YI-MapReduce Service

User Guide

CHINA TELECOM CLOUD

1. Introduction

1.1. Overview

1.1.1. Overview

YI-MapReduce, also known as YI-MR, is a product that encapsulates the latest open-source big data components. It allows for swift deployment and easy maintenance of high-performance big data components including Doris, Elasticsearch. Additionally, it provides an O&M management platform. The product is equipped with robust security verification capabilities by default, ensuring high security, excellent scalability, and convenient operation and maintenance. It supports various scenarios, such as batch data processing, streaming data processing, offline data analysis, and online queries

1.1.2. Architecture

For the component details of various YI-MapReduce cluster versions, please refer to the Version Overview.

See the figure below: YI-MapReduce Architecture

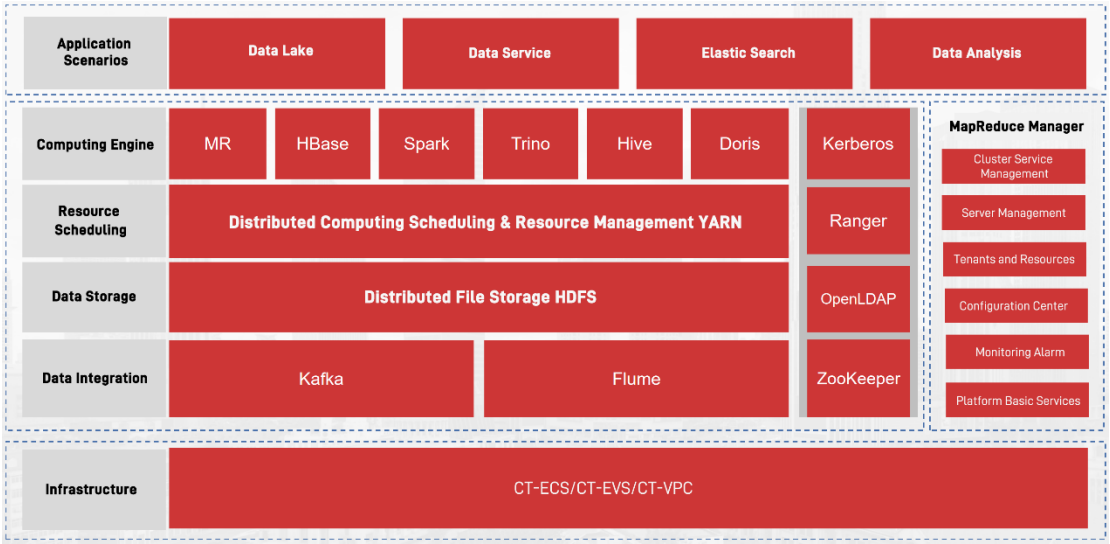

The architecture of YI-MapReduce includes capabilities at all stages of infrastructure and big data processing workflows.

● **Infrastructure**

The big data cluster, built on the eSurfing Cloud Elastic Cloud Server (CT-ECS), has high reliability and security, fully supported by the underlying virtualization.

- The Virtual Private Cloud (CT-VPC) provides each tenant with a virtual internal network, isolated from other networks by default. This, along with the accompanying security group access control, ensures network-level security.
- The Elastic Volume Service (CT-EVS) offers reliable storage capabilities with different specifications and performance.
- The Elastic Cloud Server (CT-ECS) provides an elastically scalable virtual server. Combined with the CT-VPC, security group, multiple replicas of CT-EVS data, and disaster recovery capabilities, it creates an efficient, reliable, and secure business cluster environment for customers.

● Data Storage

YI-MapReduce supports the storage of structured and unstructured data in the cluster, and supports multiple efficient formats to meet the requirements of different computing engines.

  – Doris provides real-time data warehousing service, exhibiting high concurrency and low latency.
  – Elasticsearch facilitates the retrieval and analysis of both structured and unstructured data.

● YI-MapReduce Manager

To maintain the high availability of big data component services, it is necessary to deploy the various components of the big data ecosystem based on Hadoop in a distributed manner. This deployment comes with high demands for management

and O&M complexity. YI-MapReduce offers a unified platform for O&M management, YI-MapReduce Manager, which includes capabilities for visually guided cluster deployment. In addition, YI-MapReduce Manager provides management capabilities for tenants and resources, O&M for various big data components in YI-MapReduce, and one-stop O&M capabilities such as monitoring, alarms, and configuration.

1.1.3. Product Advantages

Performance Optimization

- The MPP Doris engine, fully self-developed, implements vector execution on the Kunpeng CPU, improving overall TPC-H benchmark performance by 50%.
- Our proprietary optimization enables ad-hoc querying of billions of rows and millions of columns in milliseconds.

High Stability and Reliability

Completed more than 100 code and configuration optimizations for open-source components.

Security and Controllability

- Uses Kerberos and Ranger security technologies to authenticate and authorize components.
- Supports data permissions control at the library, table, and field levels.

Easy O&M

Enables full-link visual operations to lower the requirement for O&M, facilitating easy operations for 90% of daily O&M scenarios, and improving O&M efficiency.

1.2. Version Overview

1.2.1. Product Release Version Number

YI-MapReduce employs a version number format similar to "YI-MapReduce-a.b.c" as follows:

- "a" indicates major changes in the version.
- "b" indicates changes in some components of the version.
- "c" indicates bug fixes that are backward compatible, as well as minor changes.

1.2.2. Bundled Software and Cluster Creation for Each YI-YI-MapReduce Version

- Bundled Software: The software and its version, bundled with each release version of YI-MapReduce, remain constant.
- Cluster Creation: Once a YI-MapReduce cluster for a specific release version is created, the version of the cluster will not upgrade automatically.

1.2.3. YI-MapReduce-2.14.1

Release Time: October 2024

Components	Version
Doris	2.0.4
Elasticsearch	7.10.2
Kibana	7.10.2
Logstash	7.10.2

1.3. Product Advantages

YI-MapReduce (or "YI-MR") service has undergone extensive refinements through large-scale clusters and business operations from the Group and Cloud Company, providing annual multi-level user SLA guarantees.

The advantages of YI-MapReduce include:

1.3.1. High reliability

YI-MapReduce has implemented over 100 code and configuration optimizations on open-source components, achieving high SLA. The Master node group employs anti-affinity technology, distributing virtual machines across different physical machines to ensure high service availability.

1.3.2. Easy O&M

YI-MapReduce offers a visual O&M management platform for big data clusters and enables full-link visual operations to lower the requirement for O&M, facilitating easy operations for 90% of daily O&M scenarios, thus improving O&M efficiency.

1.3.3. Cost-effectiveness

The YI-MapReduce cluster, leveraging a diverse cloud infrastructure, offers a broad range of choices for computing and storage facilities. It supports on-demand creation and expansion and can be discarded after use, ensuring cost optimization.

1.4. Product Features

1.4.1. Multi-Tenancy

Introduction

Modern enterprises are increasingly centralizing their data clusters and transitioning to cloud computing. Diverse users execute various applications (such as analysis, search, stream processing, etc.) across multiple clusters, while storing data in a variety of types and formats. Certain users (such as banks and government institutions) attach great importance to data security and cannot allow their data to be mixed with others.

YI-MapReduce provides multi-tenancy functionality, segregating the resources of the big data cluster into independent sets. Users can "rent" the required sets of resources for application execution, work, and data storage. The big data cluster offers multiple sets of resources to cater to the diverse requirements of different users.

Advantages

Appropriate Configuration and Isolation of Resources:

Resources are isolated between tenants. One tenant's utilization of resources does not impact other tenants, ensuring each tenant can configure relevant resources in line with business needs, thereby improving resource utilization efficiency.

Measurement and Statistics of Resource Consumption:

System resources are planned and allocated based on the tenant unit. The tenant is the applicant and consumer of system resources, and the consumption of user resources is measured and counted.

Ensuring Data Security and Access Security:

In multi-tenant scenarios, data from different tenants is stored separately, and user access permissions to tenant resources are controlled, ensuring data security and access security.

1.4.2. Enhanced Security

YI-MapReduce, an eSurfing Cloud big data platform, supports industry applications such as massive data storage, large-scale data analysis, and real-time processing, providing high-level security. The product ensures safe data storage, utilization, and business operations in the following aspects:

Network Isolation

The entire system is deployed in a Virtual Private Cloud (VPC) dedicated to the user on the public cloud, providing a securely isolated network environment to ensure the security of user big data cluster business and management. By combining features such as subnet partitioning, route control, and security groups of the VPC, a highly secure and reliable network isolation environment can be provided to users.

Resource Isolation

YI-MapReduce service currently supports dedicated deployment within the eSurfing Cloud 4.0 resource pool. The cloud servers used by the service ensure resource isolation at the IaaS layer and guarantee exclusive computing resources plus exclusive storage resources in the resource combination.

Host Security

YI-MapReduce supports integration with the official eSurfing Cloud security service, supporting vulnerability scanning, security protection, application firewall, Cloud

Bastion Host (CBH), webpage tampering protection, etc. For cloud servers, the following security measures are provided:

- Security hardening of the operating system kernel
- Updating of the latest operating system patches
- Control of operating system permissions
- Management of operating system ports
- Protection against attacks on operating system protocols and ports
- Monitoring of cloud servers
- Protection against DDoS attacks
- Regular data backup
- Enhancement of login password strength

Application Security

We assure the stable performance of big data services through these measures:

- Identity Verification and Authentication
- Web Application Security
- Access Control
- Security Auditing
- Password Security

Data Security

The following measures are applied to ensure the confidentiality, integrity, and availability of vast user data.

- Disaster Recovery: Current data of the YI-MapReduce product is stored in the eSurfing Cloud Elastic Volume Service. The cloud hard disk implements a three-replica redundancy mechanism, ensuring high-level disaster recovery for the data of the upper-layer YI-MapReduce service.

- Backup: Current data of the YI-MapReduce product is stored in the eSurfing Cloud Elastic Volume Service, which features regular backup and data recovery.

1.4.3. Easy Access to Component WEB UI

Each big data component has its own WEB UI page for system management. However, due to network isolation, these pages are not easily accessible. For example, traditional methods to access the Kibana WEB UI page require the creation of an ECS and remote login to the component UI through ECS. This process is cumbersome and not user-friendly, particularly for big data novices.

YI-MapReduce provides easy access to the open-source component UI based on the internal network IP address.

Procedure

1. Log in to the YI-MapReduce console, click the name of the active cluster, and go to the cluster details page.

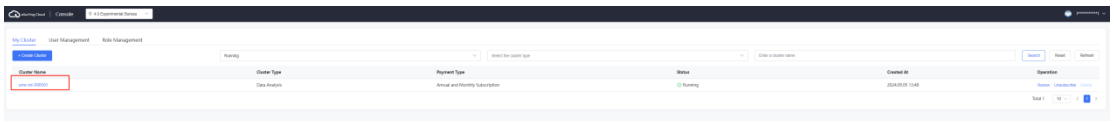

2. Click on **Access Link & Port**.

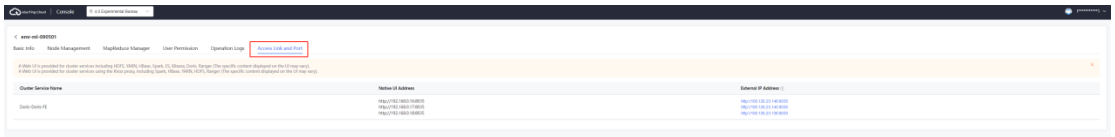

3. Cluster Service Name: Cluster services that have a WEB UI include Doris, Elasticsearch, etc. For more details, please refer to Open-Source Component Websites.

4. Native UI Address: After the cluster is deployed with services that have a WEB UI, the Native UI Address column will display relevant cluster service information by default. The native UI address is displayed as "Intranet IP Address: Port Number".

5. Obtain External Access IP Address: The external IP address is displayed as "External Address: Port Number". To successfully access the open-source component UI, you need to open the port on the security group page and set up inbound and outbound rules. Additionally, the external IP of the node where the cluster service is deployed needs to be enabled. For how to bind the external IP, please refer to Bind/Unbind EIP.

6. Access the WEB UI Address: Once the security group and external IP are successfully set up, the External IP Address column will display an accessible external IP address. You can easily access the open-source component UI by visiting the external IP address.

✧ **Note:**

The security group port needs to be opened to gain access. Please close it promptly after use to prevent security risks.

1.4.4. Enhanced Reliability

YI-MapReduce, a comprehensive open-source big data platform product offered by eSurfing Cloud , comes with a basic platform for big data storage and computing, and a platform for big data O&M management. Its main focus is on improving and tuning the reliability and performance of big data components.

System Reliability

HA Implementation for Master Nodes

In the open-source version of Component , data and computing nodes are designed based on a distributed system, so the failure of a single node does not impact the overall operation of the system. However, the potential single point of failure

occurring in master nodes operating in a centralized mode becomes the weak point in the overall system reliability.

YI-MapReduce, an eSurfing Cloud big data platform, provides a dual-machine-like mechanism for the master nodes of all components, including Doris FE, Doris BE, Elasticsearch Nodedata, and more. All adopt active/standby or load-sharing configurations, effectively avoiding the impact of single-point failure scenarios on system reliability.

Reliability Assurance in Exception Scenarios

Through the use of reliability analysis, measures to handle software and hardware exception scenarios are sorted out, enhancing the reliability of the system.

- In the event of an unexpected power outage, data reliability is guaranteed. Whether it's an accidental power outage of a single node or an unexpected power outage for the entire cluster, the system can recover the service normally after power is restored. Key data will not be lost unless the hard drive medium is damaged.
- Sub-health detection of hard drives and fault handling will not cause actual impact on the service.
- The system will automatically handle file system failures and recover affected services.
- The system will automatically handle process and node failures and recover affected services.
- The system will automatically handle network failures and recover affected services.

Node Reliability

Monitoring of Operating System Health Status

It routinely gathers data regarding the utilization rates of operating system hardware resources, including the usage status of CPU, memory, hard drive, and network resources.

Process Health Status Monitoring

YI-MapReduce provides health checks on the status of business instances and the health indicators of the processes within business instances, enabling users to promptly perceive the health status of the processes.

Automatic Hard Drive Troubleshooting

YI-MapReduce, an eSurfing Cloud big data platform, has enhanced the open source version. It is capable of monitoring the status of hard drives and file systems on each node. If an exception occurs, remove the relevant partitions immediately from the storage pool. If a hard drive returns to normal operation (typically because the user has replaced the faulty hard drive with a new one), the new hard drive will also be integrated into business operations. This significantly reduces the load on O&M personnel as hard drive replacements due to failure can be performed online. At the same time, users have the option to establish hot standby disks, which greatly decreases the repair time for faulty hard drives, thereby enhancing system reliability.

Node Disk LVM Configuration

YI-MapReduce, an eSurfing Cloud big data platform, supports configuring multiple disks into LVM (Logic Volume Management), and organizing multiple disks into a single logical volume group. The configuration of LVM avoids uneven disk usage,

maintaining balanced usage across all disks, which is particularly crucial for components such as HDFS and Kafka that can leverage multiple disk capabilities. Furthermore, LVM supports disk expansion without the need for remounting, thereby preventing service disruptions.

Data Reliability

YI-MapReduce, a big data platform product of eSurfing Cloud, utilizes the anti-affinity node group and placement group features provided by the Elastic Cloud Server (ECS). In conjunction with the rack awareness capability of Hadoop, it replicates data across multiple physical hosts, effectively preventing data loss due to physical hardware failures.

1.4.5. Cluster Management

Create Cluster

YI-MapReduce supports creating clusters with varied types, component scopes, numbers of each type of nodes, virtual machine specifications, availability zones, VPC networks, and authentication information. Based on the user's selected cluster type, version, and node specifications, YI-MapReduce aids in automatically completing the installation, deployment, and parameter optimization of enterprise-level big data platforms.

YI-MapReduce provides users with a fully controllable big data cluster. Users can set the login method for the virtual machine during its creation. The resources of the created YI-MapReduce cluster are fully allocated to the user.

YI-MapReduce Cluster Types Include Data Analysis and Elastic Search Clusters

1. Data Analysis Cluster: Apache Doris - An open-source MPP architecture OLAP analysis engine that supports sub-second data queries and multi-table joins.
2. Elastic Search Cluster: Delivers low-cost, high-performance, and reliable search and analysis service capabilities for structured/unstructured data.

YI-MapReduce Cluster Consists of Master Nodes, Core Nodes, and Task Nodes

1. Master nodes: The management nodes in a cluster, responsible for normal cluster scheduling, primarily deploying processes such as Doris FE and Elasticsearch Master. The cluster defaults to High Availability (HA) mode, with a fixed number of 3 master nodes. This type of nodes support upgrades to the configuration of these nodes, allowing them to manage larger clusters.
2. Core nodes: The computation and storage nodes in a cluster, primarily deploying processes such as Doris BE、Elasticsearch NodeData. They support configuration upgrades to meet the needs of storage data volume or computation volume expansion. The upgrade process will not affect the normal operations of the current cluster.
3. Task nodes: Purely computation nodes in a cluster. They compute data, but do not store data.

Configuration Upgrade

If the specifications (vCPU and memory) of your master or core node ECS instances fail to meet your business needs, you can enhance these ECS instance specifications through the feature of configuration upgrade.

Cluster O&M Management

MR Manager provides resource overview and supports O&M management, and configuration of cluster services, hosts, tenants and resources, monitoring and alarms.

1. Resource Overview: This feature displays the CPU, memory, network information, etc., for all hosts in this cluster, which includes metrics such as CPU usage, disk usage, memory usage, and network sending rate.
2. Cluster Service: This feature displays all the cluster services in the current cluster and lists them in a list view by component type. It also supports one-click start and one-click stop for all cluster services at the cluster service list.
3. Host: This feature, by default, displays a list of all hosts in the current cluster, allowing you to view the host information of the current O&M platform. It also lets you view role instance allocation and alarm history on each host.
4. Monitoring and Alarm: This feature supports metric queries and alarm history. The metric query feature allows for the querying of monitoring metrics at both the role instance level and host level. This supports the graphing of metric results, enabling users to intuitively understand changes in metrics. The alarm history feature allows for the querying of alarm content at the cluster service level, role instance level, and host level.
5. O&M and Configuration: This feature supports pipeline history, configuration management, configuration history, and configuration synchronization history. Pipeline history displays the operation history of all pipelines, along with the operator involved. Configuration management supports viewing the configuration files of different cluster services and performing operations such as adding, modifying, and deleting configurations. Configuration history supports viewing the configuration content of different versions of configuration files and comparing content between different versions. Configuration synchronization history supports viewing the configuration synchronization history of different environments, the operators involved in configuration synchronization, and the details of the configuration synchronization.

1.5. Application Scenarios

Big data pervades our daily lives. Fields such as finance, transportation, Internet, healthcare, energy, and government departments can all leverage the YI-MapReduce service for big data processing.

1.5.1. Offline Data Analysis Scenario

Import or incorporate massive data into the OLAP analysis engine in the form of external tables. For example, utilize the efficient, real-time, and flexible data analysis capabilities provided by Trino.

This caters to a series of business scenarios such as user profiling, crowd selection, location services, BI reports, and business analysis.

1.5.2. Online Query Scenario

Perform online analysis on PB-level structured, semi-structured, or unstructured data, which are generated from Web and mobile applications.

This allows customers' Web applications or data visualization products to access and display the analysis results in real time.

1.6. Terms

1.6.1. Node

A node in the YI-MapReduce cluster refers to a cloud server. The different types of nodes and their functions are as shown below.

Node Type	Feature
Master Node	A master node of the YI-MapReduce cluster that deploys control processes such as Doris FE and Elasticsearch Master.. The Master node group employs anti-affinity technology by default, distributing

	virtual machines across different physical machines to ensure high availability of services.
Core Node	A worker node of the YI-MapReduce cluster that stores and computes data.

1.6.2. Doris

Apache Doris is a high-performance, real-time analytical database underpinned by an MPP architecture. It stands out for its speed and ease-of-use, delivering query results from massive datasets in a sub-second response time. Doris supports both high-concurrency point query scenarios and high-throughput complex analysis scenarios. This makes it an excellent choice for report analysis, ad-hoc querying, unified data warehouse construction, and data lake federated query acceleration, among other uses. On top of Doris, users can build various applications such as user behavior analysis, AB testing platforms, log retrieval analysis, user profile analysis, and order analysis.

1.6.3. Elasticsearch

Elasticsearch is an open-source, highly scalable distributed full-text search engine that can store and retrieve data almost in real time. It centralizes your data, providing swift search, finely tuned relevance, and robust analytics capabilities.

1.7. Relationships with Other Products

1.7.1. Relationships with Other Products

Service Name	Relationships Between YI-MapReduce and Other Products	Main Interactive Features

Virtual Private Cloud (VPC)	YI-YI-MapReduce clusters are created within the subnet of the Virtual Private Cloud (VPC). VPC provides a secure, isolated network environment for user's YI-YI-MapReduce clusters through logical isolation.	Create VPC and subnets
Elastic Cloud Server (ECS)	YI-MapReduce service uses Elastic Cloud Server (ECS) as the nodes of the cluster, with each ECS being a node within the cluster.	Prepare runtime environment for cluster creation
Relational Database Service MySQL Version (RDS MySQL)	The Relational Database Service MySQL Version (RDS MySQL) is used to store metadata for YI-YI-MapReduce clusters.	Configure data connection
Identity and Access Management (IAM)	Identity and Access Management (IAM) provides authentication functions for YI-MapReduce.	Create IAM sub-users and sync user information to the console and LDAP users

1.8. Permission Management

If you need to set different access permissions for employees in your organization to the YI-MapReduce service resources that have already been created, to achieve permission isolation among different employees, you can use the Identity and Access Management (IAM), role management, and user permissions for fine-grained permission management. The IAM service provides user identity authentication. The role management service allocates menu permissions for different roles. The user permission service assigns different roles to different IAM

users, providing cluster access control. These features can help you securely control resource access.

If your existing account meets your requirements and there is no need to create separate IAM users for permission management, you can skip this section. It will not affect your use of other YI-MapReduce service functions.

1.8.1. YI-MapReduce Permission Description

By default, IAM users created by the administrator do not have any permissions. They need to be assigned actual roles to obtain corresponding permissions. This process is known as authorization. After authorization, users can operate cloud services based on the granted permissions.

Role management defines menu permission roles, which apply to all clusters; while user permissions define different roles for different IAM users, which apply to a single cluster. Therefore, role management is conducted under the administrator account, while user permissions are operated within a single cluster.

Role Management: As shown below, it includes all menu functions of YI-MapReduce. You can assign different menu functions to different roles based on personal needs.

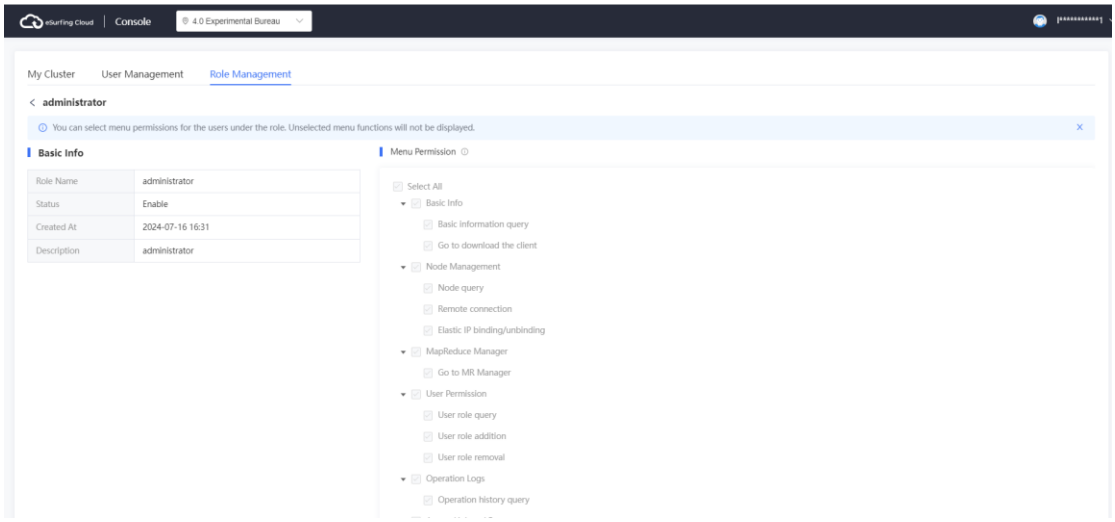

User Permissions: As shown below, users can add access permissions for different users for a specific cluster. Users can be added individually or in batches.

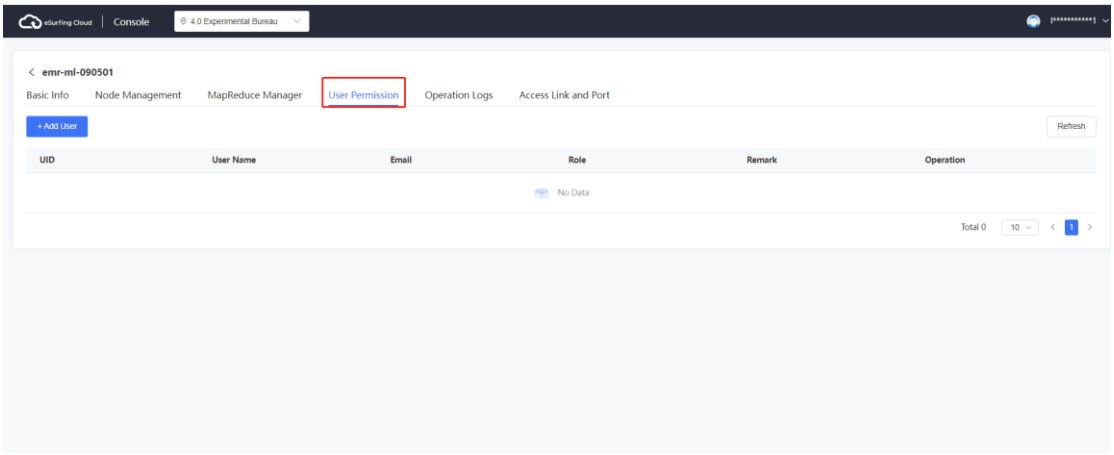

1.9. Restriction and Limitations

Before using YI-MapReduce, you need to read and understand the following usage rules carefully.

1. Clusters must be created within a VPC subnet.
2. When creating a cluster, select an existing security group from the dropdown box. After the cluster is created, do not delete or change the used security group unless necessary, as this may cause cluster exceptions and affect its usage.
    a. Do not grant unnecessary permissions of the security group used by the cluster to avoid malicious access.
    b. When creating a security group, the ports in the outbound and inbound direction rules need to be opened, and the authorization policy cannot be set to "Reject", otherwise the cluster deployment and service startup may be affected.
3. When deploying Hive/Ranger components, the database configured with metadata and the cluster need to be under the same VPC, and should be connected through an Intranet IP address.
4. Please plan the disk of the cluster nodes according to business needs. If you need to store a large amount of business data, increase the number of EVS or storage space to prevent insufficient storage space from affecting the normal operation of nodes.

5. Cluster nodes are only used to run the YI-MapReduce cluster. It is recommended to deploy other client applications and user business programs on separate ECS.

1.10. Security

1.10.1. Authentication and Access Control

Authentication

YI-MapReduce supports the Kerberos security protocol, uses LDAP as the account management system through the synchronization of IAM accounts, and performs security authentication of account information through the Kerberos service.

For a detailed introduction to the principles of Kerberos security authentication and the authentication mechanism, please refer to Principles of Security Authentication and Authentication Mechanism.

Access Control

YI-MapReduce provides a role-based access control model.

Role-Based Access Control

Built on a unified authentication system for users and roles, YI-MapReduce adheres to the account/role RBAC (Role-Based Access Control) model. This model enables permission management through roles and batch authorization management for users. It offers single sign-on capabilities, unifying the management and authentication of system users and component users.

2. Billing

2.1. Resource Node

Currently, first-class resource node is available in the **China(Hongkong2)** resource pool.

In the China(Hongkong2) resource pool, you can choose from different availability zones: AZ 1 and AZ 2.

