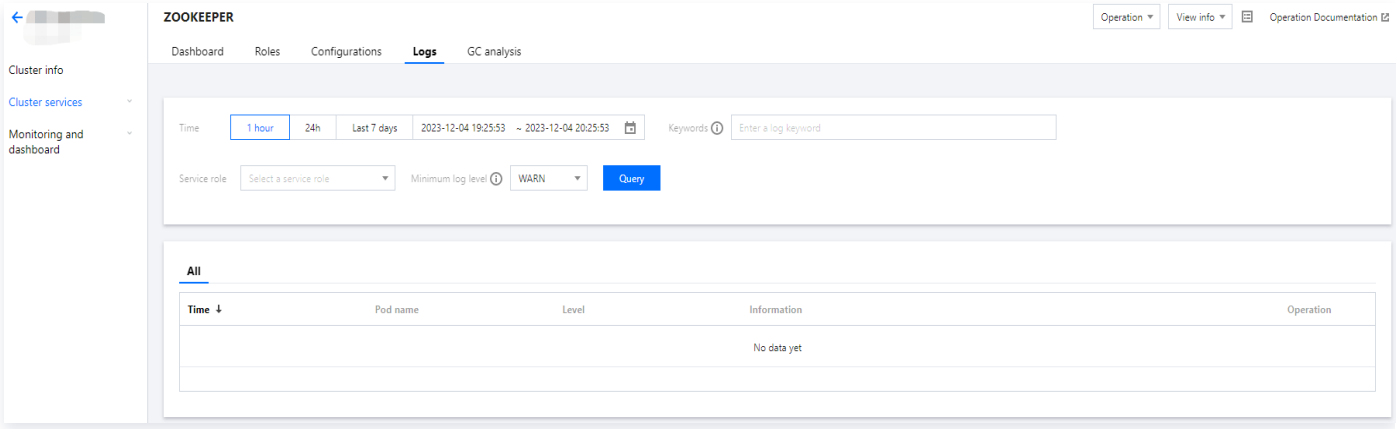

   Click on **PodName** to be redirected to the corresponding node status page. Click on **Log Source** to be redirected to the corresponding node monitoring metrics display page.
   2.1 **Keyword Description**: Full-text keyword search is supported.
   2.2 Supports search using special characters - . * > < = ! () {} /.
   2.3 Phrase search is supported, for instance: 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. This allows for the selection of a time range to view the log context information within that period.

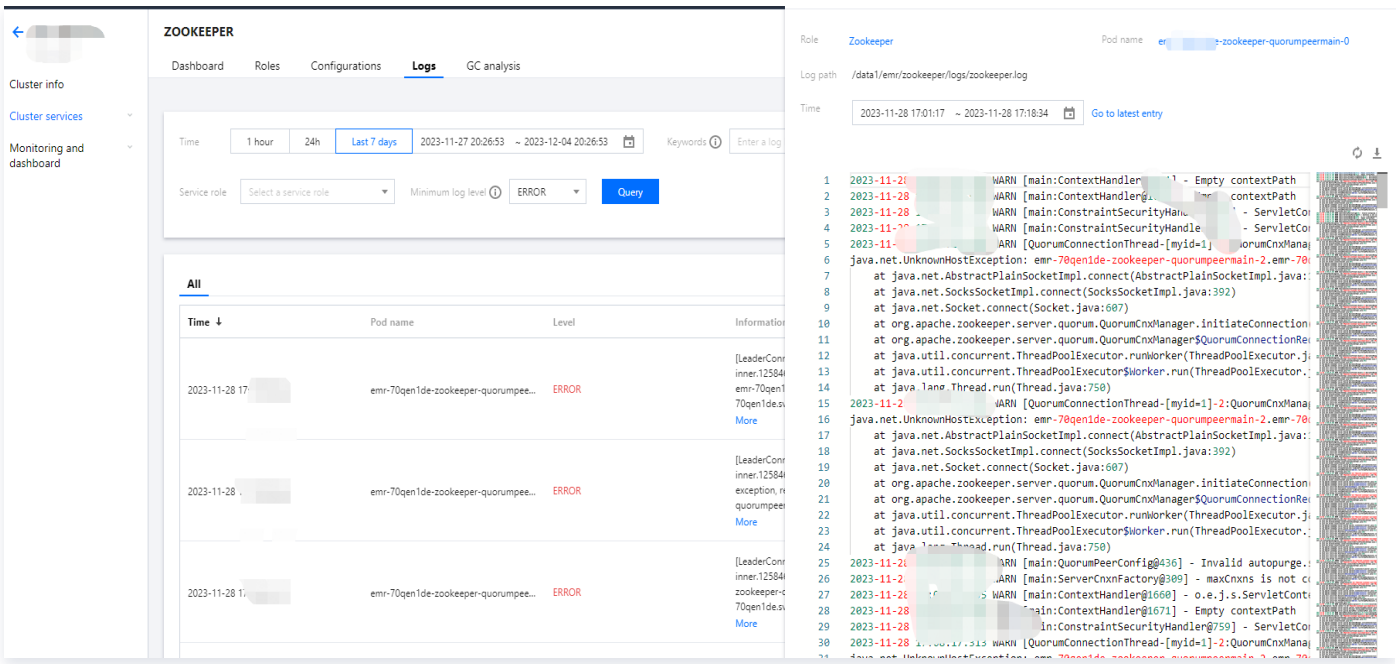

Configuring Alarm

Last updated: 2023-12-25 11:00:51

Configuration of monitoring alerts for EMR on TKE can be referenced from EMR on CVM Alert Configuration. For historical alerts, please refer to Alert History.

GC Analysis

Last updated: 2023-12-25 11:01:02

Scenario

Supports online analysis of Java process GC, by real-time collection, recording, and analysis of GC logs, assisting users in troubleshooting process abnormalities potentially caused by GC.

- GC View: Filter by service, role, node, and time to examine the corresponding GC distribution and pinpoint information.
- GC List: Allows for the filtering and selection of GC log data as needed, providing a multi-dimensional view of GC record information.

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, click on **Cluster Monitoring**, then select **JAVA Analysis > GC Online Analysis** to view the GC view and pinpoint information.
3. Click on the cluster details page **Cluster Information > Service Architecture > Service Card** or **Cluster Services > Service Name**, then select GC Analysis to choose roles and nodes as needed to view the GC view and pinpoint information.
4. It also provides GC query list information, with some column header fields supporting operations such as filtering or sorting.

The services and roles that support GC log collection are as follows:

Sevice	Role
Zookeeper	Zookeeper
Trino	Hmaster
	RegionServer
	HbaseThrift
Hive	HiveMetaStore
	HiveServer2
	HiveWebHcat
Knox	ldap
	gateway
Spark	SparkJobHistoryServer
Ranger	EmbeddedServer
	EnableUnixAuth

RSS	LivyServer
Kyuubi	KyuubiServer