2.2. Product Specifications

2.2.1. Cloud Server

General-purpose

Specification Name	Cores (vCPU)	Memory (GB)
s7.2xlarge.4	8	32
s7.4xlarge.2	16	32
s7.4xlarge.4	16	64
s7.8xlarge.2	32	64
s7.8xlarge.4	32	128

General Computing-plus

Specification Name	Cores (vCPU)	Memory (GB)
c7.2xlarge.4	8	32

c7.3xlarge.4	12	48
c7.4xlarge.2	16	32
c7.4xlarge.4	16	64
c7.6xlarge.2	24	48
c7.6xlarge.4	24	96
c7.8xlarge.2	32	64
c7.8xlarge.4	32	128
c7.12xlarge.2	48	96
c7.12xlarge.4	48	192
c7.16xlarge.2	64	128
c7.16xlarge.4	64	256
c7.24xlarge.2	96	192
c7.24xlarge.4	96	384

Memory-optimized

Specification Name	Cores (vCPU)	Memory (GB)
m7.2xlarge.8	8	64
m7.3xlarge.8	12	96
m7.4xlarge.8	16	128
m7.6xlarge.8	24	192
m7.8xlarge.8	32	256

m7.12xlarge.8	48	384
m7.16xlarge.8	64	512
m7.24xlarge.8	96	768

2.2.2. EVS

The product specifications for EVS include: general IO (SATA), high IO (SAS), ultra high IO (SSD), and general purpose SSD.

2.3. Billing Mode

2.3.1. Billing Items

The cost of purchasing an YI-MapReduce cluster includes two parts:

- YI-MapReduce service management fee
- IaaS infrastructure resource costs (cloud server, EVS)

2.3.2. Billing Method

YI-MapReduce currently supports an annual or monthly billing method.

- Annual/Monthly Subscription: The cluster fee is fully paid upfront according to the duration of cluster purchase. The minimum duration is one month, and the actual purchasable duration is subject to information displayed on the page.

2.3.3. Expiry

Upon expiration, an annual/monthly cluster enters a retention period. During this period, attempts to operate on the cluster through the YI-MapReduce management console will be unsuccessful, and related interfaces will be inaccessible. O&M services such as automated monitoring or alarm functions will also be halted. If you do not renew the service by the end of the retention period, the cluster will end its service, and the data housed in the system will be permanently deleted.

2.4. Price

The cost of the YI-MapReduce cluster consists of the YI-MapReduce service management fee and ECS product fees. The ECS product fees include the costs of Elastic Cloud Server.

the price of YI-MapReduce includes the following parts:

2.4.1. Cloud Server Price

General-purpose

Specification Name	Cores (vCPU)	Memory (GB)	Monthly Subscription (USD)	1-Year Subscription (USD)
s7.2xlarge.4	8	32	250.480532	3,005.766381
s7.4xlarge.2	16	32	353.648795	4,243.785546
s7.4xlarge.4	16	64	500.961064	6,011.532762
s7.8xlarge.2	32	64	707.297591	8,487.571092
s7.8xlarge.4	32	128	1,001.922127	12,023.065524

Note: 15% off for 1 year, 30% off for 2 years, 50% off for 3 years, 55% off for 4 years, and 60% off for 5 years.

General Computing-plus

Specification Name	Cores (vCPU)	Memory (GB)	Monthly Subscription (USD)	1-Year Subscription (USD)
c7.2xlarge.4	8	32	272.236413	3,266.836958
c7.3xlarge.4	12	48	408.354620	4,900.255437

c7.4xlarge.2	16	32	426.288827	5,115.465925
c7.4xlarge.4	16	64	544.472826	6,533.673916
c7.6xlarge.2	24	48	639.433241	7,673.198888
c7.6xlarge.4	24	96	816.709239	9,800.510874
c7.8xlarge.2	32	64	852.577654	10,230.931851
c7.8xlarge.4	32	128	1,088.945653	13,067.347832
c7.12xlarge.2	48	96	1,278.866481	15,346.397776
c7.12xlarge.4	48	192	1,633.418479	19,601.021747
c7.16xlarge.2	64	128	1,705.155308	20,461.863701
c7.16xlarge.4	64	256	2,177.891305	26,134.695663
c7.24xlarge.2	96	192	2,557.732963	30,692.795552
c7.24xlarge.4	96	384	3,266.836958	39,202.043495

Memory-optimized

Specification Name	Cores (vCPU)	Memory (GB)	Monthly Subscription (USD)	1-Year Subscription (USD)
m7.2xlarge.8	8	64	334.658971	4,015.907648
m7.3xlarge.8	12	96	501.988456	6,023.861471
m7.4xlarge.8	16	128	669.317941	8,031.815295
m7.6xlarge.8	24	192	1,003.976912	12,047.722943
m7.8xlarge.8	32	256	1,338.635883	16,063.630590

m7.12xlarge.8	48	384	2,007.953824	24,095.445885
m7.16xlarge.8	64	512	2,677.271765	32,127.261181
m7.24xlarge.8	96	768	4,015.907648	48,190.891771

2.4.2. EVS Price

Product Specifications	Monthly Subscription (USD/GB)	1-Year Subscription (USD/GB)
Common IO (SATA)	0.1270128	10.800000
High I/O (SAS)	0.08890896	7.560000
Ultra-high IO (SSD)	0.206042987	17.520000
EXTREME-SPEED SSD	0.451601067	38.400000

2.4.3. YI-MapReduce Subscription-based Service Management Price

Service management cost = ECS product cost * 0.12.

2.5. Change

2.5.1. Change Specifications

Before activating YI-MapReduce, a variety of node choices are available. You can select the appropriate master and core node instances based on business needs. After launching the cluster, YI-MapReduce offers the following methods for configuration changes.

Configuration Upgrade: If the specifications (vCPU and memory) of your master or core node ECS instances fail to meet your business needs, you can enhance these

ECS instance specifications through the feature of configuration upgrade. Please refer to Configuration Upgrade.

3. Getting Started

3.1. How to Use YI-MapReduce

YI-MapReduce, an eSurfing Cloud big data platform, provides a service that deploys and manages bigdata systems in the cloud environment. It facilitates the one-click deployment of a variety of big data clusters including Doris and Elasticsearch among others. YI-MapReduce offers tenants a fully manageable, enterprise-level big data cloud service cluster, thus simplifying the operation of various big data components such as Doris and Elasticsearch.

YI-MapReduce is user-friendly. By interconnecting multiple computers in a cluster, you can run various tasks, and process or store massive data at the petabyte scale.

The basic workflow of YI-MapReduce is as follows:

1. Create a Cluster: Visit the YI-MapReduce Service Subscription page on the official eSurfing Cloud website, select the business scenario needed and create a cluster. Users can streamline the combination of big data clusters based on the business scenario, which can be used for elastic search, data analysis, and other business requirements. After choosing the business scenario, users can further detail the specifications of the ECS instances, the number of instances, the type of data disk (with various options including common IO, high IO, ultra-high IO, and general purpose SSD), and optional services (Logstash etc.) within the cluster.
2. Manage the Cluster: Once the cluster subscription and deployment are complete, YI-MapReduce provides an enterprise-level big data cluster unified management platform, MapReduce Manager. This platform assists users in promptly grasping the health status of cluster services and hosts, acquiring key system information through graphic metric monitoring, modifying service attribute configurations to meet real business performance requirements, and performing one-click start-stop operations for clusters, services, and role instances, etc.

3. Create YI-MapReduce Operation User: The master user can create sub-users through the Identity and Access Management (IAM). With precise permission management, different role permissions can be assigned to sub-users through the YI-MapReduce console.

4. Unsubscribe: If you decide to stop using the existing cluster, you can unsubscribe it through Order Management on the official eSurfing Cloud website.

5. Renew: If your cluster is nearing its expiry, you can prolong the cluster usage period by renewing it.

6. Configuration Upgrade: If the specifications (vCPU and memory) of your master or core node ECS instances fail to meet your business needs, you can enhance these ECS instance specifications through the feature of configuration upgrade.

3.2. Create Cluster

3.2.1. Procedure

1. Log in to the YI-MapReduce Management Center.

2. Click on **Create Cluster** to access the cluster creation page.

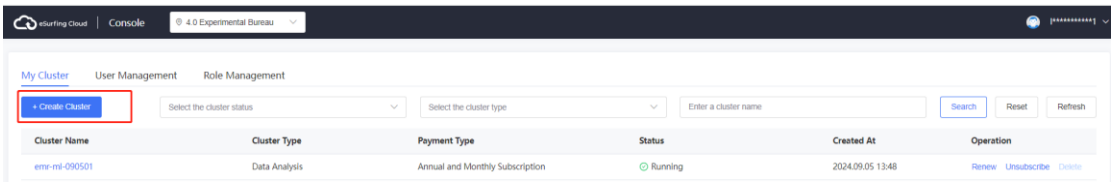

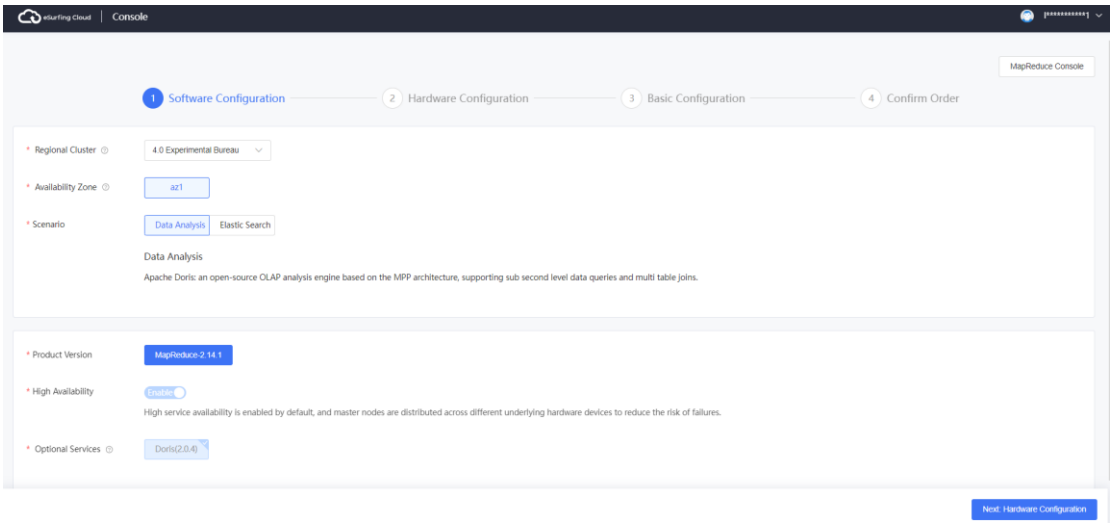

a. Regional Cluster: The default setting is suitable. It refers to the geographic location where the ECS instances of the cluster reside.

b. Availability Zone: The default setting is suitable. It refers to the selectable zones within the current region cluster.

c. Business Scenario: Choose as per your requirement.

d. Product Version: The latest version is selected by default.

e. High Service Availability: Enabled by default.

f. Optional Services: Choose the components required as per your needs. Essential components cannot be deselected.

g. Kerberos Authentication (if applicable): Enabled by default.

3. After the software configuration is complete, click **Next** to enter the hardware configuration page.

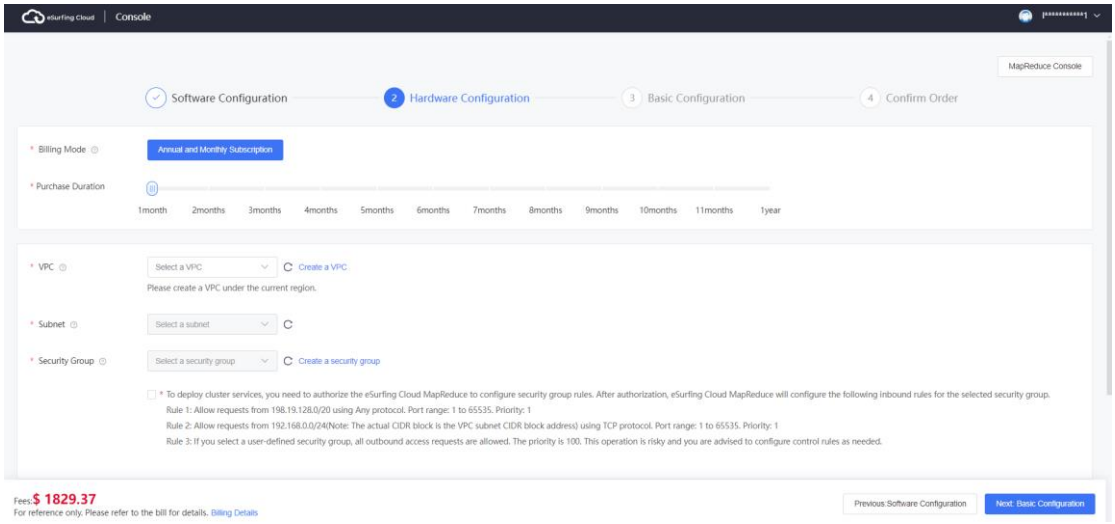

a. Billing Mode: Annual or monthly billing by default.

b. Purchase Duration: Drag the slider to choose a purchase duration (subject to the actual purchase duration during the public preview period).

c. Virtual Private Cloud: Different virtual private cloud (VPC) networks are logically isolated from each other. Select the VPC that suits your needs. If you don't have a VPC currently, click on **Create VPC** to go to the VPC creation page.

d. Subnet: After selecting a VPC, you can choose the subnet as needed.

e. Security Group: This configures network access control for the ECS instances within a cluster. When creating a security group, the ports in the outbound

and inbound direction rules need to be opened, and the authorization policy cannot be set to "Reject", otherwise the cluster deployment and service startup may be affected. Please refer to the Tips to complete the creation of the security group.

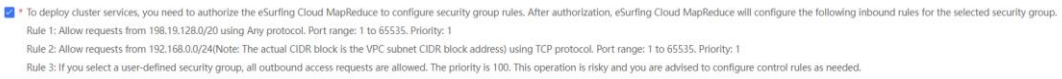

f. Node Group: Choose the specifications and quantity of the cluster nodes according to your needs.

4. After the hardware configuration is complete, click **Next** to enter the basic configurations page.

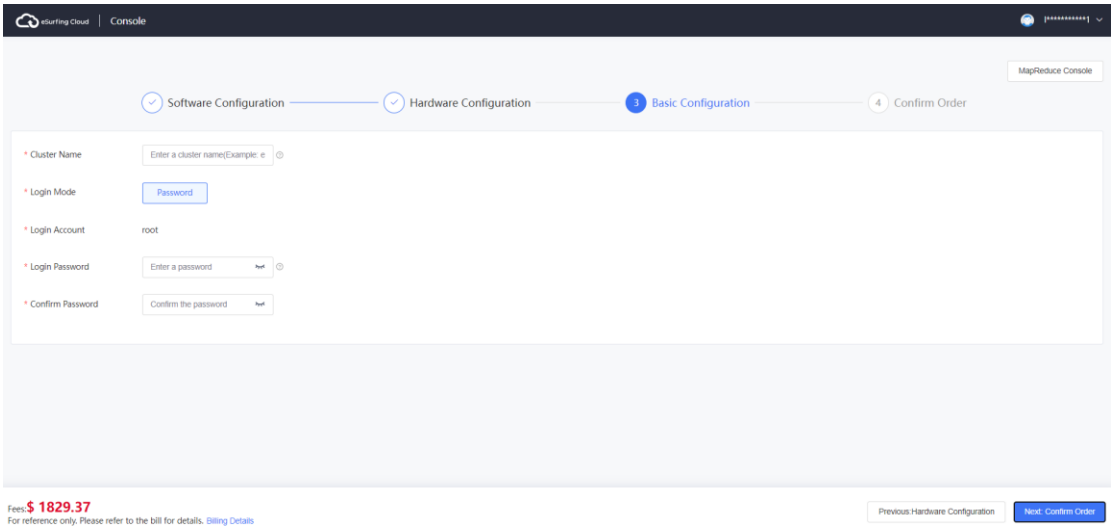

a. Cluster Name: A cluster name contains a maximum of 32 characters, including uppercase letters, lowercase letters, numbers, and special symbols.

b. Identity Credential: Refers to the administrator's password by default.

c. Login Account: The value is 'root' by default.

d. Login Password: Create a password that meets the requirements as prompted.

e. Confirm Password: This should be the same as the login password.

5. After completing the basic configurations, click the **Next** button at the bottom right corner to enter the order confirmation page.
Review the configuration details. Check "I have read and agree to the YI-MapReduce Service Agreement". Then, click **Purchase Now** to go to the

payment page. Upon successful payment, the message "Activated Successfully" will appear, and you will then be redirected to the **Management Console**.

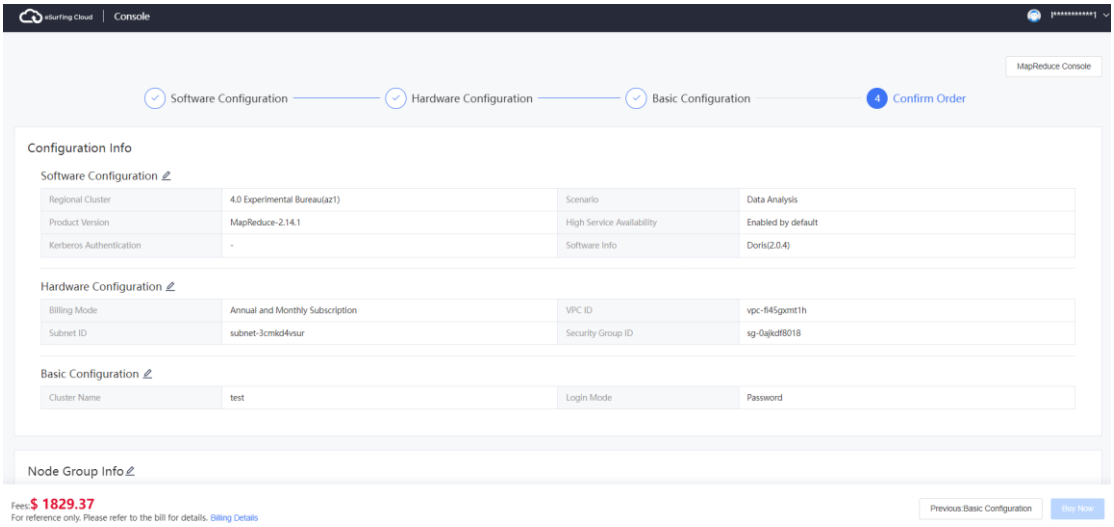

6. On the Management Console, you can view details of the cluster. Creating a cluster takes some time, and the progress can be monitored from the cluster list. Please wait a while.

3.3. Delete Cluster

If the cluster is no longer needed after job execution, you can delete the YI-MapReduce cluster.

3.3.1. Background

Typically, a cluster is deleted 15 days after the cluster expiry if the user does not renew, or if the user initiates the unsubscription.

3.3.2. Procedure

1. Log in to the YI-MapReduce management console.
2. Select the cluster to delete and click the **Delete** button in the operation column.

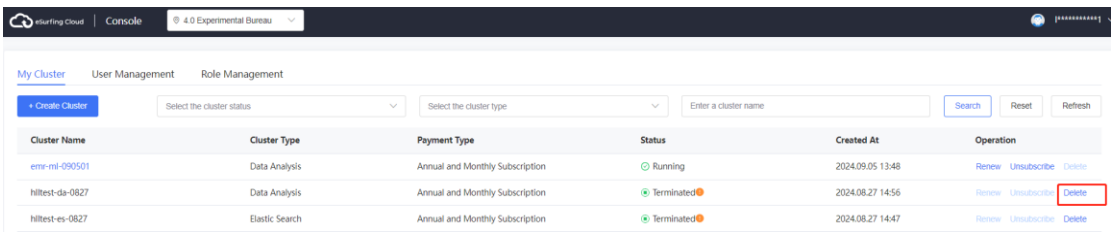

3. In the pop-up window, click **OK**.

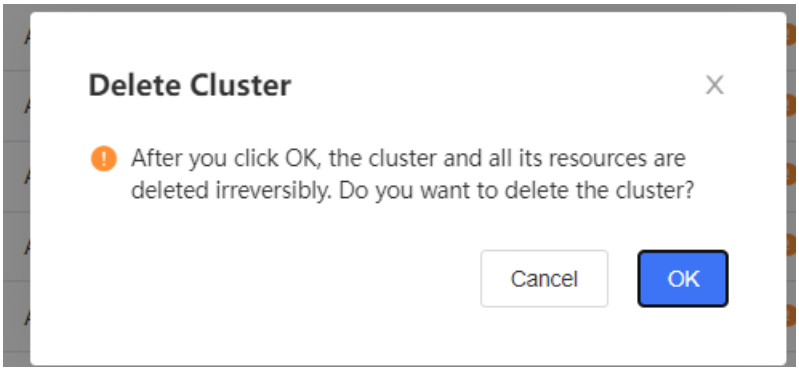

4. After deletion, the cluster will no longer appear in **My Cluster** and no further costs will be incurred.

✧ **Note:**

Deleting a cluster will result in the irreversible deletion of all resources under that cluster!

3.4. Manual Adjustment of Recommended Setting

3.4.1. Scenario

If the specifications (vCPU and memory) of the ECS instances within a node group fail to meet your business needs, you can enhance these ECS instance specifications through the feature of configuration upgrade. After upgrading the configuration, you need to manually adjust the configuration information of services such as Doris and Elasticsearch. This section guides you on how to modify configuration items on the Configuration Management page in YI-MapReduce Manager.

3.4.2. Before You Begin

You have created a cluster.

3.4.3. Procedure

In YI-MapReduce Manager, click on **O&M and Configuration**.

1. Click on **Configuration Management**.
2. Choose the **Selected Cluster Services**, click **Query**, and you can modify the configuration information on the current page.

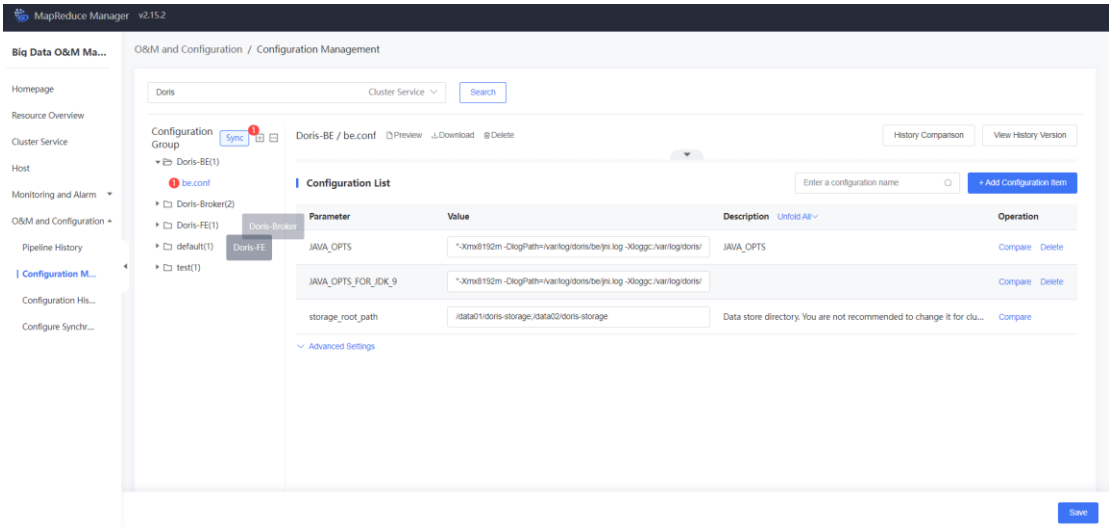

3.4.4. Suggestions for Configuration Modifications

1. Doris: After the configuration of the node where Doris resides is upgraded, it is recommended that Doris FE uses half of the node's memory.
2. Elasticsearch: For a single Elasticsearch (ES) node, it is recommended that the memory does not exceed 64G. After the configuration of the node where ES resides is upgraded, ES will automatically set the memory value according to the node conditions. Generally, no manual modification is needed, but the cluster needs to be restarted. If you wish to manually set it, you can modify the -Xms30g and -Xmx30g parameters in the jvm.options file and restart the cluster.

3.5. Table of Component Dependencies

3.5.1. Description of Component Dependencies

Component Name	Dependent Components Deployed in the Cluster
Doris	/

Elasticsearch	/
Kibana	Elasticsearch

4. User Guide

4.1. User Preparation

4.1.1. Create YI-MapReduce Operation User

If detailed permission management for your YI-MapReduce service is required, you can use the Identity and Access Management (IAM). This can be achieved by creating IAM sub-users and assigning them different role permissions in the YI-MapReduce console.

If your cloud account meets your requirements and there is no need to create separate IAM users, you can skip this section. It will not affect your use of other YI-MapReduce service functions.

This section guides you through the user authorization process. The procedure is shown as follows.

Procedure

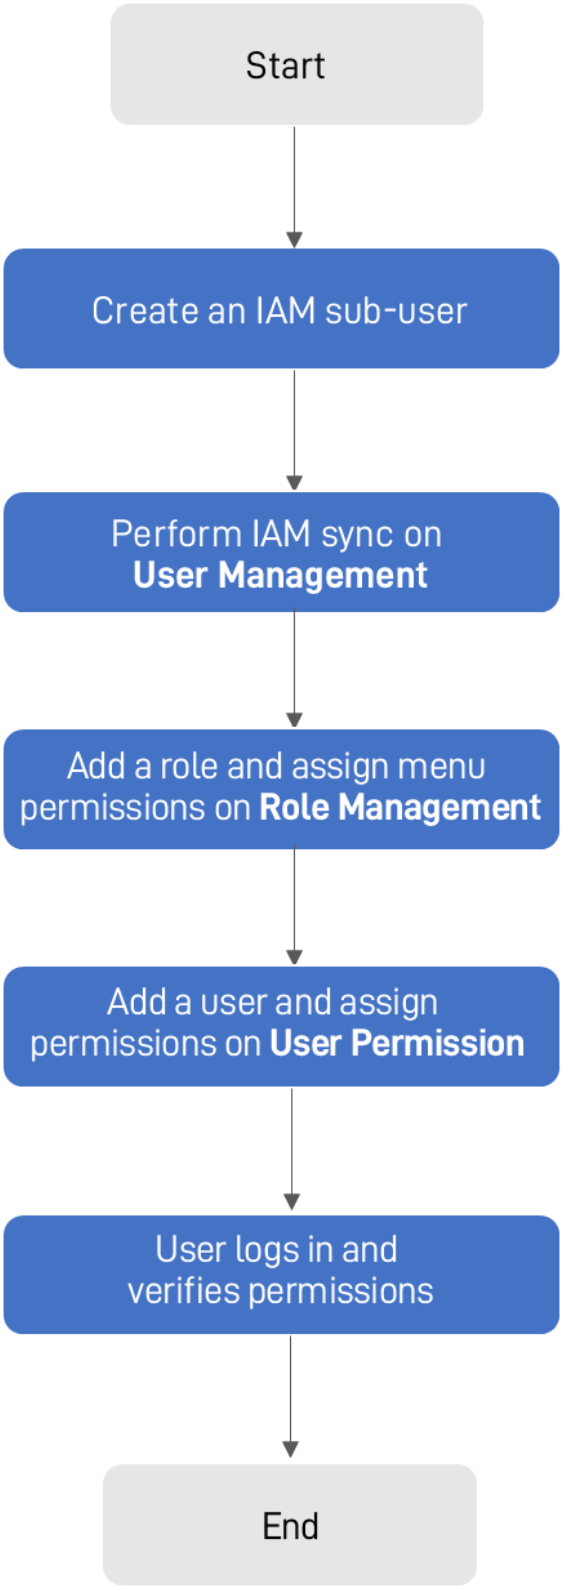
Start
Create an IAM sub-user
Perform IAM sync on **User Management**
Add a role and assign menu permissions on **Role Management**
Add a user and assign permissions on **User Permission**
User logs in and verifies permissions
End

1. Create an IAM sub-user. Go to the **User Management** page of the YI-MapReduce console. Click on the "?" next to the **IAM Sync** button. Then, click on **Access IAM** to enter the IAM user page. Click on **Create User** in the upper right corner and follow the instructions to create a sub-user.

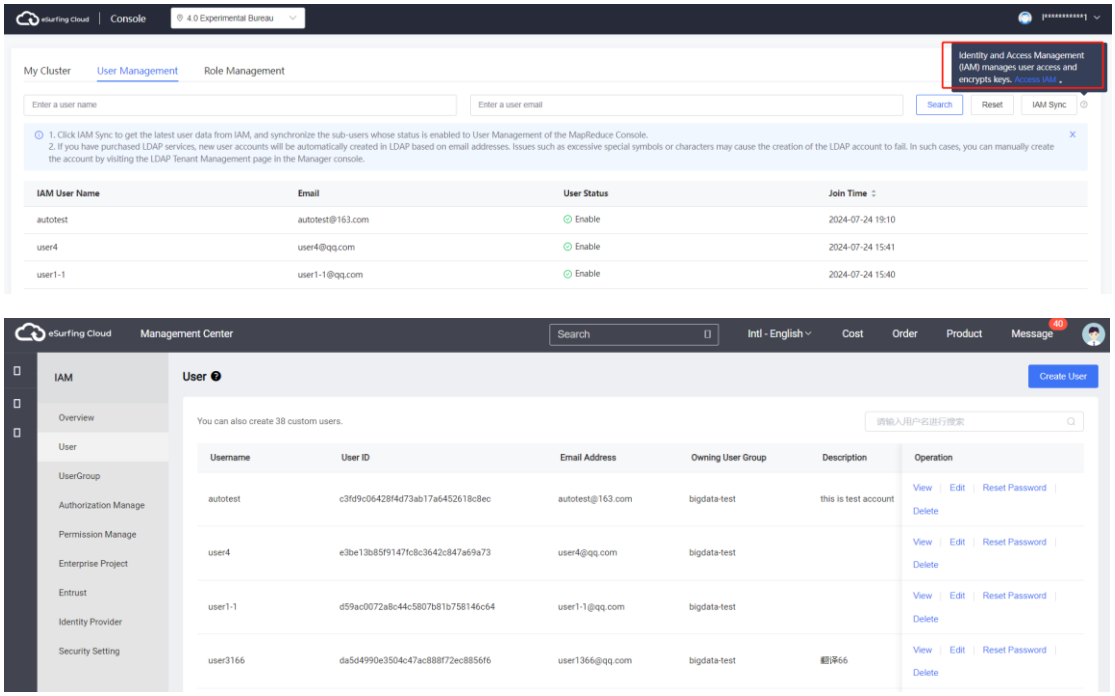

2. Click **IAM Sync** on the **User Management** page of the YI-MapReduce console to sync IAM sub-user information with the **LDAP User** in the YI-MapReduce console and YI-MapReduce Manager.

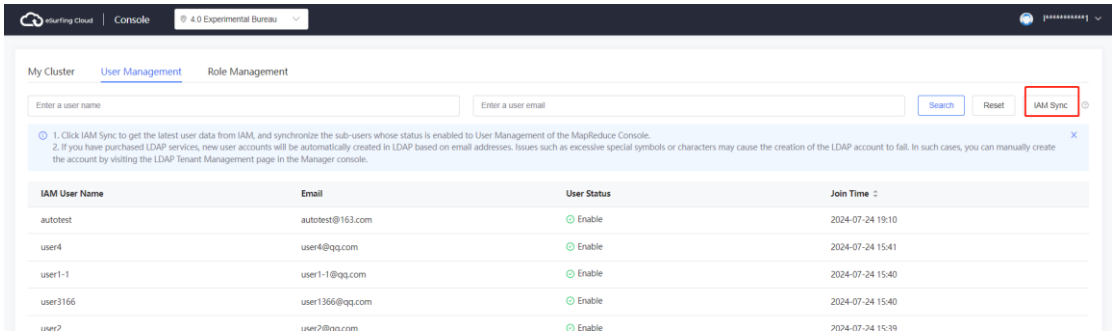

3. Add roles on the Role Management page of the YI-MapReduce console. Assign the appropriate menu function permissions to the role.

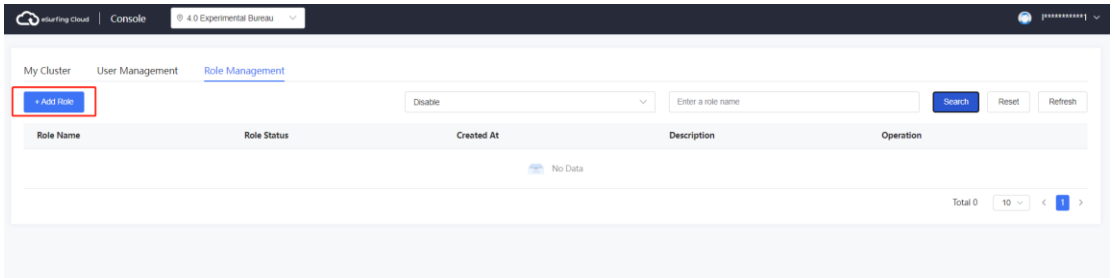

4. Add users on the User Permission page of the YI-MapReduce console. Assign the appropriate role permissions to the user.

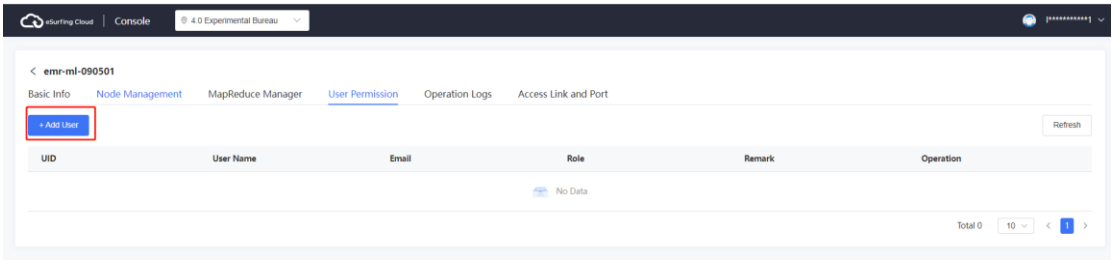

5. Then, IAM sub-users can log in to the YI-MapReduce console, access the authorized cluster, and verify permissions.

✧ **Description**

By default, IAM users created by the administrator do not have any permissions. They need to be assigned role permissions to obtain corresponding cluster permissions. This process is known as authorization. After authorization, users can operate cloud services based on the granted permissions.

4.1.2. IAM User Sync

IAM user synchronization involves syncing the created IAM sub-users into the YI-MapReduce system: one part is syncing IAM user information to the User Management service, including details such as IAM username, email, user status, add time, etc.; another part is syncing IAM users into the LDAP User service in the YI-MapReduce Manager, where LDAP usernames are denoted as "emr_email prefix". However, for the second part, the deployment, configuration, and start-up of the OpenLDAP cluster service under that cluster need to be completed; otherwise, it is not possible to sync into the LDAP User service.

Procedure

1. Click on the "?" button next to **IAM Sync** on the **User Management** page. Then, click **Access IAM** to go to the **Master Account and Member Accounts and Authorization Center** in IAM.

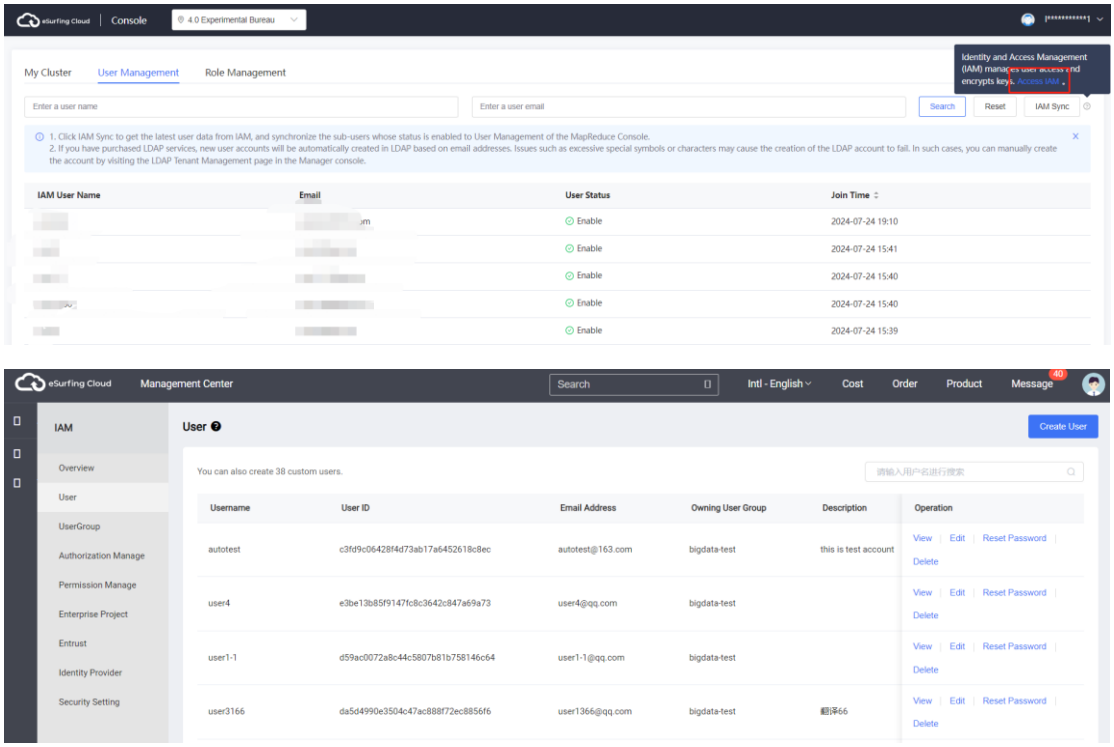

2. Click on **Create User** and follow the instructions to create an IAM user.

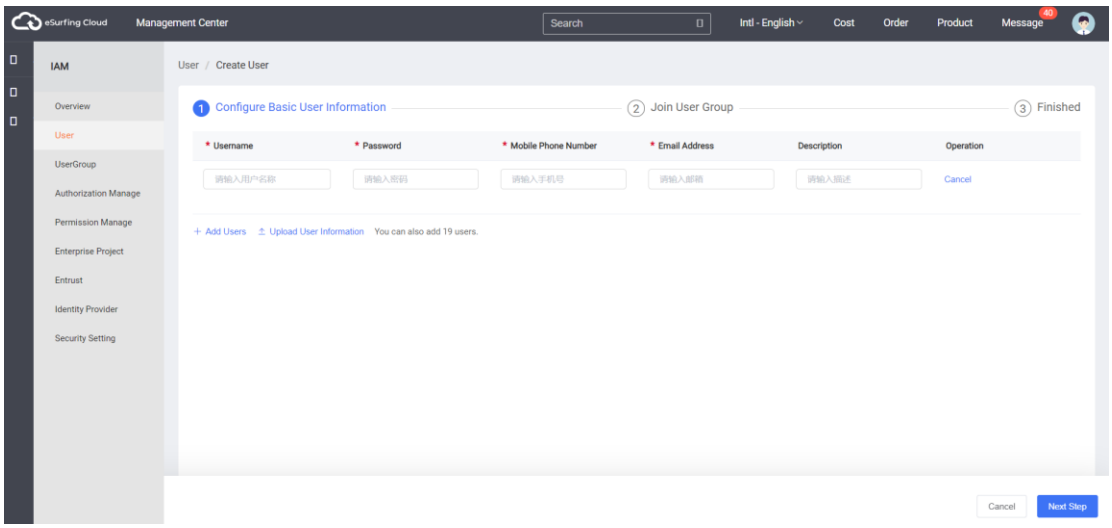

3. After creation, return to the User Management page in the YI-MapReduce console, click on **IAM Sync** to sync the IAM user information to the **User Management** page.

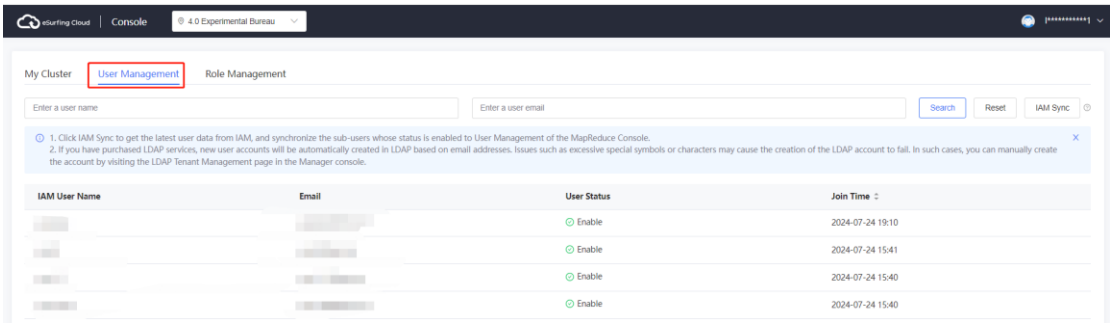

4.2. Configure Cluster

4.2.1. Introduction to Creation Methods

Quick Create Data Analysis Cluster: This method facilitates a quick creation of a data analysis cluster. It uses an open-source MPP architecture OLAP analysis engine, supports sub-second data queries and multi-table joins, assisting you in conducting data analysis tasks more efficiently.

Quick Create Elastic Search Cluster: This method facilitates the quick creation of an elastic search cluster. Through distributed computing and storage systems, it provides low-cost, high-performance, and reliable retrieval and analysis services for structured/unstructured data, assisting you in searching large data sets more conveniently.

4.2.2. Quick Create Cluster

Quick Create Data Analysis Cluster

A data analysis cluster uses Apache Doris, an open-source MPP architecture OLAP analysis engine that supports sub-second data queries and multi-table joins. Before creating a data analysis cluster, it is necessary to first create a Virtual Private Cloud.

Steps to Quick Create Data Analysis Cluster

1. Log in to the YI-MapReduce management console.
2. On the **My Cluster** page, click on **Create Cluster** to access the cluster creation page.
3. On the cluster creation page, configure the related settings of the cluster.
4. The software configuration page is displayed as shown below, with parameter descriptions as follows:

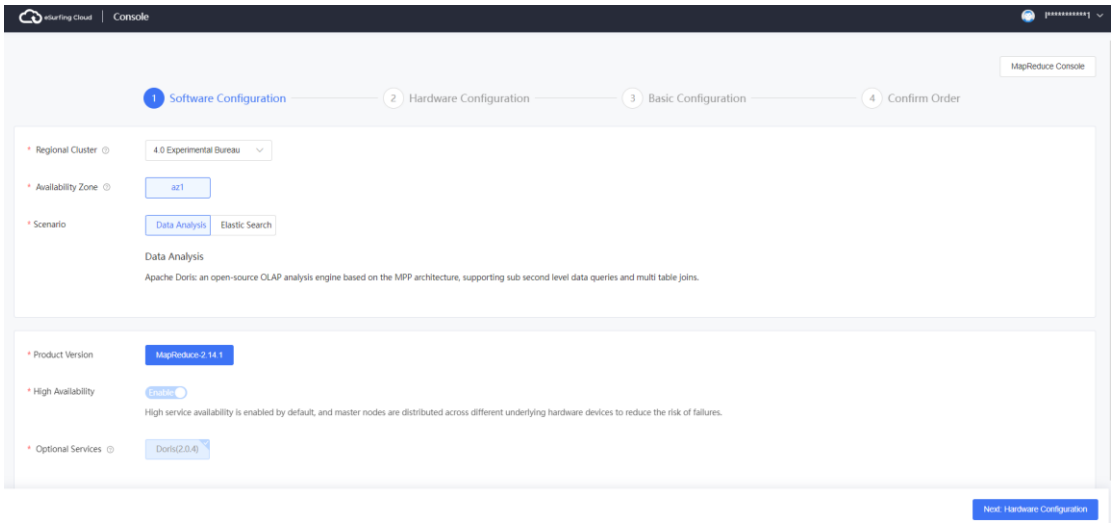

   a. Regional Cluster: This refers to the geographic location where the ECS instances of the cluster reside. Select the region and availability zone based on your requirements, or use the default value.

   b. Business Scenario: Select Data Analysis.

   c. Product Version: Choose the product version to use. The default value is suitable.

d. High Service Availability: High service availability is enabled for the YI-MapReduce service by default and cannot be disabled.

e. Optional Services: These consist of optional and required components, which are determined based on the business scenario. Based on your specific business scenario, you can select optional components. Under the data analysis cluster scenario, Doris is the only required component.

f. After the software configuration is complete, click **Next** to enter the hardware configuration page.

5. The hardware configuration page is displayed as shown below, with parameter descriptions as follows:

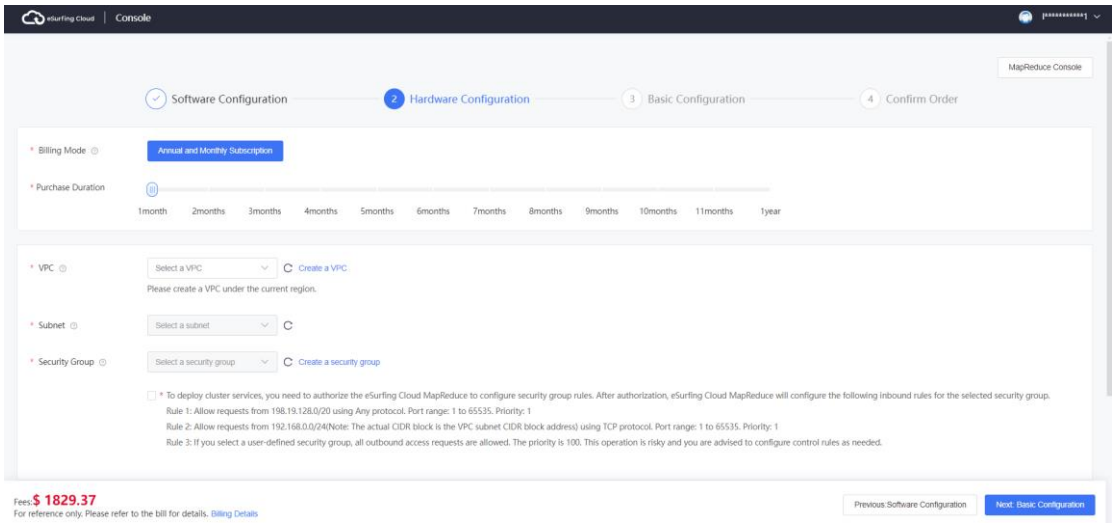

a. Billing Mode: Annual or monthly billing by default (based on the actual purchasable duration on the purchase page).

b. Purchase Duration: Drag the slider to choose a purchase duration.

c. Virtual Private Cloud: Different virtual private cloud (VPC) networks are logically isolated from each other. Select the VPC that suits your needs. If you don't have a VPC currently, click on **Create VPC** to go to the VPC creation page.

d. Subnet: After selecting a VPC, you can choose the subnet as needed.

e. Security Group: This configures network access control for the ECS instances within a cluster. When creating a security group, the ports in the outbound and inbound direction rules need to be opened, and the authorization policy cannot be set to "Reject", otherwise the cluster deployment and service startup may be affected. For details, see "How to use a custom

security group to create a YI-MapReduce cluster?" under Cluster Creation in the FAQs.

f. Node Group: Choose the specifications and quantity of the cluster nodes according to your needs, including the selection of node group types, option configurations, EVS parameters, and performance.

g. After completing the hardware configuration, the cost of the current configuration is calculated in the bottom left corner. Click the **Next** button in the bottom right corner to proceed to the basic configurations page.

6. The basic configurations page is displayed as shown below, with parameter descriptions as follows:

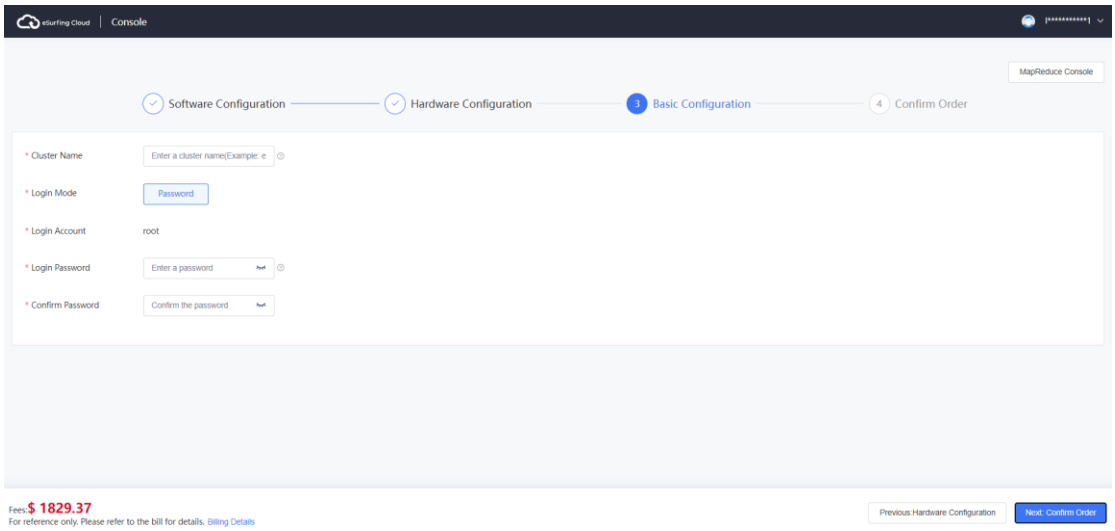

a. Cluster Name: A cluster name contains a maximum of 32 characters, including uppercase letters, lowercase letters, numbers, and special symbols.

b. Identity Credential: The default password.

c. Login Account: The value is 'root' by default.

d. Login Password: Create a password that meets the requirements as prompted.

e. Confirm Password: This should be the same as the login password.

f. After completing the basic configurations, click the **Next** button at the bottom right corner to enter the order confirmation page.

7. The Confirm Order page is as shown below, with descriptions for the configuration info as follows:

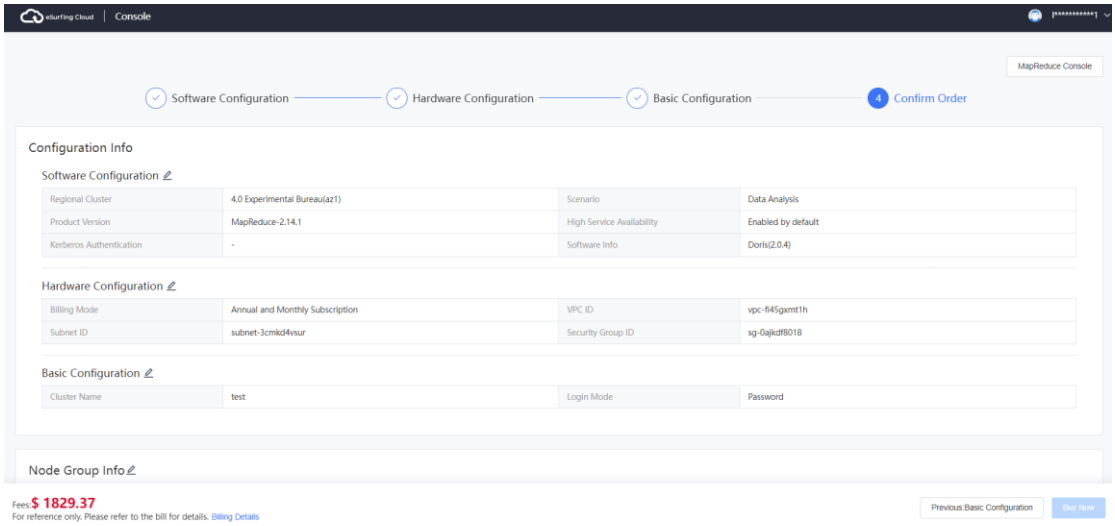

a. Software Configuration: Check the content for any errors. Click **Edit**

button to go back to the Software Configuration page and modify cluster information.

b. Hardware Configuration: Check the content for any errors. Click the

Edit button to go back to the Hardware Configuration page and modify cluster information.

c. Basic Configurations: Check the content for any errors. Click the

Edit button to go back to the Basic Configurations page and modify cluster information.

d. Node Group Info: Check the content for any errors. Click the

Edit button to go back to the Hardware Configuration page and modify cluster information.

e. Check "I have read and agree to the relevant agreements" and the **Purchase Now** button will become highlighted. Click it to proceed with payment. After purchase, you can rent the cluster with the current configuration.

Quick Create Elastic Search Cluster

Elastic search cluster delivers low-cost, high-performance, and reliable search and analysis service capabilities for structured/unstructured data. Before creating an elastic search cluster, it is necessary to first create a Virtual Private Cloud.

Steps to Quick Create Elastic Search Cluster

1. Log in to the YI-MapReduce management console.
2. On the **My Cluster** page, click on **Create Cluster** to access the cluster creation page.
3. On the cluster creation page, configure the related settings of the cluster.
4. The software configuration page is displayed as shown below, with parameter descriptions as follows:

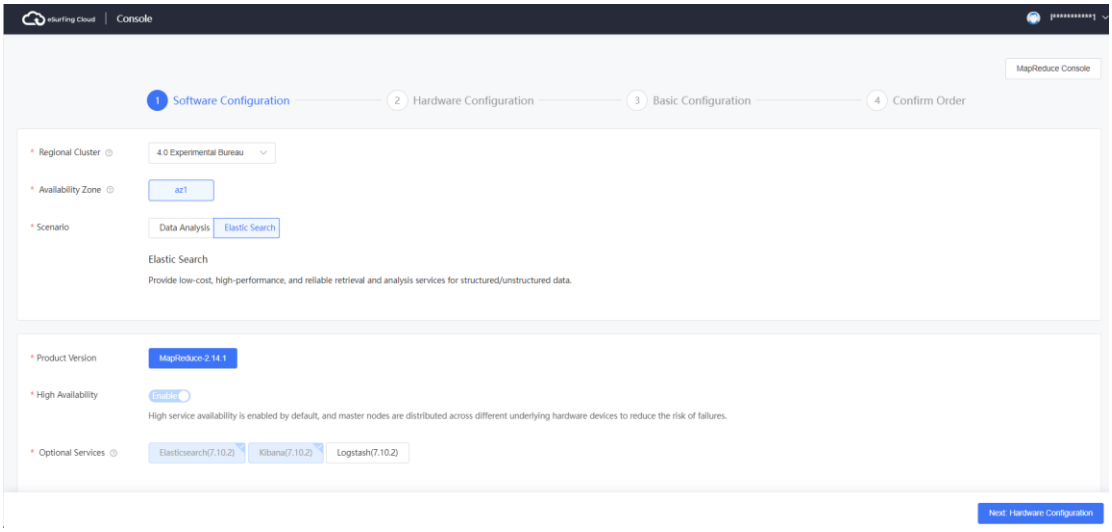

   a. Regional Cluster: This refers to the geographic location where the ECS instances of the cluster reside. Select the region and availability zone based on your requirements, or use the default value.

   b. Business Scenario: Select Elastic Search.

   c. Product Version: Choose the product version to use. The default value is suitable.

   d. High Service Availability: High service availability mode is enabled for the YI-MapReduce service by default and cannot be disabled.

e. Optional Services: These consist of optional and required components, which are determined based on the business scenario. Based on your specific business scenario, you can select optional components. The default optional components for the elastic search cluster scenario are shown below, with gray indicating mandatory.

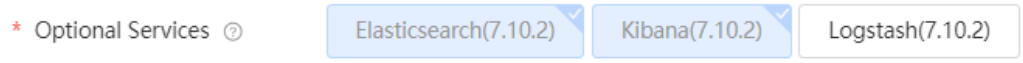

f. After the software configuration is complete, click **Next** to enter the hardware configuration page.

5. The hardware configuration page is displayed as shown below, with parameter descriptions as follows:

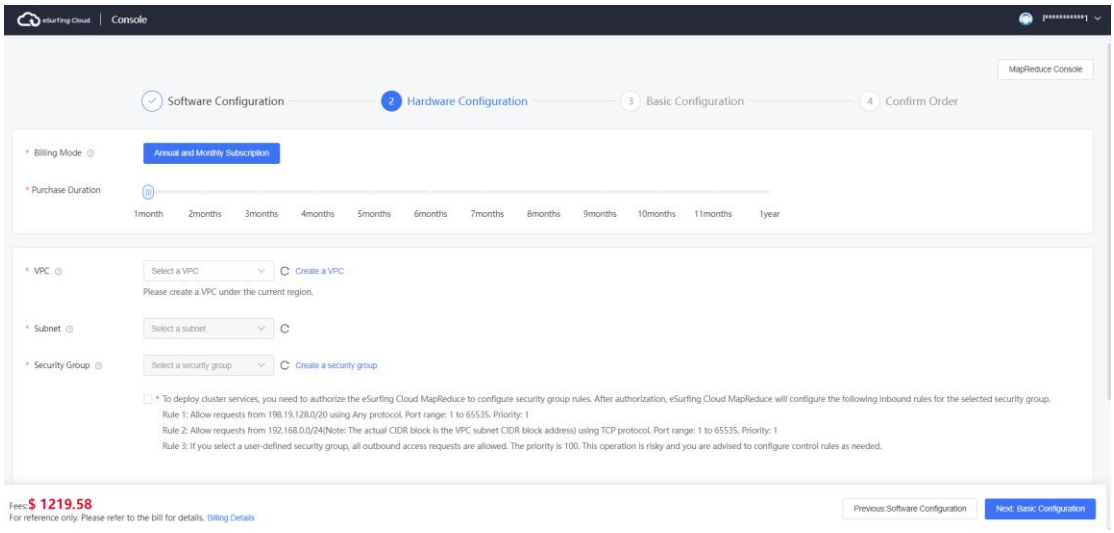

a. Billing Mode: Annual or monthly billing by default (based on the actual purchasable duration on the purchase page).

b. Purchase Duration: Drag the slider to choose a purchase duration.

c. Virtual Private Cloud: Different virtual private cloud (VPC) networks are logically isolated from each other. Select the VPC that suits your needs. If you don't have a VPC currently, click on **Create VPC** to go to the VPC creation page.

d. Subnet: After selecting a VPC, you can choose the subnet as needed.

e. Security Group: This configures network access control for the ECS instances within a cluster. When creating a security group, the ports in the outbound and inbound direction rules need to be opened, and the authorization policy cannot be set to "Reject", otherwise the cluster deployment and

service startup may be affected. For details, see "How to use a custom security group to create a YI-MapReduce cluster?" under Cluster Creation in the FAQs.

f. Node Group: Choose the specifications and quantity of the cluster nodes according to your needs, including the selection of node group types, option configurations, EVS parameters, and performance.

g. After completing the hardware configuration, the cost of the current configuration is calculated in the bottom left corner. Click the **Next** button in the bottom right corner to proceed to the basic configurations page.

6. The basic configurations page is displayed as shown below, with parameter descriptions as follows:

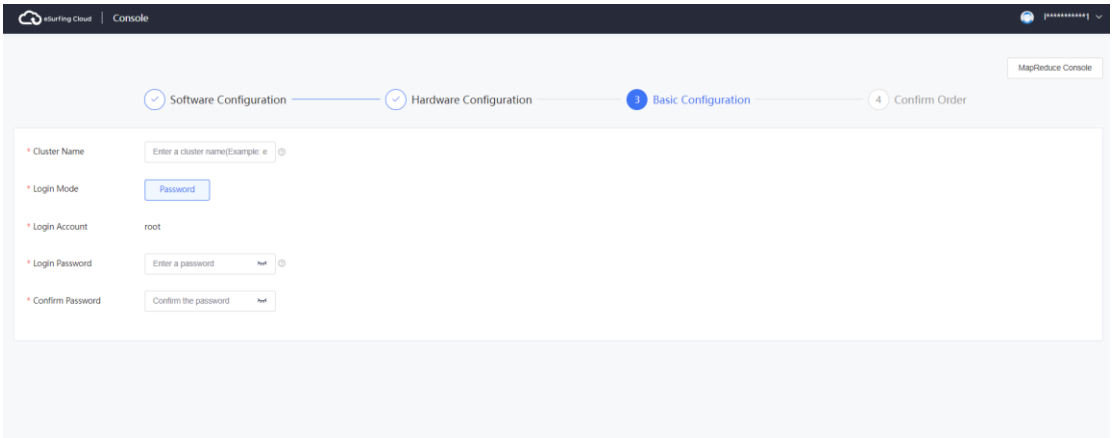

a. Cluster Name: A cluster name contains a maximum of 32 characters, including uppercase letters, lowercase letters, numbers, and special symbols.

b. Identity Credential: The default password.

c. Login Account: The value is 'root' by default.

d. Login Password: Create a password that meets the requirements as prompted.

e. Confirm Password: This should be the same as the login password.

f. After completing the basic configurations, click the **Next** button at the bottom right corner to enter the order confirmation page.

7. The Confirm Order page is as shown below, with descriptions for the configuration info as follows:

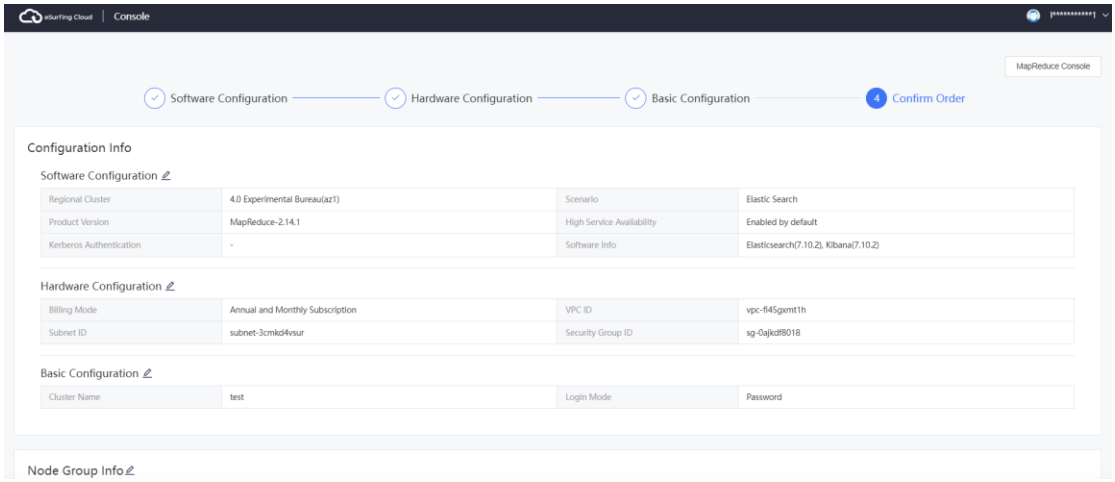

a. Software Configuration: Check the content for any errors. Click **Edit**

   button to go back to the Software Configuration page and modify cluster information.

b. Hardware Configuration: Check the content for any errors. Click the

   Edit button to go back to the Hardware Configuration page and modify cluster information.

c. Basic Configurations: Check the content for any errors. Click the

   Edit button to go back to the Basic Configurations page and modify cluster information.

d. Node Group Info: Check the content for any errors. Click the

   Edit button to go back to the Hardware Configuration page and modify cluster information.

e. Check "I have read and agree to the relevant agreements" and the **Purchase Now** button will become highlighted. Click it to proceed with payment. After purchase, you can rent the cluster with the current configuration.

4.3. Management Cluster

4.3.1. Log in to Cluster

Introduction to YI-MapReduce Cluster Nodes

Remote Login

Nodes within a YI-MapReduce cluster support two methods for remote login, interface login and SSH login:

- Interface Login: Go to the YI-MapReduce console -> enter the Node Management page. Use the remote login feature on the page and log in to the Linux interface of each node in the cluster.
- SSH Login: Only applicable to Linux elastic cloud server (ECS). You can use tools such as PuTTY and XSHELL for remote login to the ECS. Prior to this method, the ECS needs to bind an Elastic IP address. For detailed steps, please refer to Bind/Unbind Elastic IP.

✧ **Description**

For detailed steps of accessing cluster nodes using the password method, please refer to Log in to Cluster Nodes (SSH password method).

Node Types and Functions

A node in the YI-MapReduce cluster refers to an ECS. The different types of nodes and their functions are as shown below.

Node Type	Feature
Master Node	The master node of the YI-MapReduce cluster supervises and monitors the cluster. Select **My Cluster** in the YI-MapReduce console, choose a running cluster and click on its name to enter the cluster information page. Click **Node Management** to view information of

	the nodes. Nodes can log in either using the VNC method by clicking **Remote Login**, or through SSH login.
Core Node	A worker node of the YI-MapReduce cluster that processes and analyzes data, as well as stores process data.

Log in to Cluster Nodes

This section provides guidance on how to utilize the remote node connection (VNC method) offered by the YI-MapReduce console and the password method (SSH method) to log in to nodes within a YI-MapReduce cluster. The main use case for remote connection is in emergency O&M scenarios where related maintenance operations are required on ECS. For other cases, it is highly recommended to use the SSH method for logging into cluster nodes.

✧ **Description**

If the SSH method is required for logging into cluster nodes, you need to manually add an inbound rule in the security group rules of the cluster where the source address is "client public IPV4 address/32" and the port is 22. For detailed steps, please refer to Help Center > Virtual Private Cloud > Security Group > [Adding a Security Group Rule].

Log in to Elastic Cloud Server (VNC Method)

1. Log in to the YI-MapReduce service console.
2. Select **Cluster List > My Cluster**, choose a running cluster and click on its name to enter the cluster information page.
3. Select **Node Management**. Click to expand node group and it will display the node list.
4. In the **Operation** column on the right side of the target node, click **Remote Connection**.

5. In the new tab of the VNC mode web page, you can see the command line interface of the server, similar to the login mode of a Linux server. Enter the root account and password.

✧ **Description**

If you chose the "Password" login method when creating the cluster, the password required is the one you set when creating the cluster.

Log in to Elastic Cloud Server (SSH Key Method)

For Local Windows Operating System Users

If your local system operates on Windows and you need to log in to a Linux ECS, you can follow the steps below. This example uses Xshell.

1. Log in to the YI-MapReduce service console.
2. Select **Cluster List > My Cluster**, choose a running cluster and click on its name to enter the cluster information page.
3. Select **Node Management**. Click to expand node group and it will display the node list.
4. In the **Operation** column to the right of the target node, click **More** and select **Bind EIP**. You can choose an available and unbound elastic public IP for this node, or click **+ Create an EIP**. This will redirect you to the EIP console where you can purchase an EIP, and then complete the EIP binding. If you have already bound an EIP, you may skip this step.
5. Run Xshell.
6. Select New Session.
7. Name: Choose a session name that is suitable and easy for management.
8. Protocol: Use the default "SSH".
9. Host: Enter the EIP to which the ECS is bound.
10. Port No: Use the default "22", as shown below. Click on **Session**.

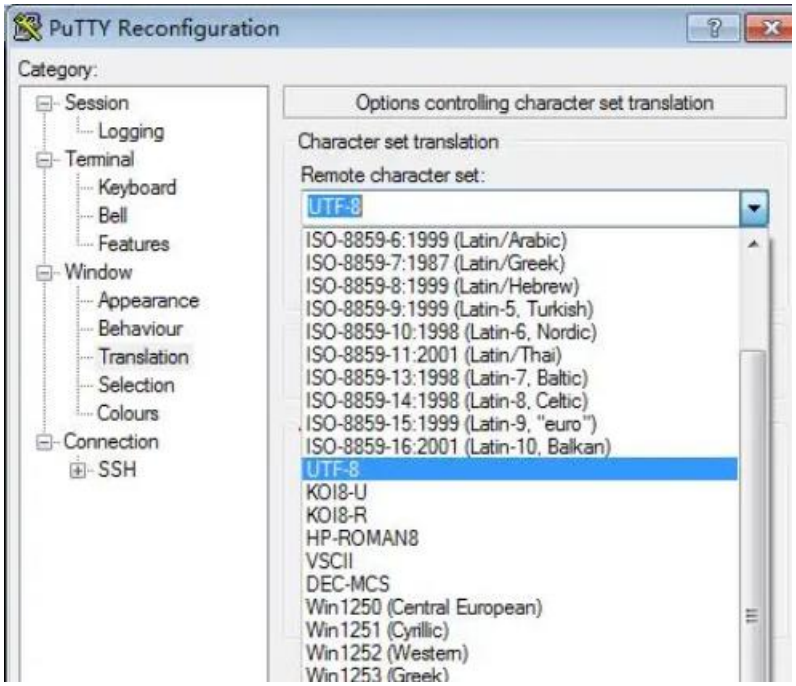

11. Click **Connect**. Follow the instructions to input and save the login username ("root" by default) and password to complete the remote login.

For Local Linux Operating System Users

If your local system operates on Linux and you need to log in to a Linux ECS, you can follow the steps below.

Run the following command to log in to the ECS.

```
ssh default username@EIP
```

Assuming the default username for the Linux ECS is "root", and the EIP is 123.123.123.123, the command should be:

```
ssh root@123.123.123.123
```

4.3.2. Cluster Overview

Introduction to Cluster List

After logging into the YI-MapReduce console, you can view the list of clusters under your account and their current status.

Cluster Status

The table below illustrates the statuses of a cluster.

Status	Description
Starting	The cluster creation is in progress and can be tracked.
Running	The cluster is successfully established and all its components are functioning properly.
Abnormally Terminated	The cluster startup failed or was terminated due to an error. You can view the cause of termination at Reason for Failure.
Releasing	This indicates that the cluster is in the process of unsubscription.
Terminated	The cluster has expired for more than 15 days or has been unsubscribed.
Frozen	For a cluster that has expired, resources are retained for 15 days, during which its status is marked as frozen.

View Cluster Status

Once the cluster creation is complete, you can log in to the YI-MapReduce console to check the list and status of all clusters in your account.

Within **My Cluster**, you can see all clusters that you have created. The default view displays up to 10 cluster information entries. If you have many clusters, you can modify the maximum number of cluster entries displayed on the page at the bottom right, or you can navigate through different pages.

Description of the parameters in the cluster list:

Parameter	Parameter Description

Cluster Name	The cluster name set by the user.
Cluster Type	The cluster type selected by the user during creation.
Payment Type	Annual or monthly by default.
Status	Current status of the cluster. Starting: The cluster creation is in progress. You can view the specific progress at View Process. Running: The cluster has been successfully created and all components are operating normally. Abnormally Terminated: The cluster startup failed or was terminated due to an error. You can view the cause of termination at Reason for Failure. Releasing: The cluster is in the process of unsubscription. Termination: The cluster has expired for more than 15 days or has been unsubscribed. Frozen: For a cluster that has expired, resources are retained for 15 days during which its status is marked as frozen.
Creation Time	The time when the cluster was created.

View Basic Cluster Information

After logging into the YI-MapReduce console, go to the cluster list page. Click on the cluster name needed to enter its basic information page.

The basic information parameters for the cluster mainly include details about the cluster, software, and network hardware.

Cluster Information

Parameter	Description
Cluster Name	The name of the cluster. Click **Edit** to modify its name.
Creation Time	The time when the cluster was created.
Run Time	The length of time the cluster has been running.
Expiration Time	The time when the cluster expires.
High Availability	Enabled by default.
Payment Type	Annual or monthly by default.
Kerberos Authentication	Enabled by default, based on requirements of business scenarios.

Software Info

Parameter	Description
Product Version	The version of YI-MapReduce used by the cluster.
Software Info	The components used and their respective versions used during cluster creation. For detailed information, click on the component name you want to view.
Scenario	The business scenario selected when creating the cluster.

Network Hardware

Parameter	Description
Regional Cluster	The region where the cluster is located.

Hardware Configuration	The hardware information for all nodes in the cluster, including MASTER, CORE, and TASK (if any).
VPC ID	The ID of the VPC used by the cluster. Click  to copy the VPC ID. Click on the VPC ID to jump to and view the detailed information of this VPC.
Security Group ID	The ID of the security group used by the cluster. Click  to copy the Security Group ID. Click on the security group ID to jump to and view the detailed information of the security group.

Manage Component and Host Monitoring

In daily operations, users can use YI-MapReduce Manager to monitor the status and metric information of all components (including role instances) and hosts:

- Status information includes running status, health status, and role instance status statistics.
- Metric information includes the main monitoring metrics for each component.

Manage Component Monitoring

1. Go to the YI-MapReduce Manager page and select **Cluster Service** from the left menu.
2. Click on any cluster in the Cluster Service list to enter.
3. Choose the Cluster Status page to view the health status of instance groups.

Status	Description
Alarm	Exceptions have occurred in the cluster instance service.
Fault	The cluster instance service has encountered a failure.
Normal	The cluster instance service is running.

Not monitored	The cluster instance does not have Prometheus installed. The service is not being monitored.

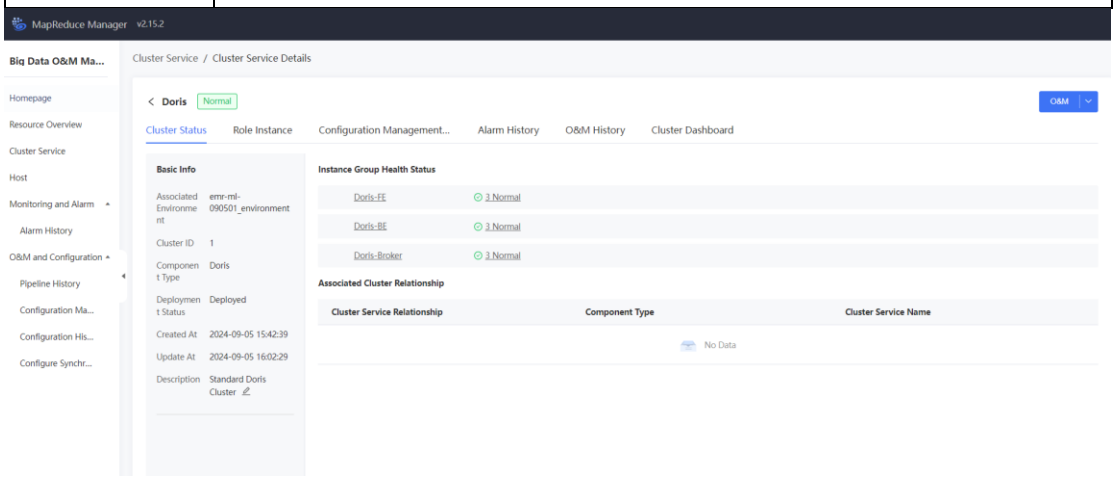

4. Click on the health status of any instance to go to the Role Instance page. It displays the role of each instance component, host IP, host name, host status, startup status, and health status, grouped by instance.

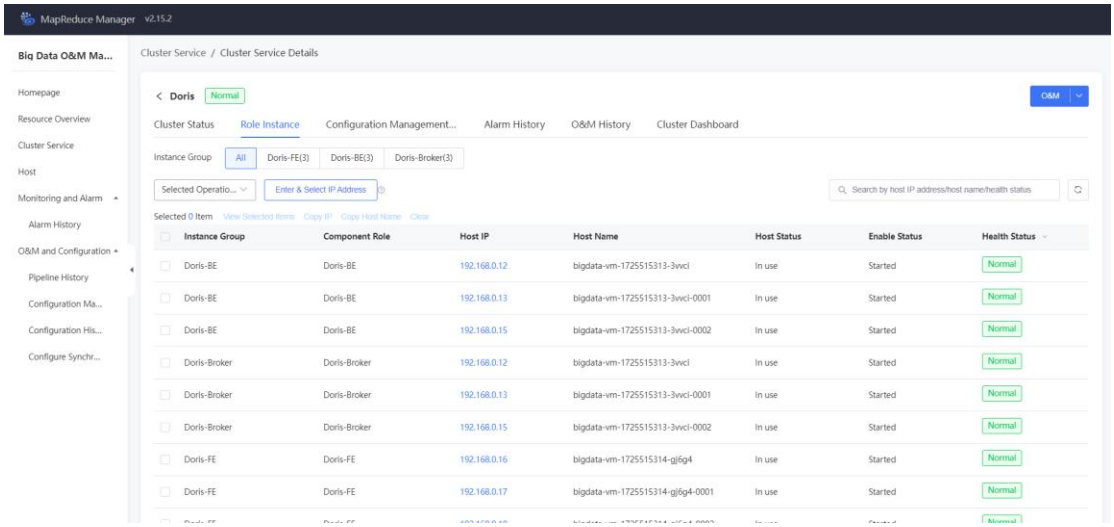

Manage Host Monitoring

1. Go to the YI-MapReduce Manager page and select Host from the left menu.
2. The host page shows various details such as Host IP, Host Name, Initialization Status, Host Status, Health Status, Host Label, Memory, CPU Model, Host Configuration, Operating System, Description, etc.

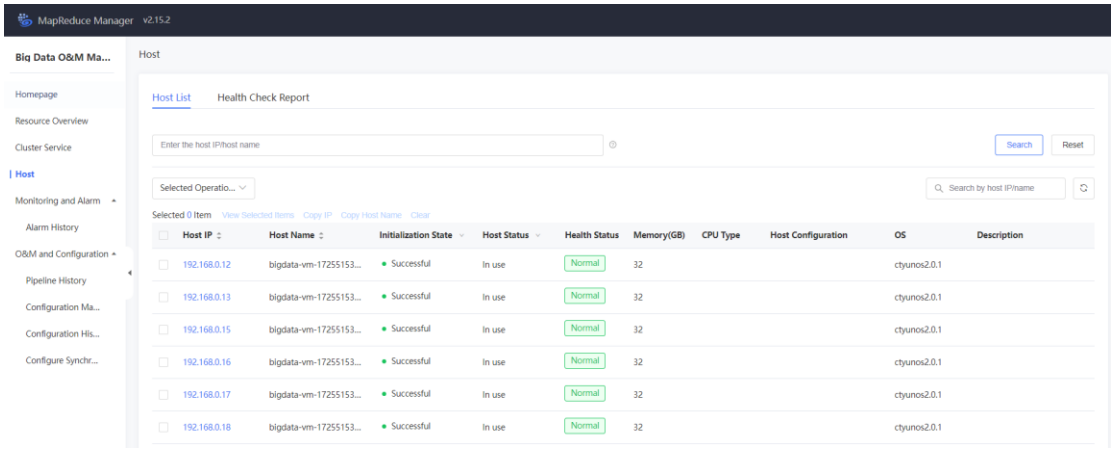

3. The Health Status column displays the health status of the host.

Status	Description
Alarm	Exceptions have occurred in the cluster instance service.
Fault	The cluster instance service has encountered a failure.
Normal	The cluster instance service is running.
Not monitored	The cluster instance does not have Prometheus installed. The service is not being monitored.

4. Click on the **Host IP** to go to the Host Status information bar, which displays the host status, instance list, and alarm history.

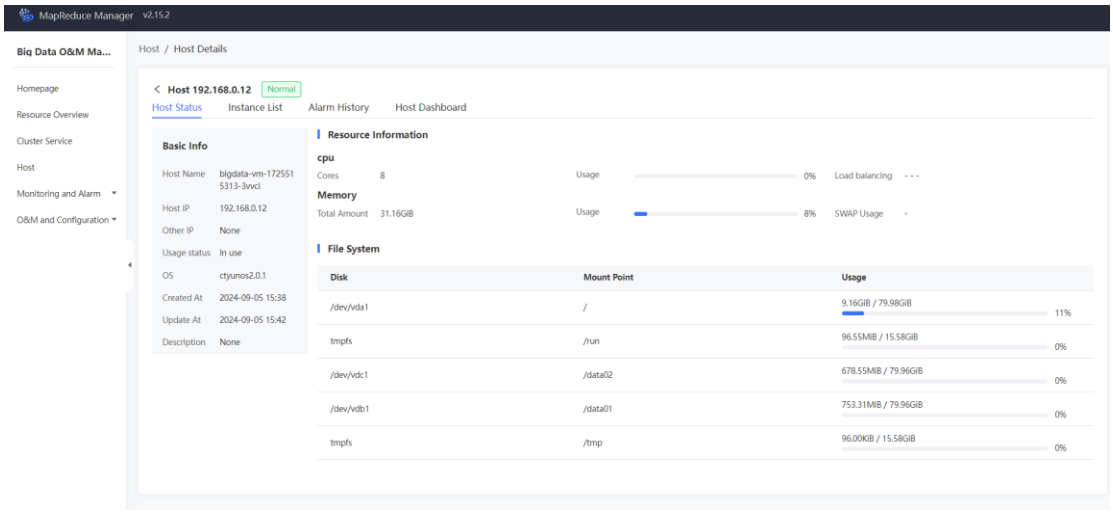

5. Click on **Alarm History** to view the historical alarms associated with the specific host IP.

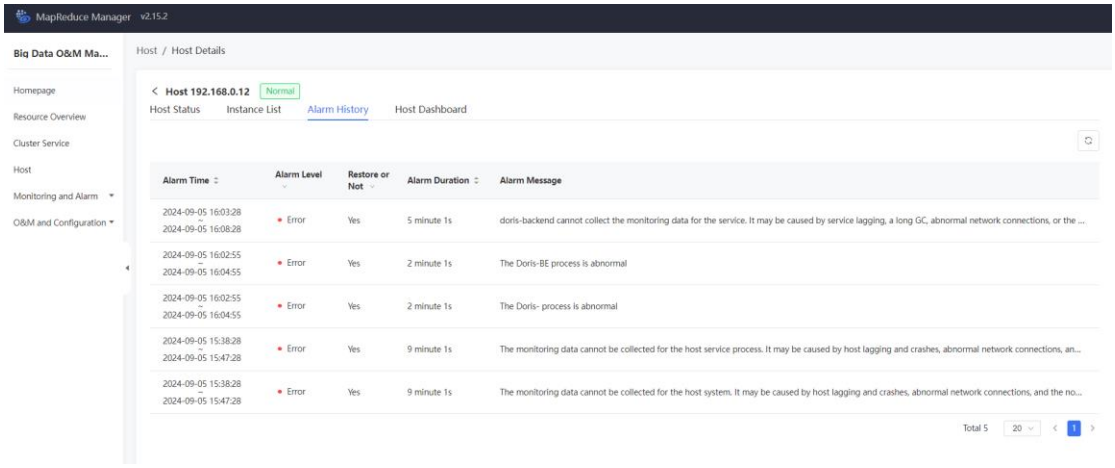

4.3.3. Health Check

Health Check

Health Check

Scenario

This guide aids users in performing regular health checks on the cluster as part of routine maintenance. This process ensures that all cluster parameters, configurations, and monitoring systems operate without exception, thereby facilitating long-term stability.

Procedure

A cluster health check involves examining every component at the Manager service level. The service level focuses on the ability of components to provide regular services, the status of alarms, and various inspection metrics for each component.

Procedure

Cluster Health Check

1. Initiate a health check for all services manually.

On the cluster details page, click **More Operations > Start Cluster Health Check** in the upper right corner of the cluster service list.

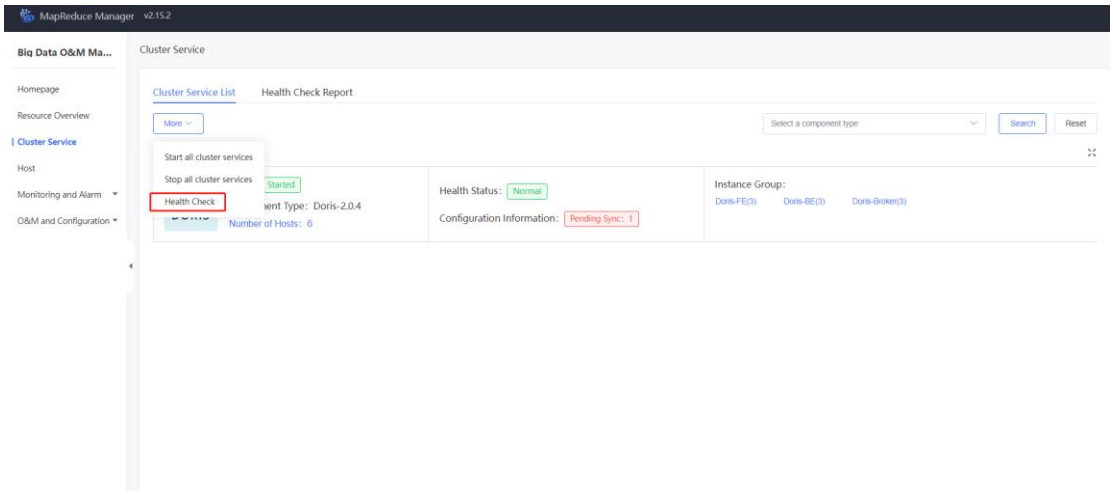

2. Set up a periodic health check for all services.
   Go to **Cluster Service > Health Check Report** > click the **Settings** button in the upper right corner of the list. A pop-up for periodic task settings will appear. This allows users to set intervals based on days/weeks/months, and it can be paired with a specific time. After successful creation, a health check will be initiated regularly for the supported component services.

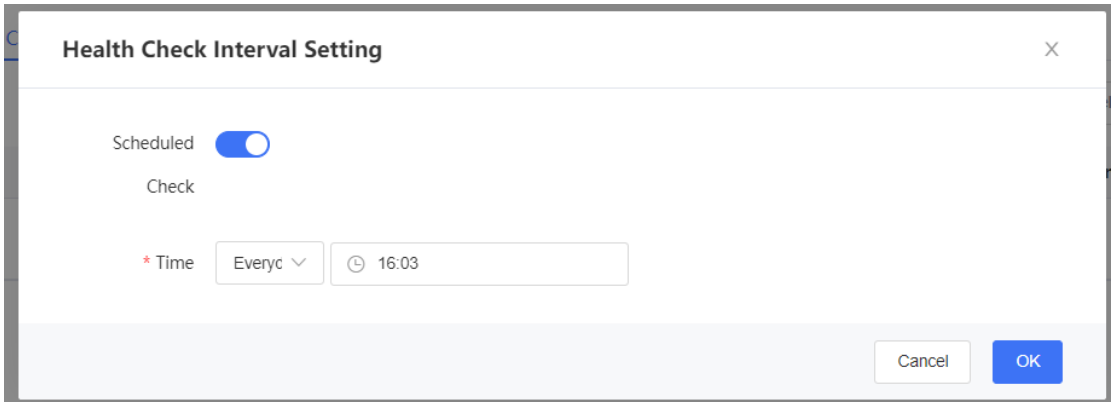

✧ Description

- In the YI-MapReduce Manager interface, select Cluster Settings > More Operations > Health Check to initiate a cluster health check. The health check launched from the cluster list page targets all component services in the list.
- After clicking on **Health Check**, you will be redirected to the Health Check Report tab, which displays the cluster health check list. The most recent initiated health check is listed at the top. Expand the first-level list to view the check details for each component.

- Click **Export Report** under the first-level list operations to export the report. This supports exporting the report as a CSV file, allowing users to review the health report locally.

View and Export Check Reports

Scenario

For a detailed analysis of health check outcomes, you can view and export health check results on YI-MapReduce.

Procedure

The scope of platform health checks includes the health checks at the Manager service level.

Cluster health checks may cover three aspects: the service status of each object checked, alarm information, and related metrics differentiated for each component.

Before You Begin

A health check has been conducted.

Procedure

1. On the cluster details page, click on **Management Operations > View Cluster Health Check Report** to view the health check report.

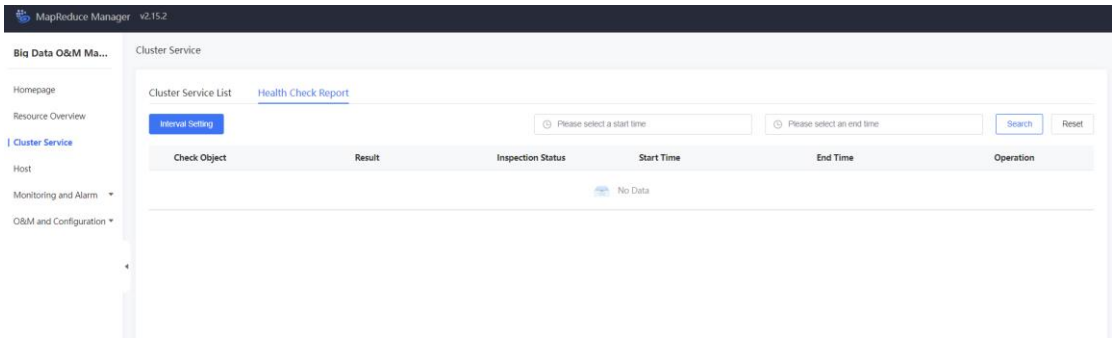

2. On the health check report panel, click **Export Report** to export the health check report. After downloading, you can view the complete information of the

check items locally. Both cluster and host health check reports are exported in CSV format.

View YI-MapReduce Service Operation Logs

Operation Log Access

1. Under the **Cluster Name** column on the cluster list, click on the name of the cluster whose log you want to view, to access the cluster information page.

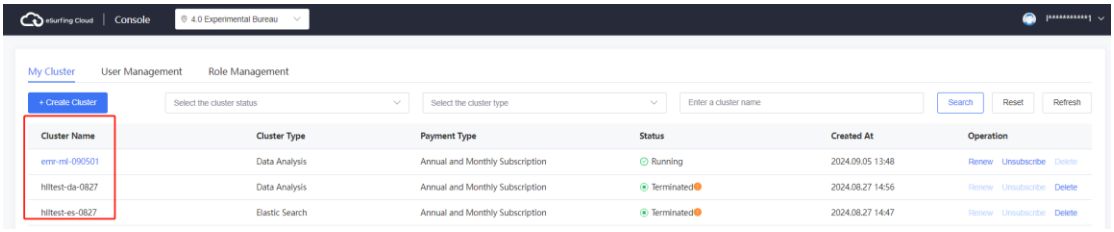

2. On the cluster information page, click on **Operation Logs** as shown below to go to the operation logs page.

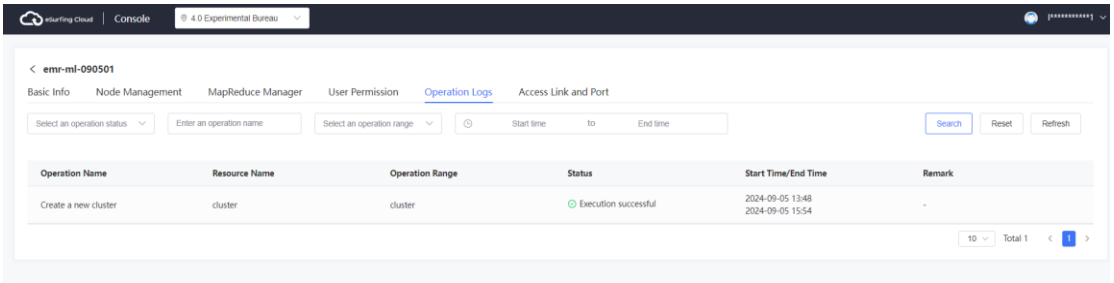

Operation Type

Currently, the YI-MapReduce service operation logs include a filter box to assist users in quickly locating the source of the issue. After selecting the filter conditions, you can click the **Query** button to search the log or the **Reset** button to clear the filter conditions.

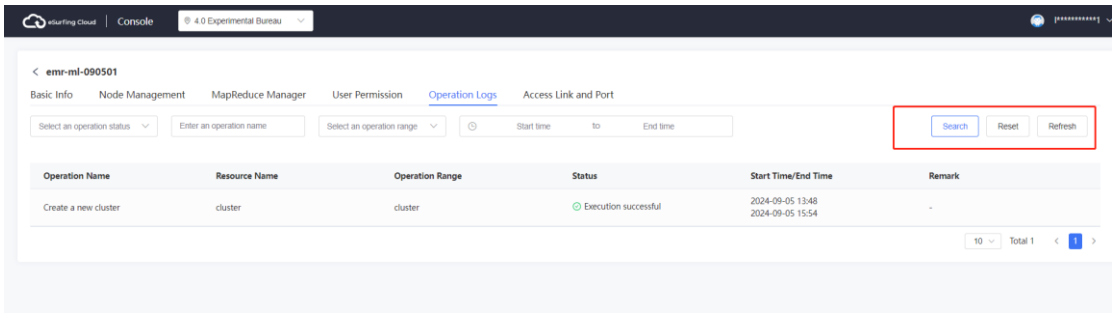

1. Operation Status: The operation status filter box is as shown below. Click to select filter conditions, including "Execution successful", "Failed", and "In Progress".

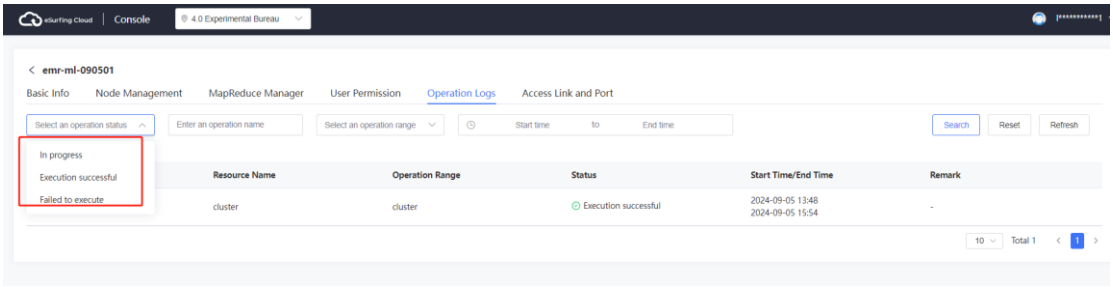

2. Time Range: The time range filter box is as shown below. Click **Start Date** and **End Date** to set filter conditions.

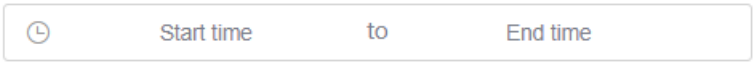

Log Field

Description of log field parameters.

Parameter	Parameter Description
Action Name	The name of the executed operation, such as creating a new cluster, upgrading configurations, etc.
Status	The status of the operation, including succeeded, failed, in progress.
Operation Range	The range of the operation, including cluster, node, node group, etc.
Operation User	The user who executes the operation.
Start Time/End Time	The start and end times of the operation.
Remarks	Notes explaining the failure of operations.

4.3.4. Node Management

Manage Node Operations

Scenario

In the case of a node failure, you may need to stop all cluster services on YI-MapReduce and perform maintenance checks on the node. Once the issue is resolved, you can restart all cluster services on the node to restore service functionality.

Procedure

1. Log in to the YI-MapReduce console, click the name of the active cluster on **My Cluster**, and go to the cluster details page.
2. On the cluster details page, click **YI-MapReduce Manager** and click **Go to YI-MapReduce Manager**.
3. On the **Cluster Service** page of the Manager, select **More Operations**, and then click on **Stop All Cluster Services** to perform the corresponding operation.

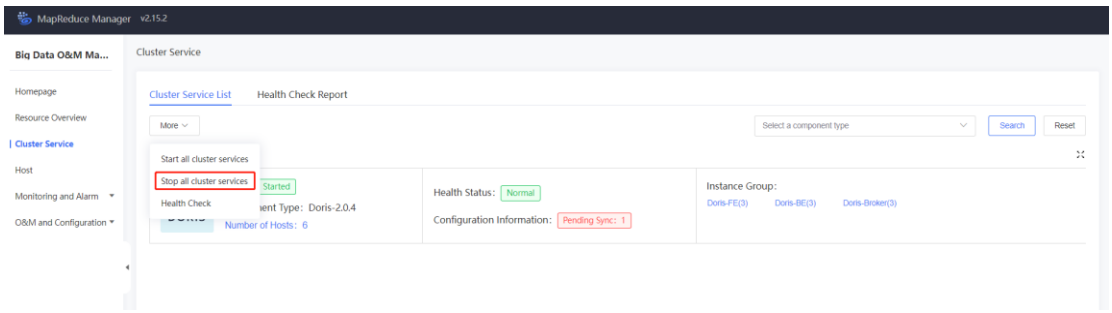

Remote Connection

Scenario

This section primarily introduces how to log in to the ECS using the remote login feature (i.e., VNC method) provided by the console. This allows users to check the operational status or any issues with the operating system of the instance.

Before You Begin

The ECS instance has a pre-set login password.

Procedure

1. Log in to the YI-MapReduce console, click the name of the active cluster, and go to the cluster details page.
2. Go to the **Node Management** page, click the drop-down button in the **Node Group Name** column to expand the corresponding node information, click **Remote Connection** in the **Operations** column, and a shell command window will appear on the new page.

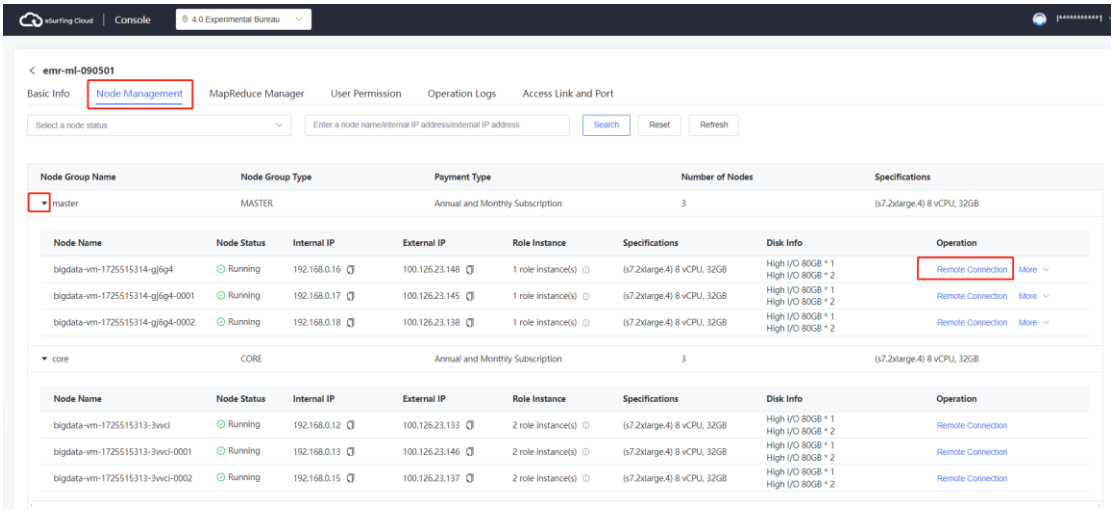

3. You can execute remote login through the command line as prompted.

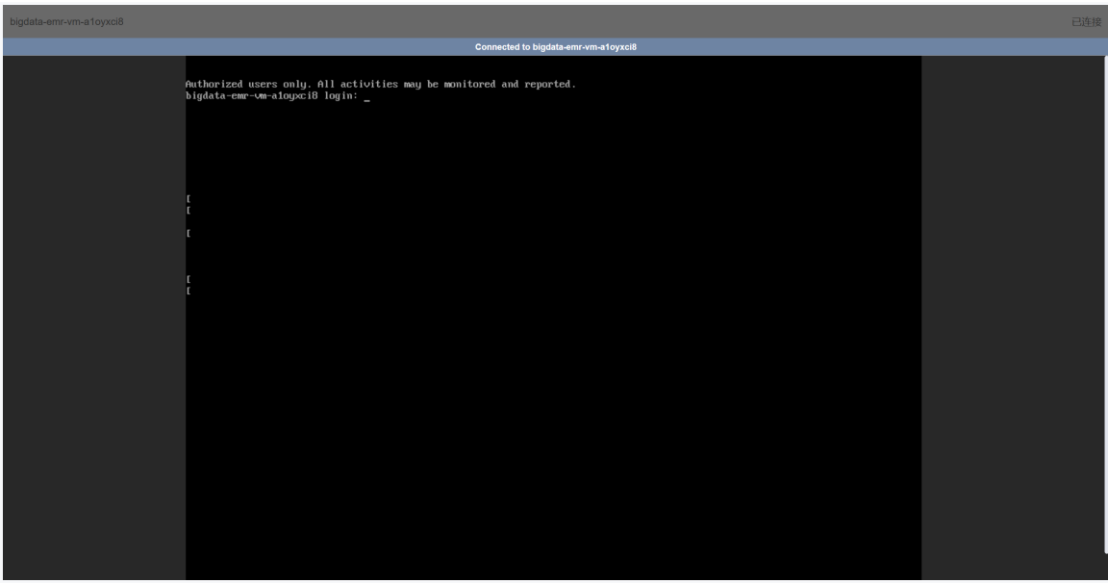

Bind/Unbind EIP

Background

By default, all nodes in the created YI-MapReduce cluster are assigned an internal IP. Users are required to manually generate an external IP.

Steps to Bind EIP

1. Log in to the YI-MapReduce console. On the **My Cluster** page, click on the specific cluster name to go to the cluster details page.

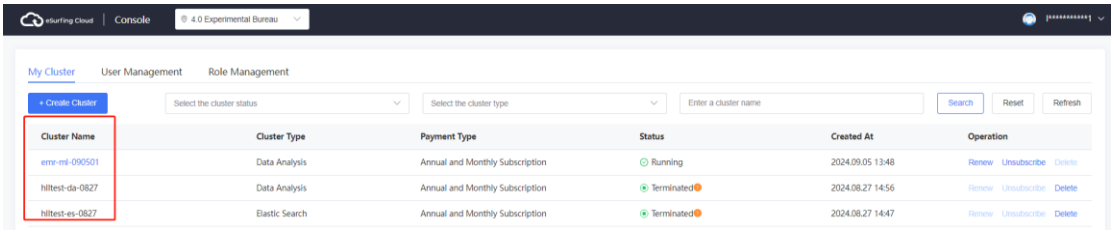

2. On the **Node Management** page, click the drop-down button in the **Node Group Name** column to expand the corresponding node information. Click **More** in the **Operations** column.

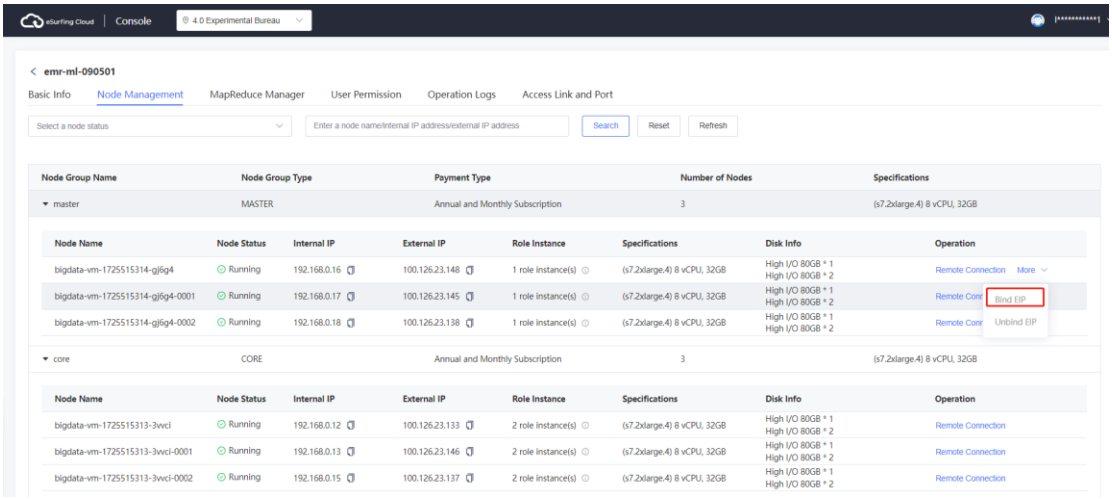

3. Binding: Click on **Bind EIP**. In the Bind EIP pop-up, if there are no available EIP under your account, you need to click the **+ Create an EIP** button to be redirected to a new page for creation. If there are available EIPs under your account, you can select the IP from the EIP dropdown box, and click **OK** to bind. Please note: The bound EIP will not be automatically released with the node. You need to go to the EIP console for further actions.

Steps to Unbind EIP

1. Log in to the YI-MapReduce console. On the **My Cluster** page, click on the specific cluster name to go to the cluster details page.

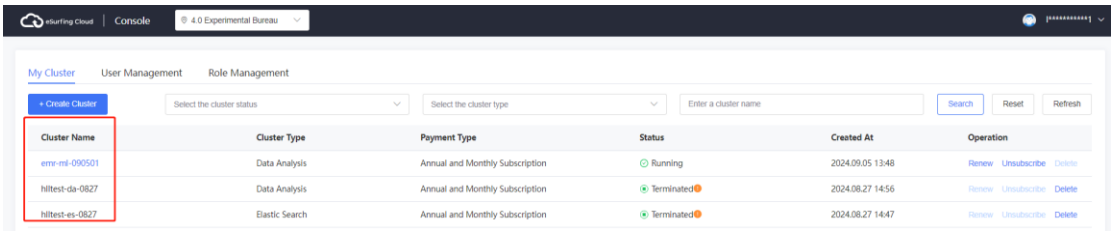

2. On the **Node Management** page, click the drop-down button in the **Node Group Name** column to expand the corresponding node information. Click **More** in the **Operations** column.

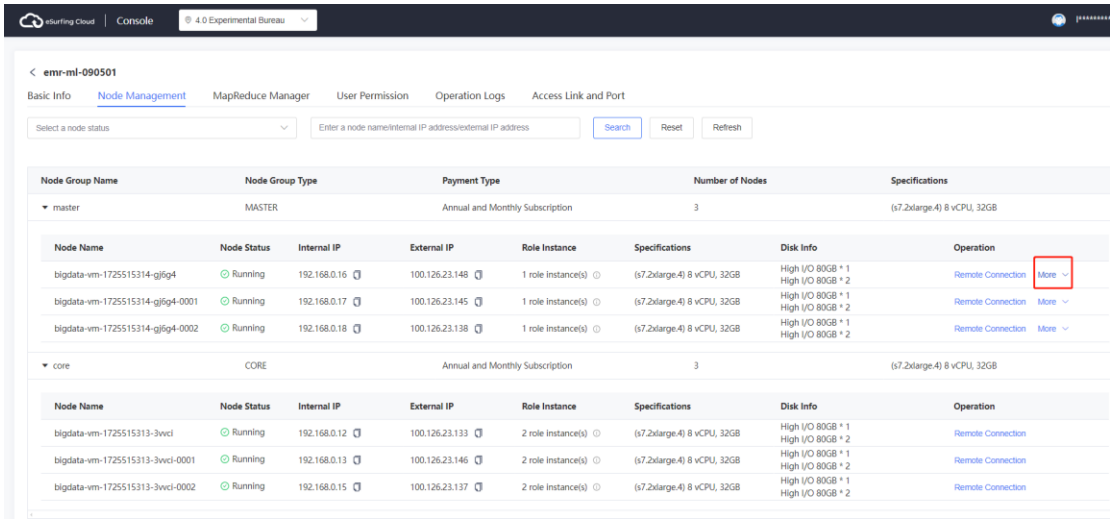

3. Unbinding: Click on **Unbind EIP**. In the Unbind EIP pop-up box, click **OK** to unbind. Please note: The unbound EIP will not be automatically released with the unbinding operation. You need to go to the Elastic IP Console for further actions.

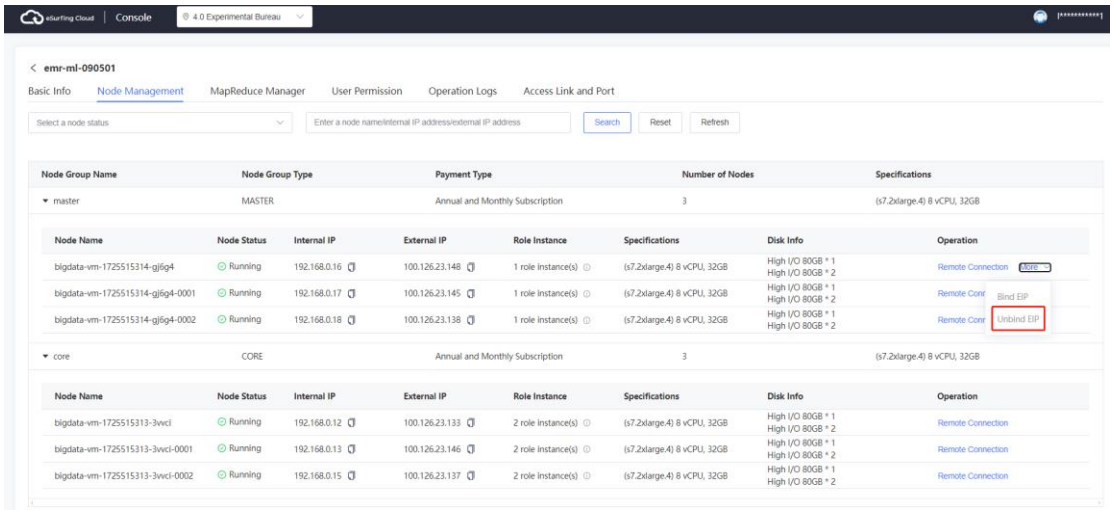

4.3.5. Component Management

Object Management Introduction

YI-MapReduce (or "YI-MR") clusters include a variety of different basic objects. The descriptions of these distinct objects are as follows:

Object	Description	Example

Cluster Service	Collection of features capable of executing specific tasks	HDFS cluster service, YARN cluster service.
Service Roles	Functional instances that constitute a complete service, usually distinguished by roles	HDFS comprises NameNode, DataNode, HDFSClient, ZKFC, JournalNode. Each plays various roles within the component, communicating and coordinating to provide distributed file storage capacity.
Instance Groups	Groups comprising instances with identical roles, typically one instance group corresponds to one role, and the name generally aligns with the role name	JournalNode instance group, DataNode instance group.
Role Instances	Specific instances of service roles operating on host nodes	KerberosAdmin operating on Host2, KerberosServer operating on Host3.
Host	Physical or virtual hosts used for running role instances	Host1 to Host5.
Cluster	Logical entities composed of multiple hosts capable of providing various services	A cluster named Cluster1 composed of 5 hosts (Host1 to Host5) providing services such as HDFS and YARN.

View Configuration

On YI-MapReduce (or "YI-MR"), users can view the configurations of cluster services.

Procedure

1. Log in to the YI-MapReduce management console.

2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Management** on the menu.
5. Choose the specific cluster service in the query conditions to view the associated configurations of that cluster service. As shown in the figure:

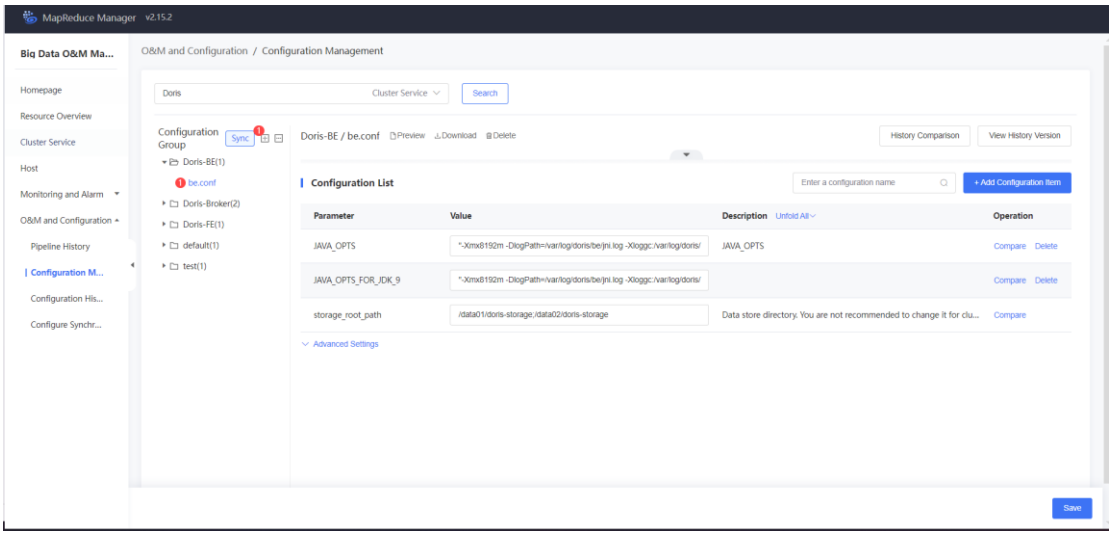

Manage Services

On YI-MapReduce (or ""YI-MR"), users can perform the following operations:

- Start Cluster Service
- Stop Cluster Service

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. On the **YI-MapReduce Manager** interface, click the **Cluster Service** menu.

5. Choose the specific cluster service. Click on the cluster service icon to enter the cluster service details page.
6. Click on the **O&M Operations** button, then on **Start Cluster Service, Stop Cluster Service, or Roll-restart Cluster Service** to perform the corresponding operations. There are dependencies between cluster services. When performing start, stop, rolling restart operations on a certain cluster service, services that have dependencies on this cluster service will be affected. The specific impacts are as follows:
    a. When starting a certain cluster service, the lower-level cluster services that it depends on must first be started for the cluster service to function effectively.
    b. When stopping a certain cluster service, the upper-level cluster services that depend on it will not be able to provide functions.
    c. When rolling restart a certain cluster service, the upper-level cluster services that depend on it and are started need to be restarted to take effect.

Configure Service Parameters

On YI-MapReduce (or " YI-MR"), users can view and modify the default configurations of cluster services easily according to actual business scenarios.

Steps to Modify Cluster Service Configurations

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. On the **YI-MapReduce Manager** interface, click the **Cluster Service** menu.
5. Choose the specific cluster service. Click on the cluster service icon to enter the cluster service details page.
6. Click the **Configuration Management** tab, select the configuration file to be modified, make the necessary changes, and then click the **Save Changes** button at the bottom right corner.

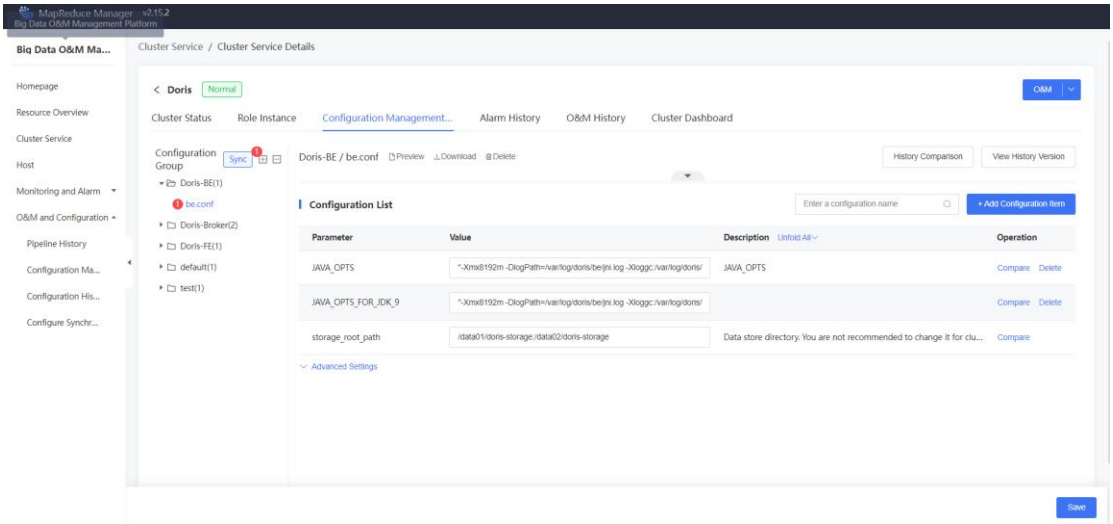

7. Click the **Sync** button to synchronize the modified configuration file.
8. Once the configuration sync is complete, re-enter the service details page.
9. Click on **O&M Operations**, then click **Roll-restart Cluster**. The modified configurations take effect after the cluster service restarts.

Sync Service Configurations

Scenario

If some cluster services have configuration files that are pending synchronization, you can use the configuration synchronization feature to carry out this task.

Impact on System

Once the cluster service configuration is synchronized, a rolling restart of this configuration is necessary for the changes to take effect. During the restart, the relevant cluster service is temporarily unavailable.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.

3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. On the **YI-MapReduce Manager** interface, click the **Cluster Service** menu.
5. Choose the specific cluster service. Click on the cluster service icon to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Click the **Sync** button to synchronize the configuration file pending synchronization.
8. Once the configuration sync is complete, re-enter the service details page.
9. Click on **O&M Operations**, then click **Roll-restart Cluster**. The modified configurations take effect after the cluster service restarts.

Manage Role Instances

On YI-MapReduce (or "MapReduce"), users can perform the following operations on specific role instances:

- Start Instance
- Stop Instance
- Restart Instance
- Set Maintenance
- Cancel Maintenance

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. On the **YI-MapReduce Manager** interface, click the **Cluster Service** menu.

5. Choose the specific cluster service. Click on the cluster service icon to enter the cluster service details page.
6. Click on the **Role Instance** tab.
7. Check the box in front of the role instance that requires operation.
8. Click **Operate on Selected**, and then click **Start Instance, Stop Instance**, **Restart Instance**, **Roll-stop Instance**, **Roll-restart Instance**, **Set Maintenance**, or **Cancel Maintenance** to execute the corresponding operation.

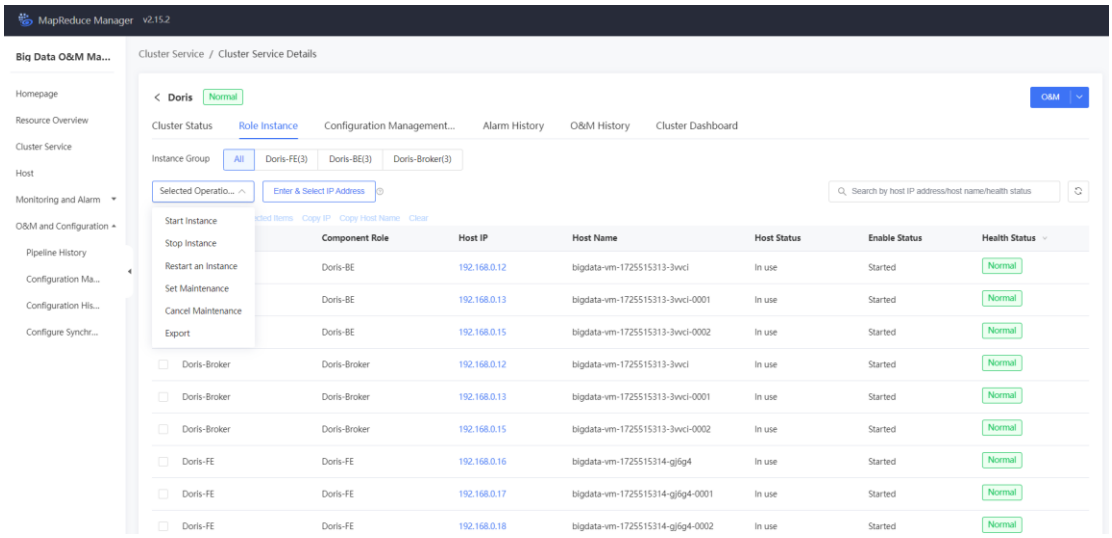

Start and Stop Cluster

A cluster is a set of service components. Users can start or stop all services within the cluster.

Start and Stop Cluster

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. On the **YI-MapReduce Manager** interface, click the **Cluster Service** menu.
5. Click **More Operations**, then choose either **Start All Cluster Services** or **Stop All Cluster Services** to execute the corresponding operation.

Download Cluster Configurations

Scenario

To meet practical business needs, you can download configuration files to quickly update the configurations of cluster services.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. On the **YI-MapReduce Manager** interface, click the **Cluster Service** menu.
5. Choose the specific cluster service. Click on the cluster service icon to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Choose and click on the name of the configuration file to download. Click the **Download** button to proceed.

4.3.6. Alarm Management

View Alarm List

View Alarm List

Alarm History presents all alarm information within the cluster.

By default, the alarm history is organized by time, prioritizing the most recent alarms.

Descriptions for fields in the alarm message are as follows:

Parameter	Parameter Description

Alarm Time	The time at which the alarm was generated.
Alarm Level	Disaster: Cluster service is unavailable. Serious: A fault that might impact the service has occurred in the cluster service. Error: An exception has occurred in the cluster service that requires additional handling, but the cluster service continues to provide normal service. Warning: Whether the alarm will affect the service operation remains uncertain. It indicates potential risks that require further monitoring and troubleshooting.
Restore or Not	Indicates whether the alarm still exists.
Alarm Duration	The duration from the occurrence of the alarm to the present.
Host IP	The IP address of the host where the alarm occurred.
Alarm Message	The alarm message.
Associated Cluster Service	The cluster service or instance where the alarm occurred.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **Monitoring and Alarm> Alarm History** on the menu.
5. Select the query conditions for the specified cluster service, role instance, or host IP. Then, click the **Query** button to view related alarm information.

4.4. Access Web Pages of Open-Source Components Hosted on YI-MapReduce Cluster

4.4.1. Open-Source Component Port List

Background

With the number of managing components increasing, port conflicts can occur, especially during the testing and validation phases when different nodes are fully mixed. These conflicts may prevent some services from starting.

Method

Common fixed ports for components are identified through the netstat command corresponding to PID, official configuration files, and current component configuration files.

Result

Common ports for each component are identified. During the initialization of Linux machines, these common ports are set in net.ipv4.ip_local_reserved_ports to reserve these ports.

List of Common Components

Common components include Doris and ElasticSearch, and more.

Common Doris Ports

Version: 2.0.4

Configuration Parameter	Default Port	Current Configuration Port	Port Modification	Port Description
be_port	9060	9060		Port of the thrift server on BE, used to receive requests from FE.
brpc_port	8060	8060		The brpc port on BE, used for communication between BEs.
edit_log_port	9010	9010		bdbje port
heartbeat_service_port	9050	9050		Heartbeat service port on BE (thrift), used to receive heartbeats from FE.
http_port	8030	8030	8035	FE http port. Currently, all FE http ports must be the same. It has been modified to 8035 in the production environment.

https_port	8050			FE http port. Currently, all FE http ports must be the same.
query_port	9030	9030		Doris FE query connection port through MySQL protocol
rpc_port	9020	9020		Port of FE Thrift Server
single_replica_load_brpc_port	9070			RPC port for communication between Master and Slave replicas in the single replica data import function. After the Master replica completes the flush, it notifies the Slave replica to synchronize data through RPC, and the Slave replica notifies the

				Master replica after the data synchronization is completed. The system opens a separate BRPC thread pool for communication between Master and Slave replicas during the single replica data import process, to avoid data synchronization between replicas competing for thread resources with data distribution and query tasks when import concurrency is high.
single_replica_load_download_port	8050			Port for the Slave replica to download data files

				from the Master replica through HTTP in the single replica data import function. The system opens a separate HTTP thread pool for the Slave replica to download data files from the Master replica during the single replica data import process, to avoid the Slave replica downloading data files competing for thread resources with other http tasks when import concurrency is high.

webserver_port	8040	8045	8045	Service port of the http server on BE

Common ElasticSearch Ports

Official reference:

https://www.elastic.co/guide/en/elasticsearch/reference/7.10/modules-network.html

Version: Elasticsearch-7.10.2

Configuration Parameter	Default Port	Current Configuration Port	Port Description
com.sun.management.jmxremote.port	None	9400	es jmx port
http.port	None	9200	es http port
transport.port	None	9300	es transport port

4.4.2. Access Through EIP

For easy user access to open-source component websites, the YI-MapReduce cluster enables access to hosted open-source components on the cluster by assigning an EIP. This approach is more user-friendly and efficient, and is recommended for accessing open-source component websites.

Bind EIP

1. Under the **Cluster Name** column on the cluster list, click on the name of the cluster that requires IP binding, to access the cluster information page.

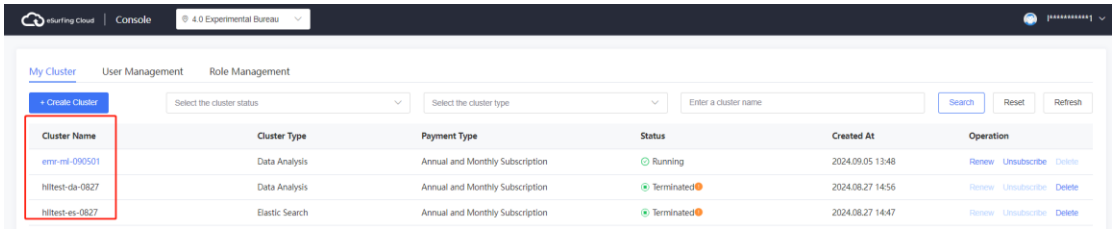

2. Determine the node where the component that requires public access is located. Click on **Access Link and Port** to find the component that requires public access in the Cluster Service Name column. Its native UI address contains the Intranet IP address of the node.

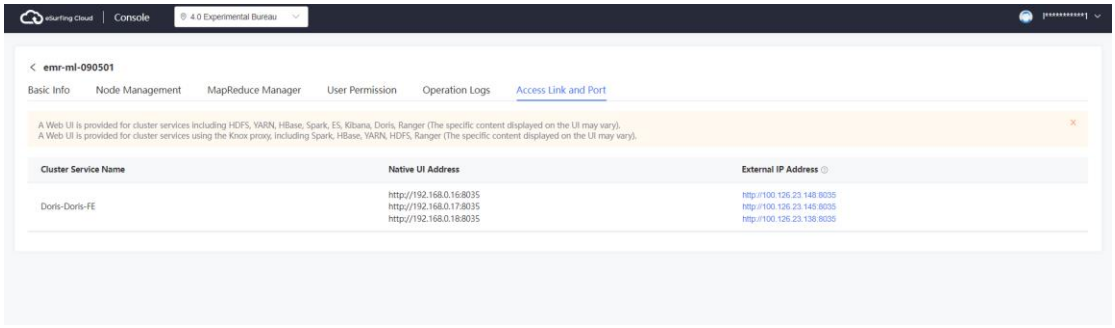

3. Click on **Node Management** and then the drop-down button in the **Node Group Name** column to expand the corresponding node information and determine the node that requires public access. Click

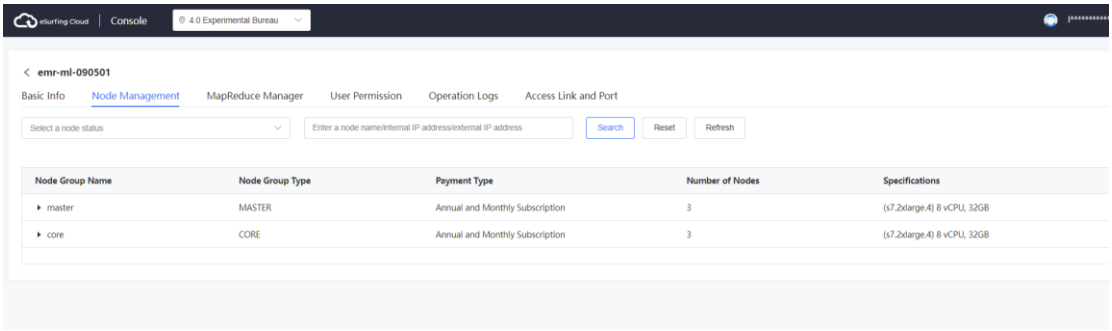

More in the **Operations** column.

4. Click on **Bind EIP**. In the **Bind EIP** pop-up, if there are no available EIP under your account, you need to click the **+ Create an EIP** button to be redirected to a new page for creation. If there are available EIPs under your account, you can select the IP from the EIP dropdown box.

   Description

   The bound EIP will not be automatically released with the node. You need to go to the EIP console for further actions.

5. After binding the EIP, you can add or modify the cluster security group rules, and adjust the port access permissions according to your need. Then, you can access the website of the open-source component in the "Public IP: Port" format.

4.5. YI-MapReduce Manager Instructions

4.5.1. Getting Started

Introduction to YI-MapReduce Manager

Overview

YI-MapReduce Manager, independently developed by eSurfing Cloud, primarily offers O&M capabilities for big data components. It manages various system resources and data assets visually and process-orientedly, supporting automated O&M scheduling, unified configuration file management, unified O&M monitoring, and functions such as component cluster service management, multi-tenant management, and resource management. Besides enhancing the efficiency of O&M personnel, it provides comprehensive and professional big data O&M capabilities for big data O&M workers, greatly reducing the threshold for operating big data platforms.

- It offers cluster service management, allowing users to quickly grasp data information of the environment, cluster, host, and component services.
- It provides O&M automation management to increase the efficiency of big data operations and reduce labor costs.

Browser Support

Google Chrome browser is recommended.

System Interface Introduction

the YI-MapReduce Manager provides a unified platform for cluster management, facilitating efficient and intuitive operation and maintenance of the cluster.

See the figure below: YI-MapReduce System Interface.

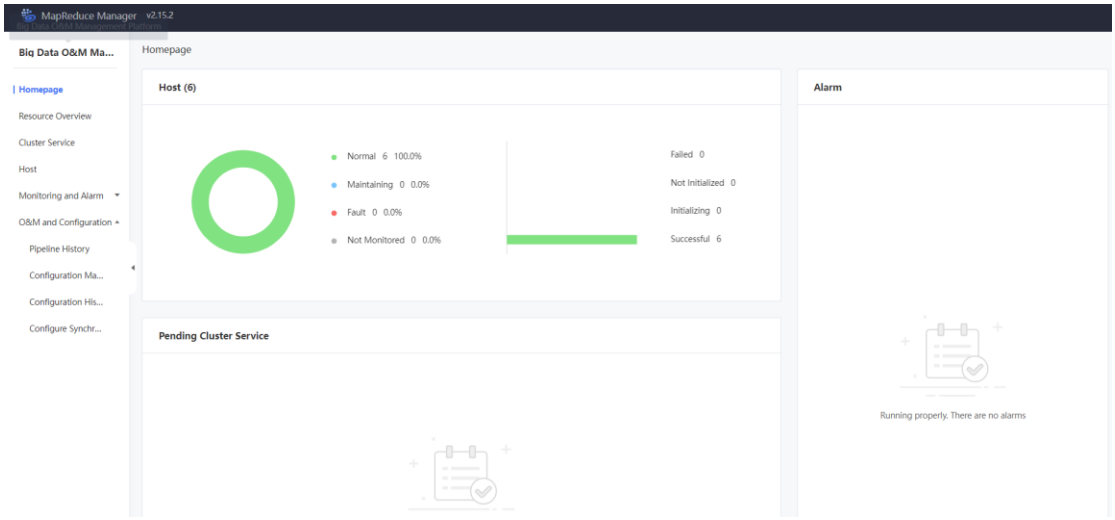

The left side consists of the navigation menu area, and the right side is for display and operations.

The detailed functions of the navigation menu area are as shown in the table below:

Menu Item	Description
Homepage	Displays summary of the health status and initialization of the platform host; Information on clusters with faults and alarms; Platform alarm messages.
Resource Overview	Displays information such as CPU, memory, and network of all hosts.
Cluster Service	Displays all clusters and lists the cluster service types; Supports one-click start and stop of all clusters in the cluster list; Supports actions such as start, stop, restart, and roll-restart

	of a single cluster; Supports actions such as start, stop, restart, roll-stop, and roll-restart of role instances; Supports viewing the O&M history, alarm history, and monitoring dashboard of the cluster.
Host	Displays all hosts of the current platform; Supports actions such as setting a host to maintenance mode and canceling maintenance for hosts; Supports viewing basic information of a single host, host resource usage information, and file system information on the host; Supports viewing role instances installed on a single host and alarm history.
Monitoring and Alarm	Supports querying monitoring indicators at the role instance level and host level; Supports the graphing of metric results, enabling users to intuitively understand changes in metrics; Supports the querying of alarm content at the cluster service level, role instance level, and host level.
O&M and Configuration	Supports actions such as modification, synchronization, and rollback of configuration files for each cluster service; Supports viewing the O&M history of all pipelines and the operator.

Check the Version of YI-MapReduce Manager

Follow these steps:

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab.

4. Click **Go to YI-MapReduce Manager** to enter the Manager interface. The version number of the Manager is displayed in the lower-left corner of the interface.

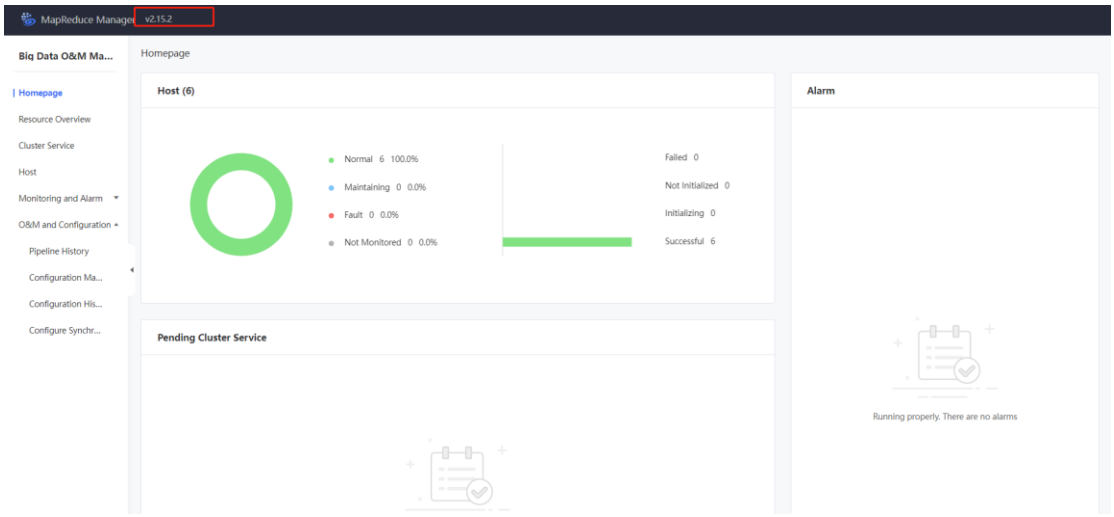

4.5.2. Homepage

Homepage Overview

Upon entering the YI-MapReduce Manager, you are taken to the Homepage, as shown below:

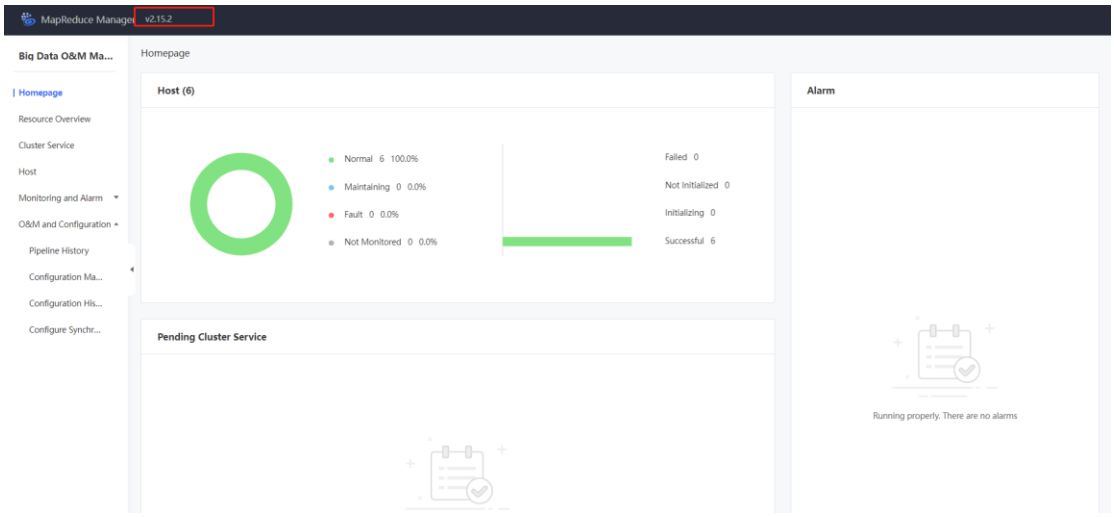

- Host: Displays summary of the health status and initialization of the platform host.
- Clusters Pending Processing: Displays information about clusters that have faults or alarms.

- Alarms: Displays alarms that require attention and allows viewing of alarms at different levels.

4.5.3. Resource Overview

Resource Overview

Once in the YI-MapReduce Manager, click the **Resource Overview** menu to go to the resource overview page.

The resource overview displays the information for all hosts on the current platform over the last five hours, including CPU utilization, disk utilization, memory utilization, network send speed, and network receive speed. As shown in the figure:

4.5.4. Cluster Service

Cluster Service Management Overview

Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service page, as shown below:

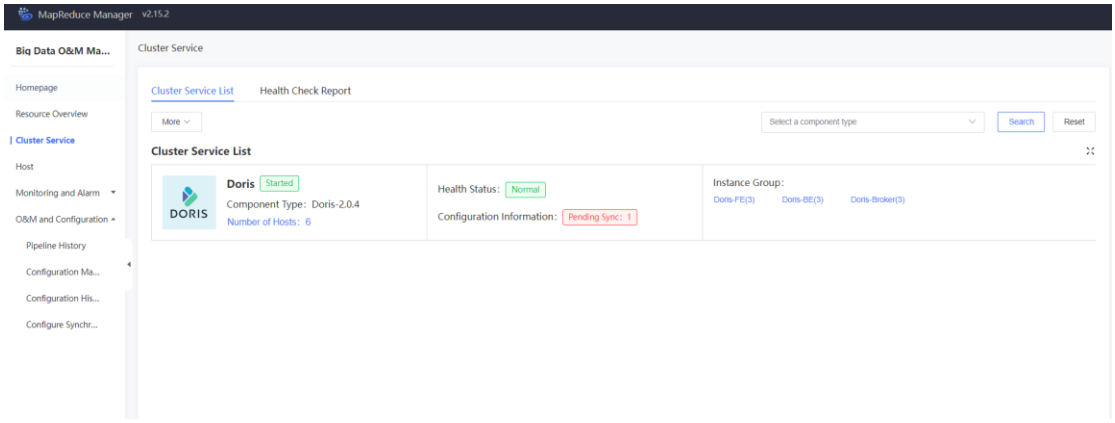

- The query function at the top of the cluster services list allows querying by component type.
- The cluster services list is sorted by component type, displayed in a list view.
- In the top right corner of the cluster services list, the following operations are supported: start all cluster services, stop all cluster services, perform a health check.
- Click on the cluster service icon or the name of the cluster service to enter the cluster details page.
- Click on the number of hosts to go to the host menu page.
- Click on **Instance Groups** to go to the **Cluster Details - Role Instance** page.

Cluster Service Details Page Overview

Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.

Click on the name of a specific cluster service to view its details, as shown below:

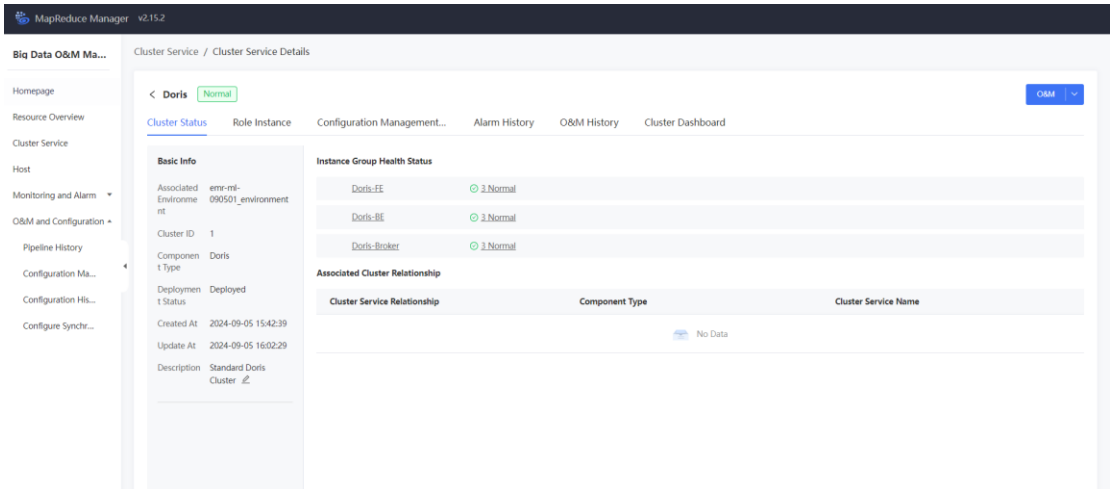

- The upper section of the details page displays the name and health status of the cluster service.
- The upper right corner of the details page includes O&M operations for managing the cluster service, such as starting the cluster service, stopping the cluster service, and performing a rolling restart of the cluster service.
- The upper section of the details page presents the following information in order: the status of the cluster, role instances, configuration management, O&M history, alarm history, the cluster dashboard, and more.

Start and Stop All Cluster Services

Start All Cluster Services

Scenario

Once-click start all cluster services.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.

4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Click on **More Operations > Start All Cluster Services**. A confirmation pop-up will appear.
6. Click **OK**

   to initiate the process of starting all cluster services. Wait for all cluster services to complete their start-up process, then click **Finish**.

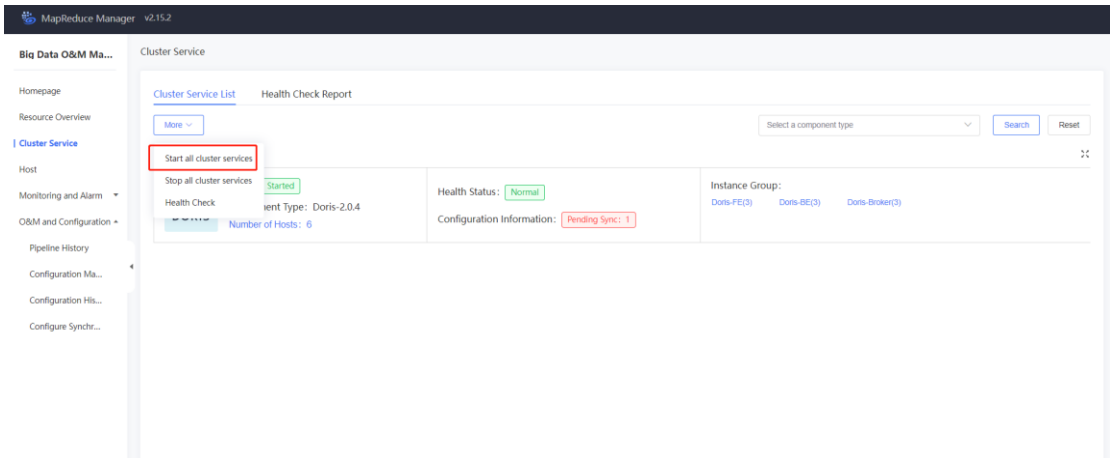

✧ **Note:**

- When starting all cluster services, they are initiated sequentially according to cluster dependency relationships.
- During the execution process, you can cancel the current operation. However, cluster services that have already started cannot be halted.

Stop All Cluster Services

Scenario

Once-click stop all cluster services.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Click on **More Operations > Stop All Cluster Services**. A confirmation pop-up will appear.
6. Click

   to initiate the process of stopping all cluster services. Wait for all cluster services to complete their start-up process, then click **Finish**.

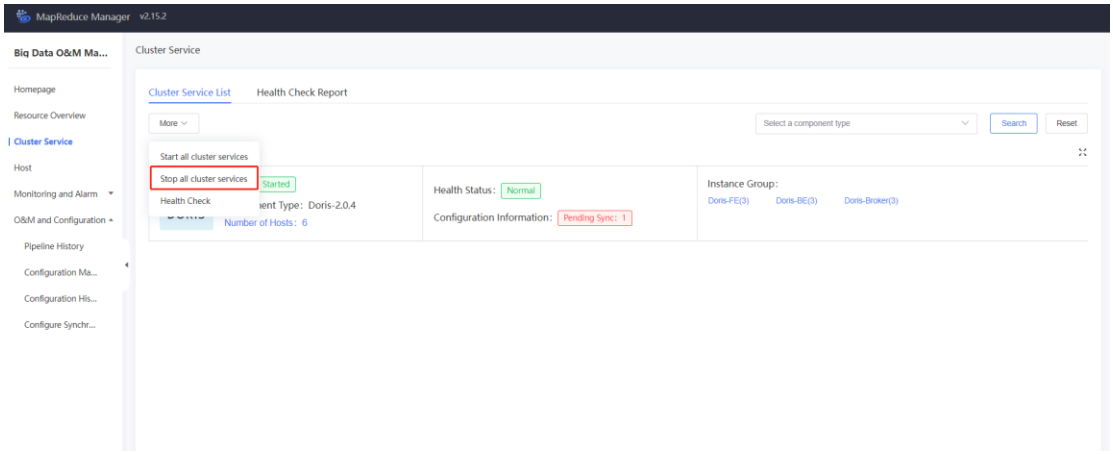

✧ **Note:**

- When stopping all cluster services, they are initiated sequentially according to cluster dependency relationships.
- During the execution process, you can cancel the current operation. However, cluster services that have already stopped cannot be halted.

Start and Stop a Single Cluster Service

Start Cluster Service

Scenario

Once-click start a single cluster service.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click on **O&M Operations > Start Cluster Service**. A confirmation pop-up will appear.
7. Click OK to enter the start cluster service page, showing a list of instances that the present cluster service can start. As shown in the figure:

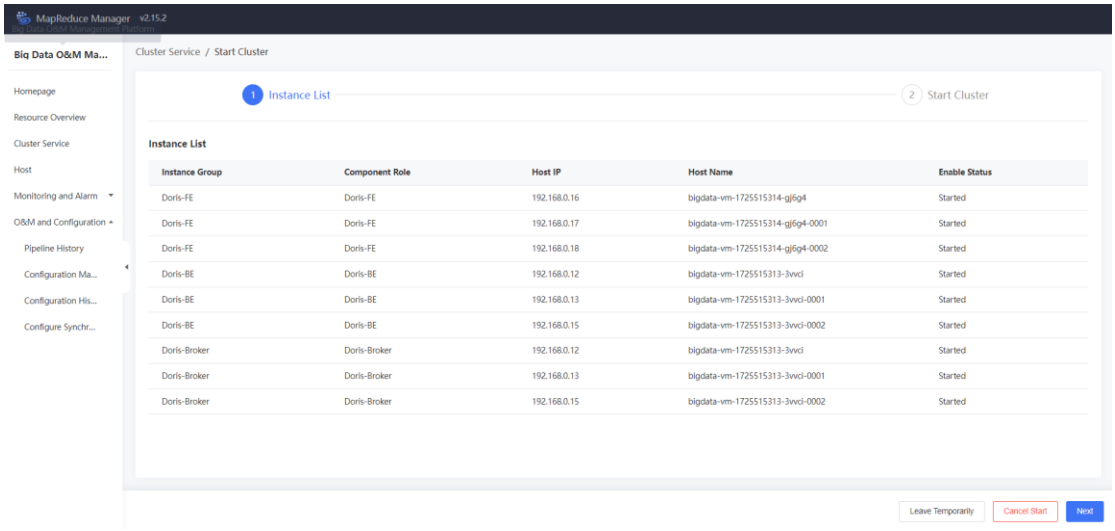

8. Click **Next** to start instances in the order of instance groups. Once all instances have successfully started, click **Finish**. As shown in the figure:

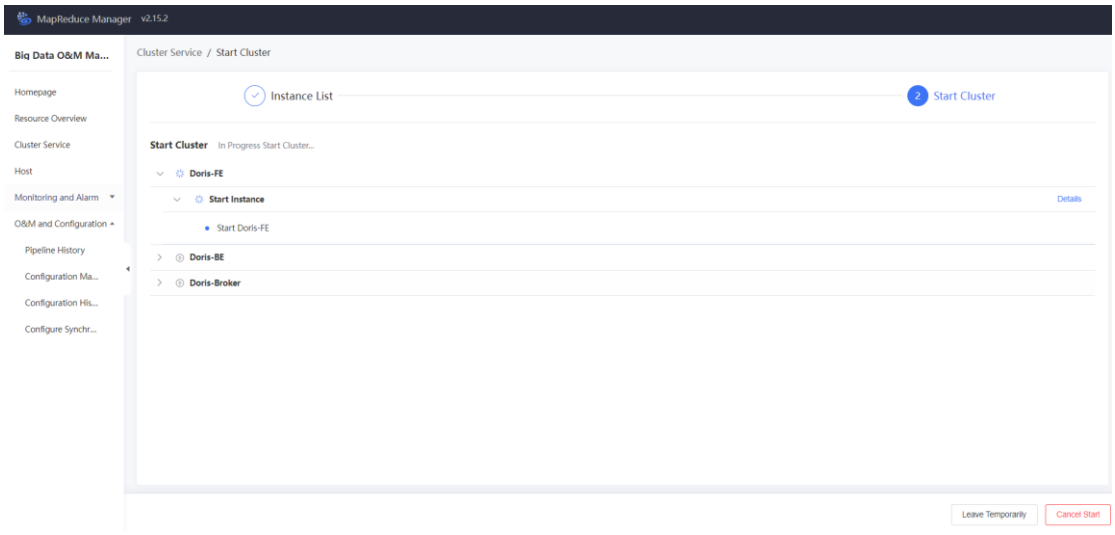

Stop Cluster Service

Scenario

Once-click stop a single cluster service.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.

3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click on **O&M Operations > Stop Cluster Service**. A confirmation pop-up will appear.
7. Click **OK** to enter the stop cluster service page, showing a list of instances that the present cluster service can stop.
8. Click **Next** to stop instances in the order of instance groups. Once all instances have stopped, click **Finish**.

Start, Stop, Restart Instance

Start Instance

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click on the **Role Instance** tab.
7. Check the box in front of the role instance that requires operation.
8. Click on **Operate on Selected > Start Instance**. A confirmation pop-up will appear.

9. Click **OK**

to enter the start instance page, showing a list of instances that the present cluster service can start. As shown in the figure:

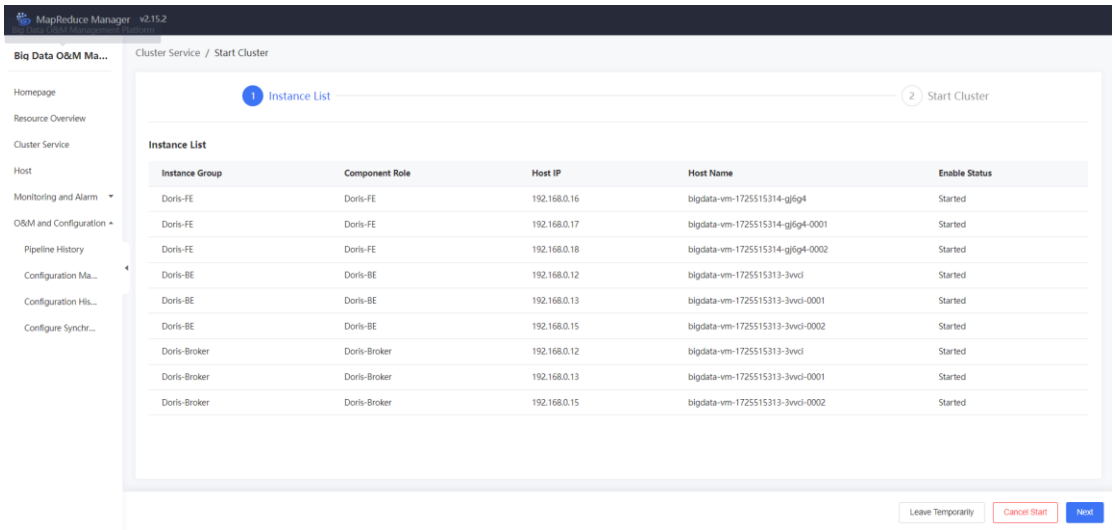

10. Click **Next** to start instances in the order of instance groups. Once all instances have successfully started, click **Finish**. As shown in the figure:

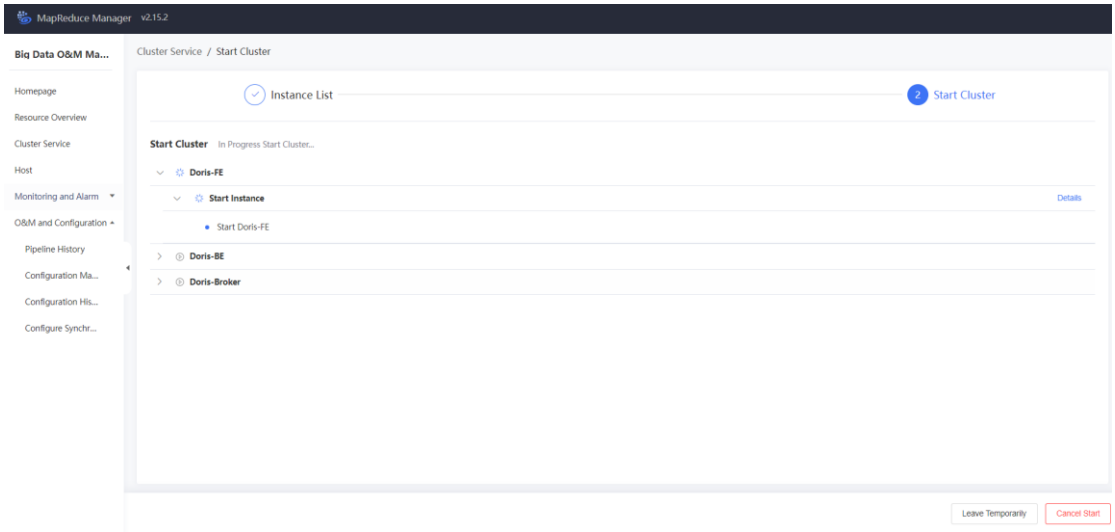

Stop Instance

Procedure

1. Log in to the YI-MapReduce management console.

2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click on the **Role Instance** tab.
7. Check the box in front of the role instance that requires operation.
8. Click on **Operate on Selected > Stop Instance**. A confirmation pop-up will appear.
9. Click **OK** to enter the stop instance page, showing a list of instances that the present cluster service can stop.
10. Click **Next** to stop instances in the order of instance groups. Once all instances have stopped, click **Finish**.

Restart Instance

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click on the **Role Instance** tab.
7. Check the box in front of the role instance that requires operation.

8. Click on **Operate on Selected > Restart Instance**. A confirmation pop-up will appear.
9. Click **OK** to enter the restart instance page, showing a list of instances that the present cluster service can restart.
10. Click **Next** to perform instance restart in the order of instance groups. Once all instances have completed the restart, click **Finish**.

Instance Maintenance and Cancellation

Set Maintenance

Scenario

Users may need to set a specific instance into maintenance mode.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click on the **Role Instance** tab.
7. Check the box in front of the role instance that requires operation.
8. Click on **Operate on Selected > Set Maintenance**. A set maintenance pop-up will appear. As shown in the figure:

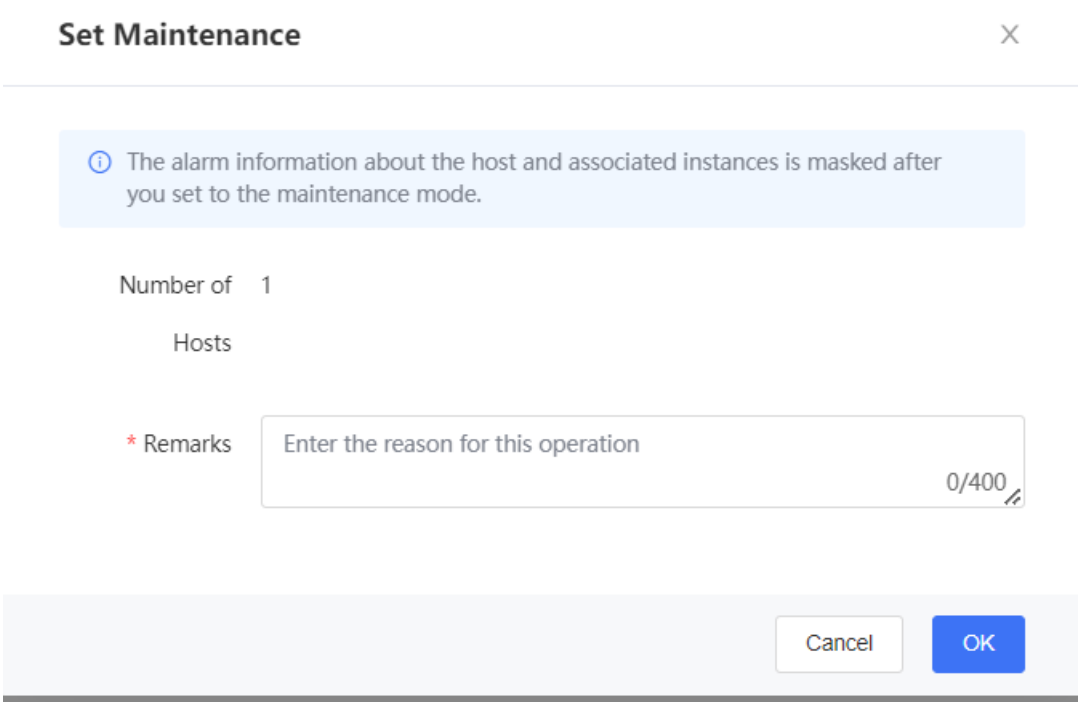

9. Enter notes and click

   OK.

   Note:

   – After setting the instance to maintenance, the health status of this instance will be displayed as in maintenance.

   – After setting the instance to maintenance, the alarm messages of this instance will be blocked.

Cancel Maintenance

Scenario

Users may need to cancel the maintenance status for a specific instance.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.

3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click on the **Role Instance** tab.
7. Check the box in front of the role instance that requires operation.
8. Click **Operate on Selected > Cancel Maintenance**. The cancel maintenance operation pop-up will appear. As shown in the figure:

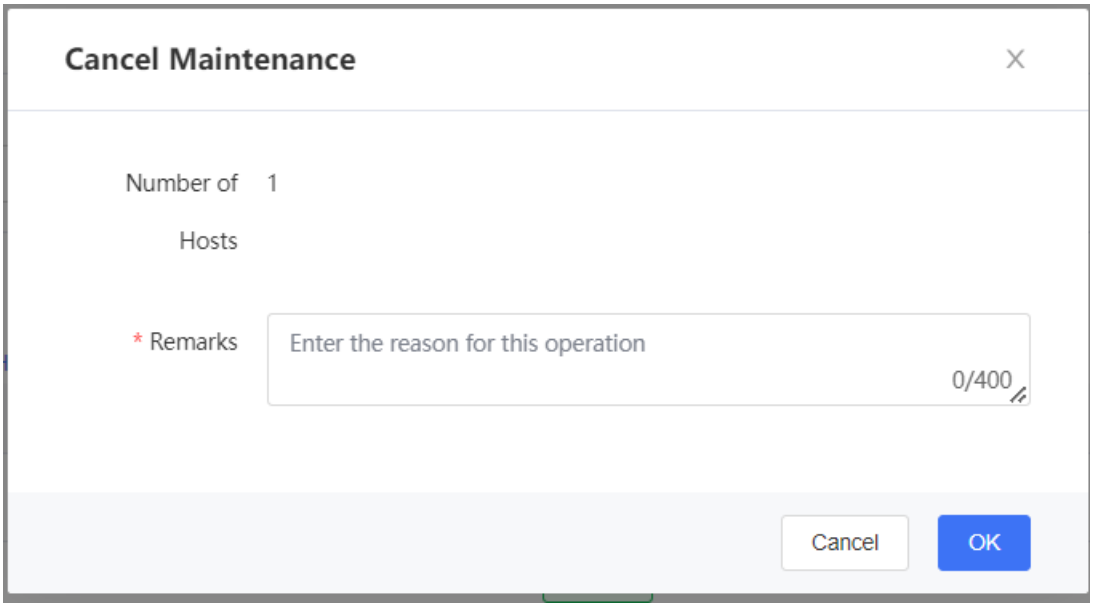

9. Enter notes and click **OK**.

.

View Cluster Service Configurations

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.

3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab to view all configurations for this cluster service. As shown in the figure:

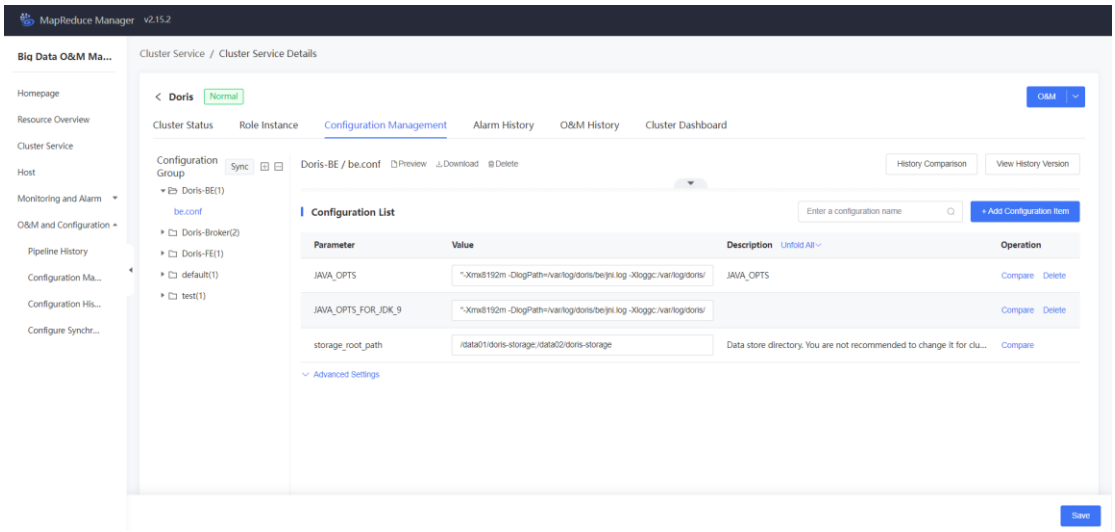

Add Configuration Group

Scenario

Users can classify hosts within a cluster service based on different grouping criteria.

All hosts within a configuration group share a common set of configurations.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.

3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Click on **Add Configuration Group** to go to the add configuration group page.
8. Choose the grouping method as needed, select the host, and click the **Save** button at the bottom right corner.

As shown in the figure:

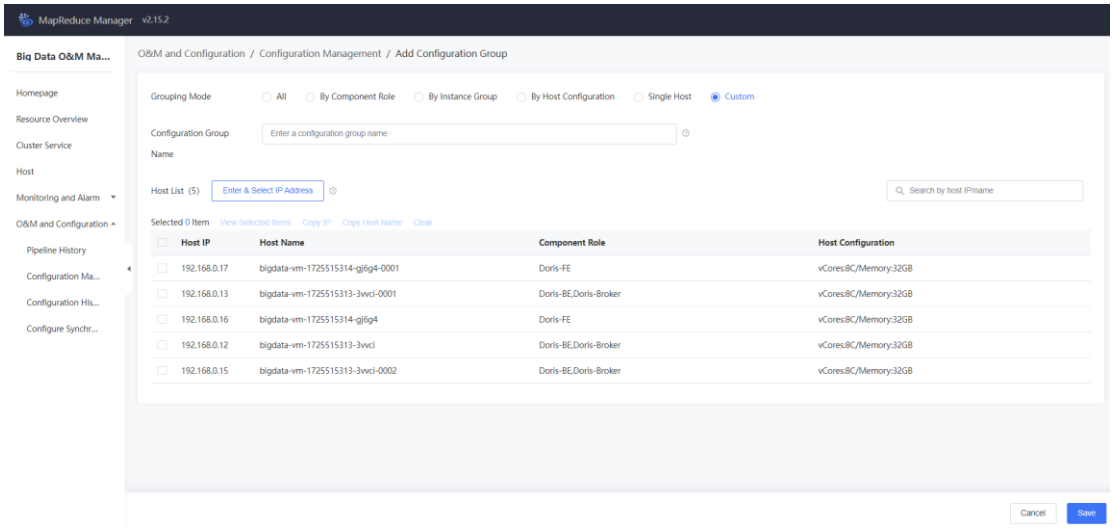

✧ **Note:**

- When the grouping method is "All", "By Component Role", "By Instance Group", or "By Host Configuration", the configuration group name is predefined and cannot be changed.

- When the grouping method is "Single Host", you need to select the host to be grouped. The configuration group name is the name of the selected host and cannot be changed.

- When the grouping method is "Custom", you need to select the host to be grouped. The configuration group name can be customized but must not be the same as other configuration group names.

View Configuration Group Details

Scenario

Users can view details of different configuration groups to learn about the host information contained within them.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Choose the configuration group you want to view. Hover over the configuration group name, and click on the **Details** icon on the right to access the configuration group details page. As shown in the figure:

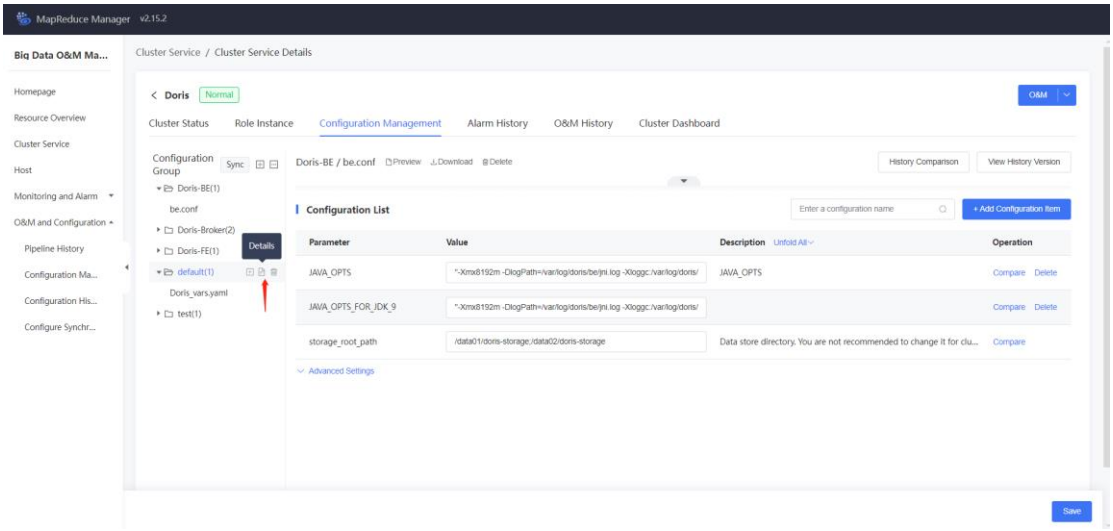

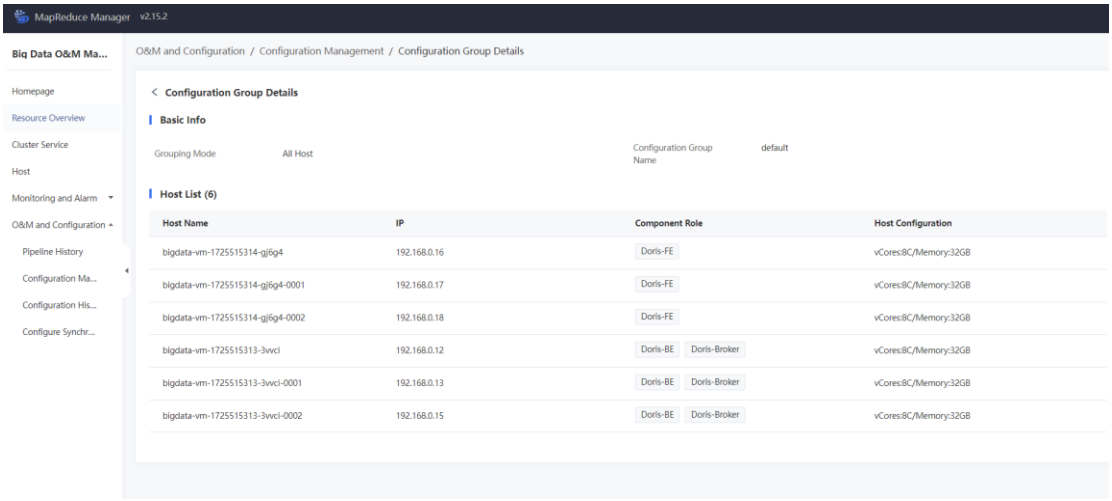

Delete Configuration Group

Scenario

Users can remove configuration groups no longer needed.

Before You Begin

The number of configuration files in the configuration group must be 0 to proceed with the deletion.

Procedure

1. Log in to the YI-MapReduce management console.

2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Choose the configuration group you want to delete. Hover over the configuration group name, and click on the **Delete** icon on the right. A deletion confirmation pop-up will appear.

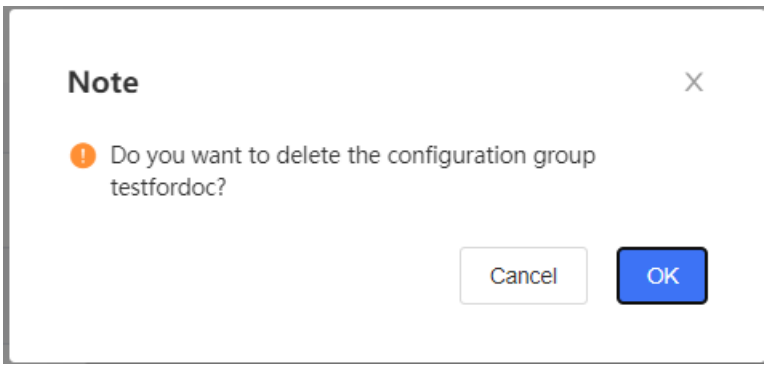

8. Click **OK** to delete the configuration group.

Add Configuration

YI-MapReduce Manager supports four methods to add new configurations.

Method 1: Create New

Scenario

When there is no need to share configurations with other clusters, and the configuration template does not meet usage requirements, users can add a completely new configuration file for the cluster service.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Choose the configuration group for adding new configurations. Hover over the configuration group name, and click on the **Add Configuration** icon on the right. The right side of the page will show the new configurations. As shown in the figure:

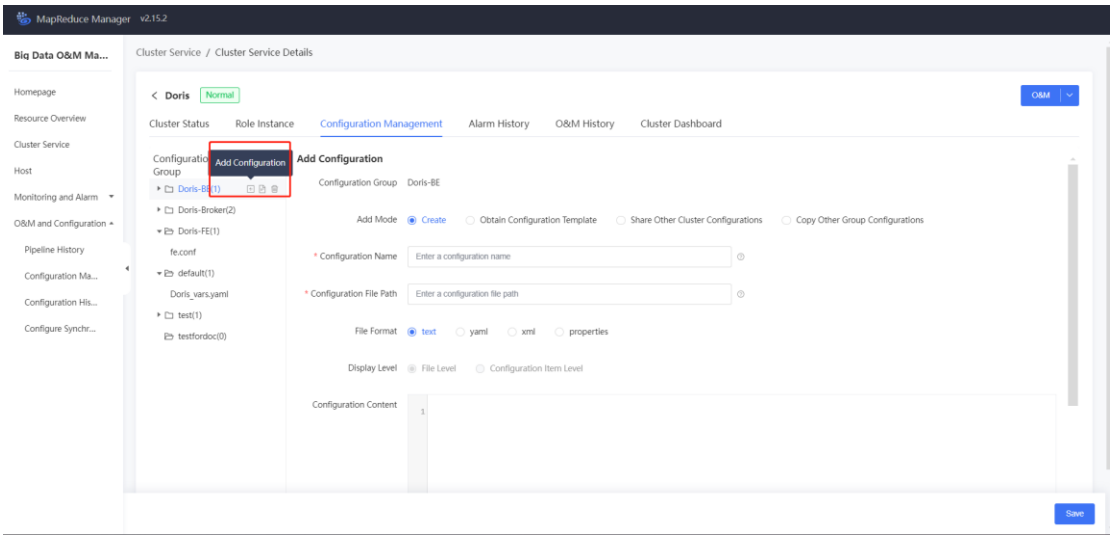

8. Select Create New as the method for adding, enter the configuration name, and configuration file path, choose the file format, and input the configuration content or add configuration items.
9. Click **Save** to complete the configuration addition.

Cautions

1. Configuration Name: Only English letters, numbers, and symbols like "_-." are supported. It cannot start with a number.
2. File Path: Only English letters, numbers, and symbols like "_-." are supported. The configuration path must be included in the trustlist.
3. When the file format is text, you need only to input the configuration content. As shown in the figure:

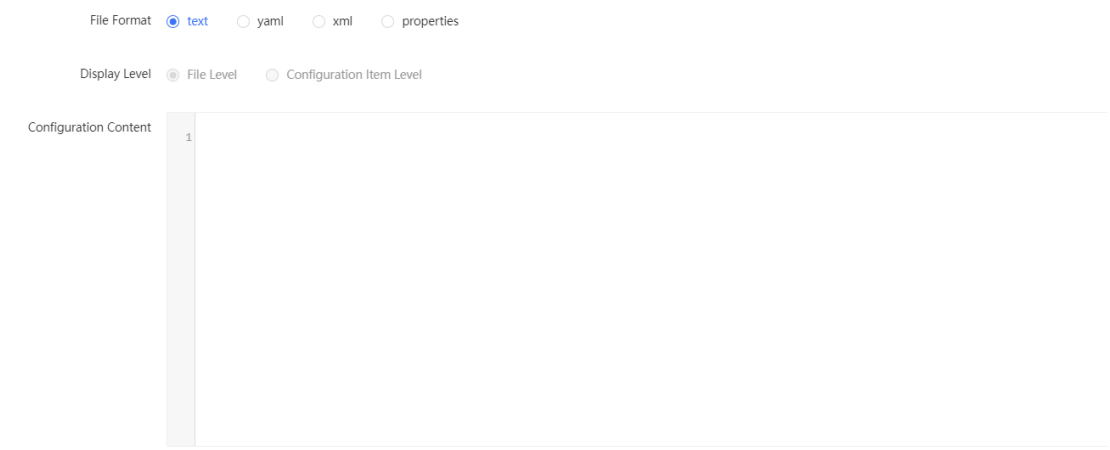

4. For non-text file formats like yaml, xml, or properties, you are required to: select the display level, choose the configuration rule group, and set the configuration items. As shown in the figure:

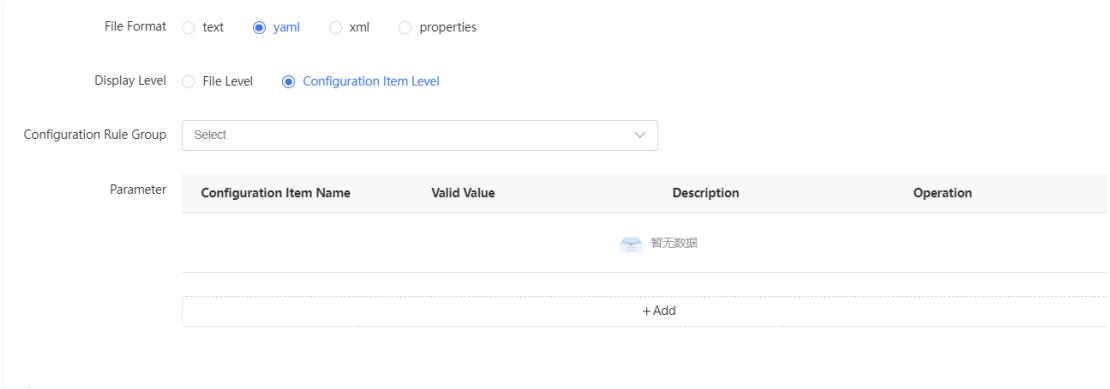

5. If the display level is set at the configuration item level, the configuration will be displayed as individual configuration items. As shown in the figure:

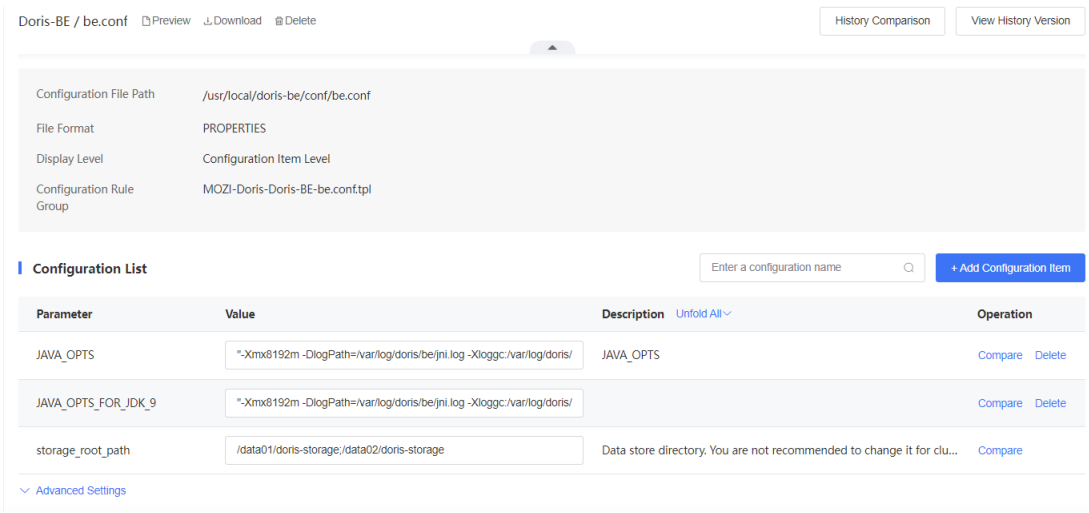

6. If it is set at the file level, the configuration will be displayed as text. As shown in the figure:

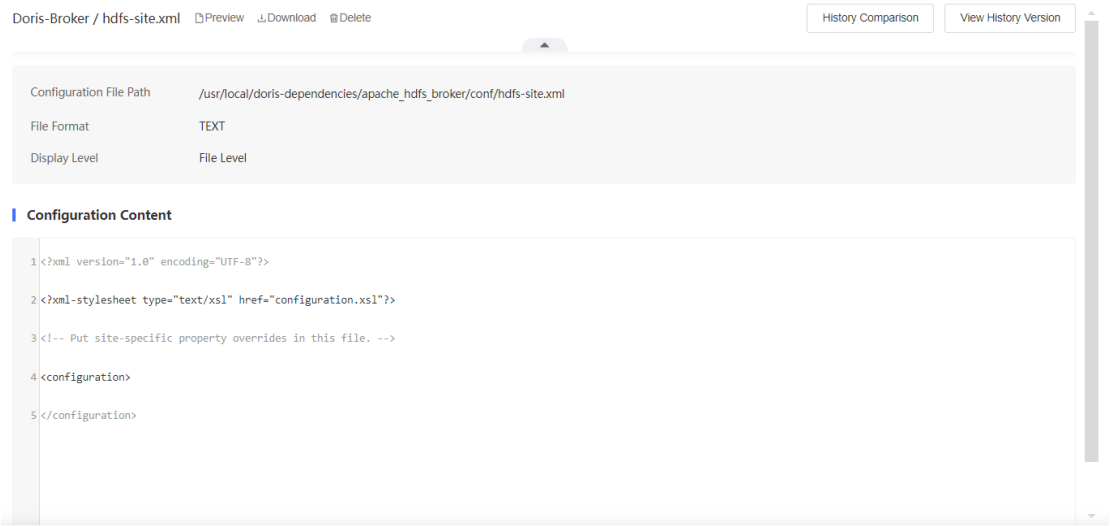

7. Configuration Rule Group: The configuration rule maintained in the system. After selecting the configuration rule group, the configuration items below will display the configuration items included in the configuration rule group. You can edit the configuration values and descriptions of the configuration items brought out by the configuration rule group and continue to add new configuration items.

Method 2: Retrieve Configuration Templates

Scenario

Users may want to use configuration templates directly when adding new configurations to the cluster service.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Choose the configuration group for adding new configurations. Hover over the configuration group name, and click on the **Add Configuration** icon on the right. The right side of the page will show the new configurations.
8. Choose Retrieve Configuration Template as the adding method. As shown in the figure:

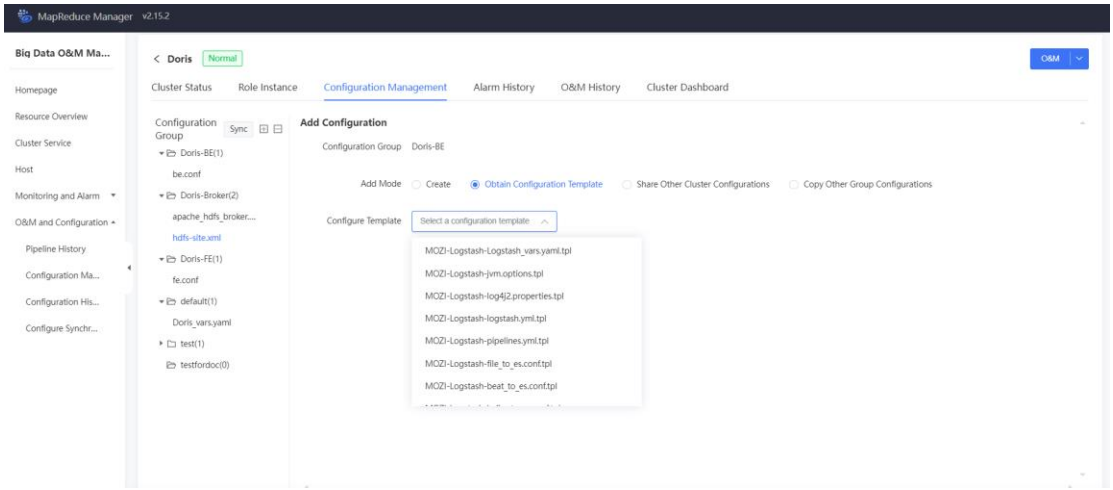

9. Select the configuration name.
10. Click **Save** to complete the configuration addition.

Method 3: Share Configuration from Other Clusters

Scenario

When the service of the current cluster relies on the configuration of other clusters, users can directly share the configuration of other cluster services.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Choose the configuration group for adding new configurations. Hover over the configuration group name, and click on the **Add Configuration** icon on the right. The right side of the page will show the new configurations.
8. Choose Share Configuration from Other Clusters as the adding method. As shown in the figure:

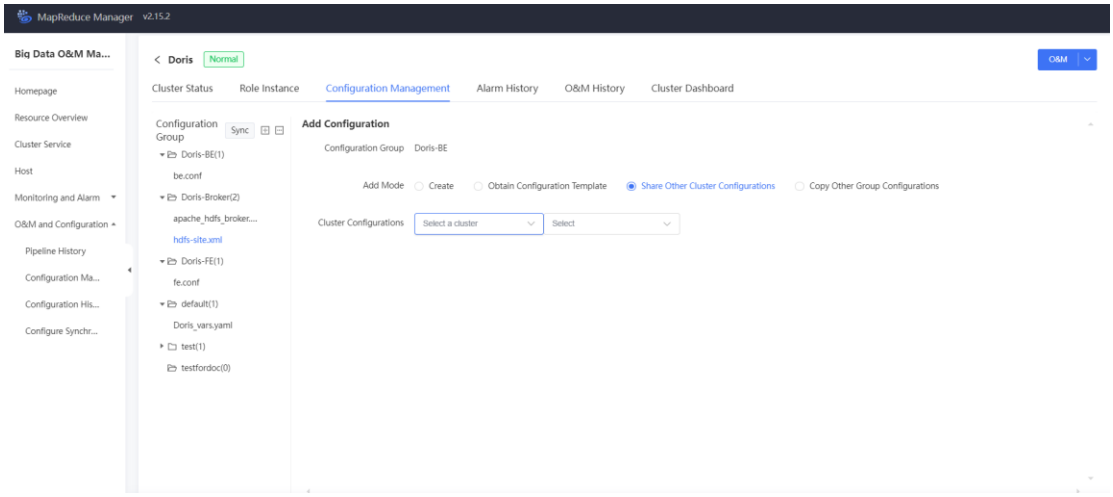

9. Select the cluster and configuration you want to share.
10. Click **Save** to complete the configuration addition.

Method 4: Copy Configuration from Other Groups

Scenario

If the desired configuration is already present in another group, there is no need to create a new one. Instead, it can be directly copied from the other group to the current group.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.

7. Choose the configuration group for adding new configurations. Hover over the configuration group name, and click on the **Add Configuration** icon on the right. The right side of the page will show the new configurations.

8. Choose Copy Configuration from Other Groups as the adding method. As shown in the figure:

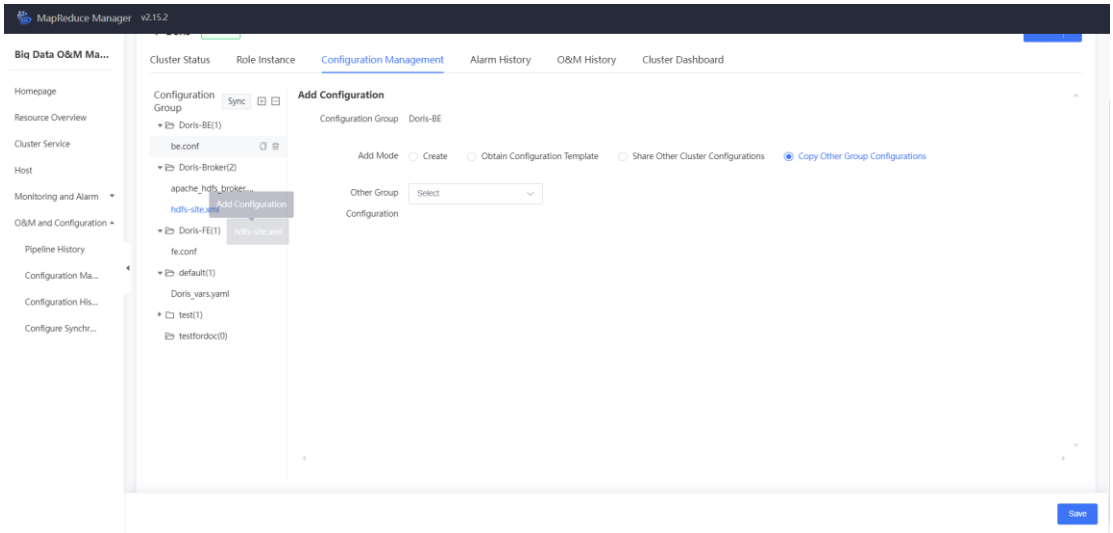

9. Select the configuration group and configuration to copy.

10. Click **Save** to complete the configuration addition.

Cautions

Modifying the copied configuration file, or altering the original configuration from which it was copied, does not impact either configuration.

View Configuration Details

Scenario

Users can view the details of a certain configuration.

Procedure

1. Log in to the YI-MapReduce management console.

2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.

3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Click on the configuration group, and then on the name of the configuration you want to view. Details of this configuration will be displayed on the right side of the page. As shown in the figure:

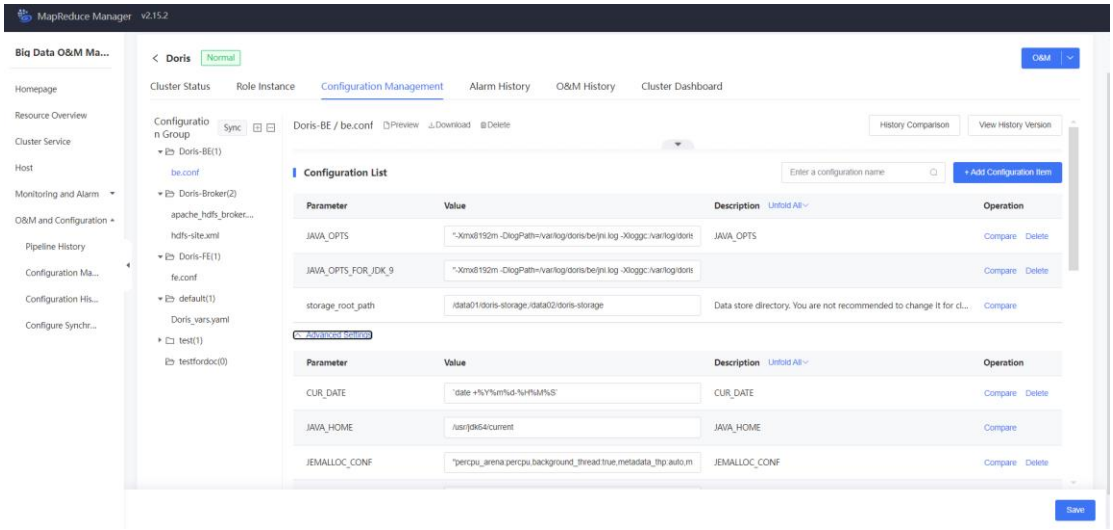

Edit Configuration

Scenario

Users can modify configuration files.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.

3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Click on the configuration group, and then on the name of the configuration you want to modify. Details of this configuration will be displayed on the right side of the page. As shown in the figure:

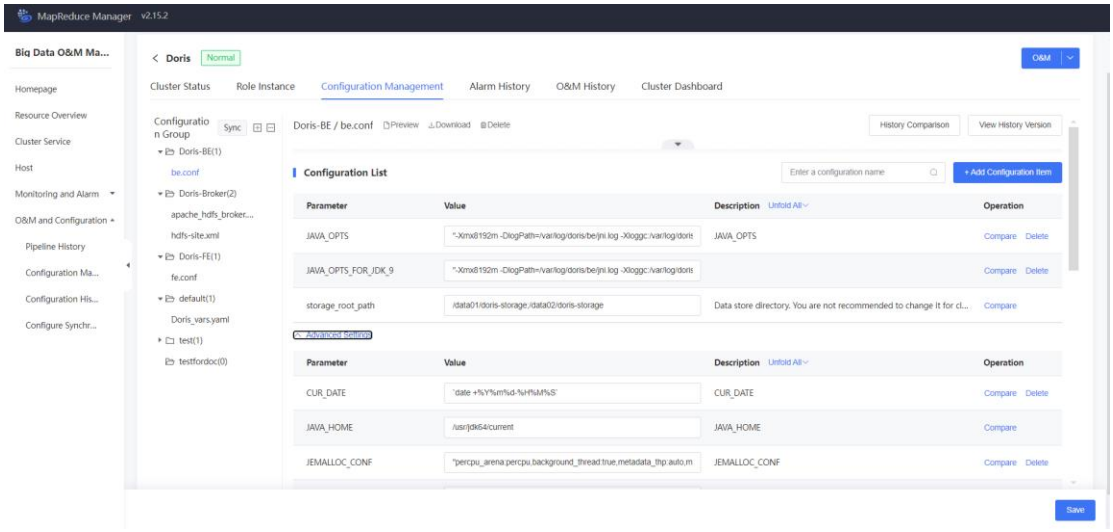

8. You can modify the configuration file directly. After completing the modifications, click **Save**.

Cautions

1. Display level: File level.

- Direct editing or modification is allowed in the text box.
- After modifications, if the current configuration file has not been synchronized, the number 1 will appear before the configuration name. As shown in the figure:

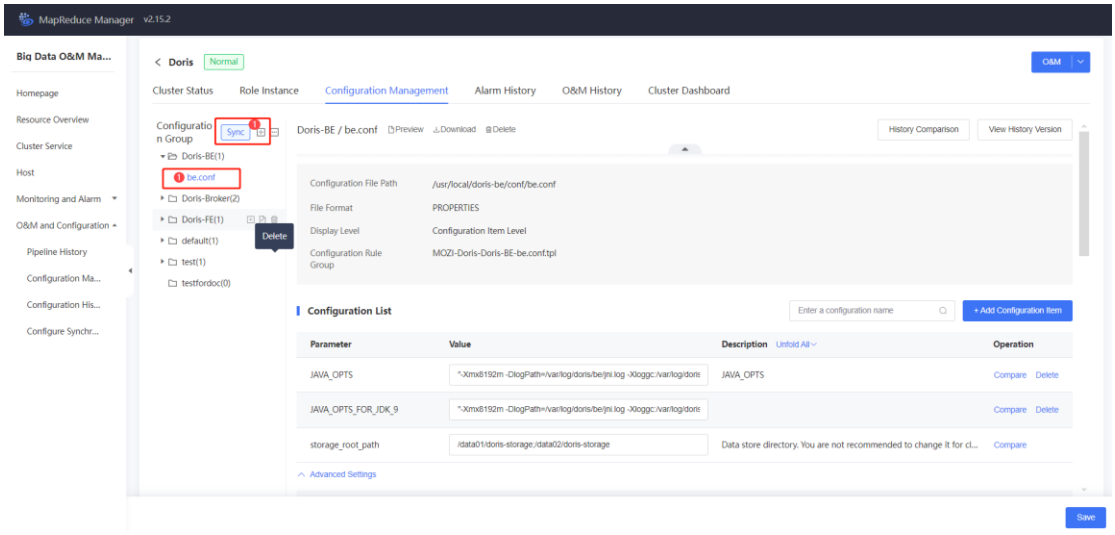

2. Display level: Configuration item level.

- Configuration items need to be modified individually.
- After modifications, if the current configuration file has not been synchronized, the number of modified configuration items will be displayed before the configuration name, and a modification mark will appear before the name of the modified configuration item. As shown in the figure:

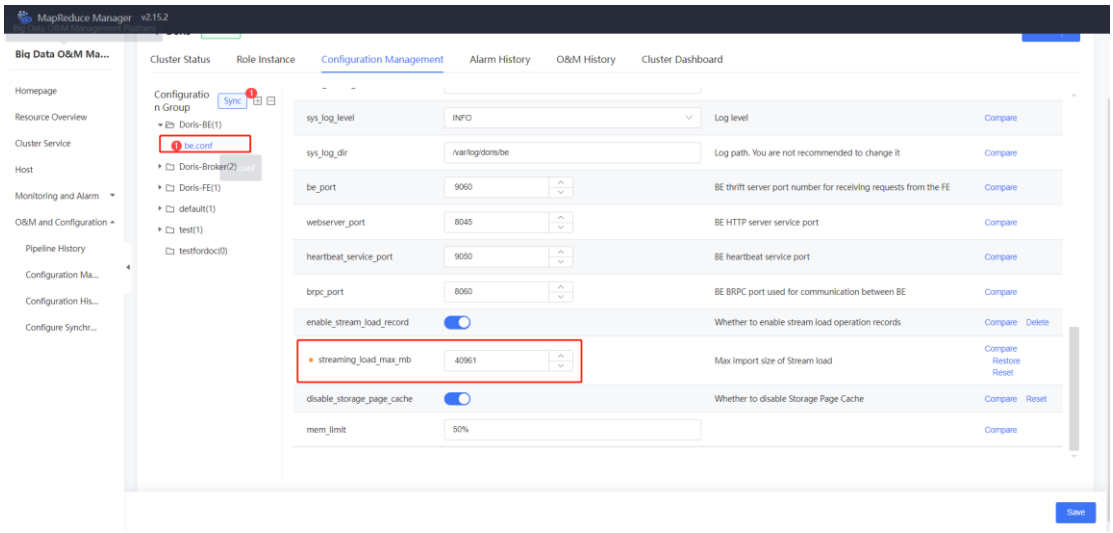

Configuration Item Comparison

Scenario

Users can view the values of different configuration groups for the same configuration item within the same cluster.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Click on the configuration group, and then on the name of the configuration file. Details of this configuration will be displayed on the right side of the page.
8. Select the configuration item to compare, and click the **Configuration Item Comparison** button on the right side of the name of the configuration item. A configuration item comparison dialog box will appear. As shown in the figure:

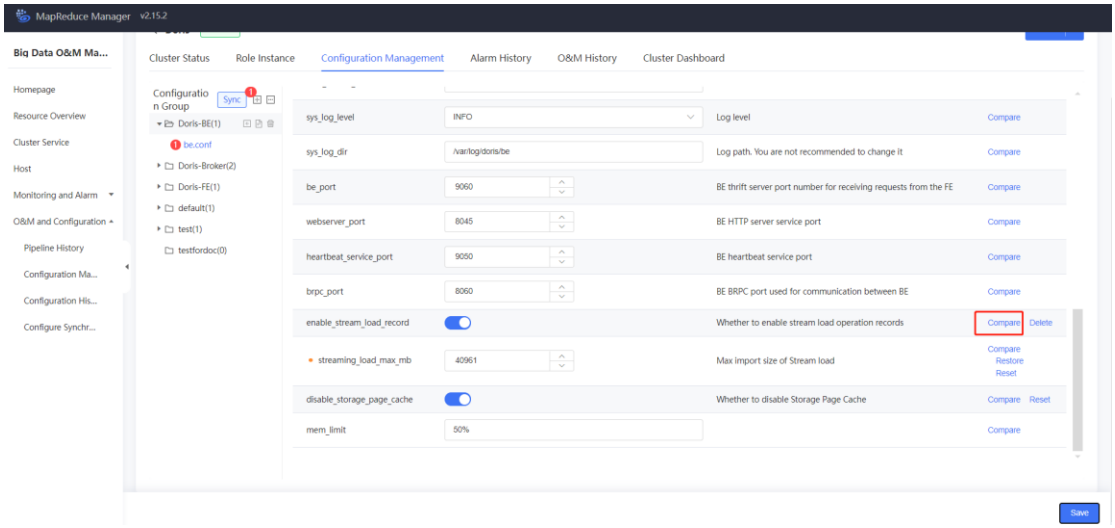

Configuration File Preview

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Click on the configuration group, and then on the name of the configuration file. Details of this configuration will be displayed on the right side of the page.
8. Click **File Preview** and a configuration file preview dialog box will appear. As shown in the figure:

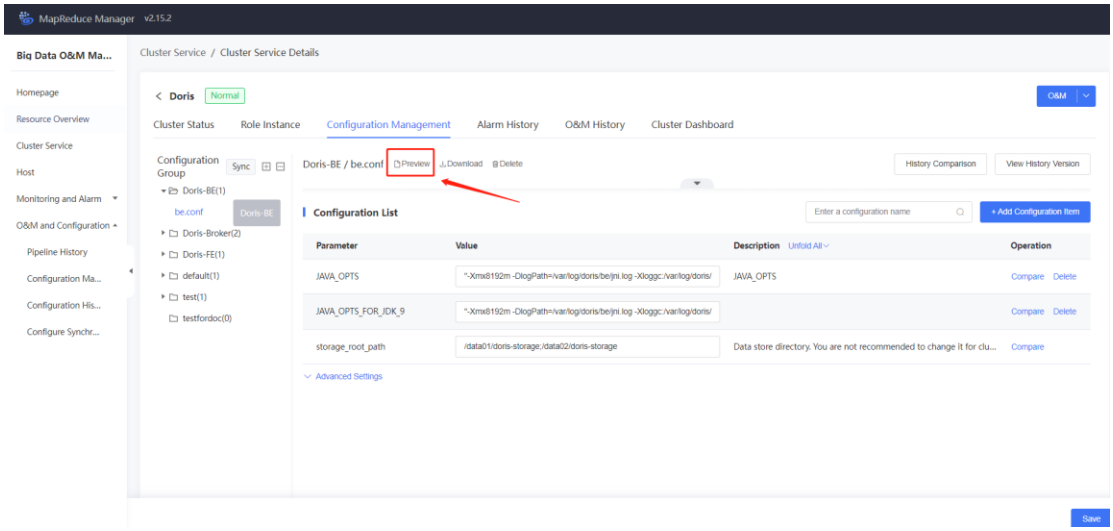

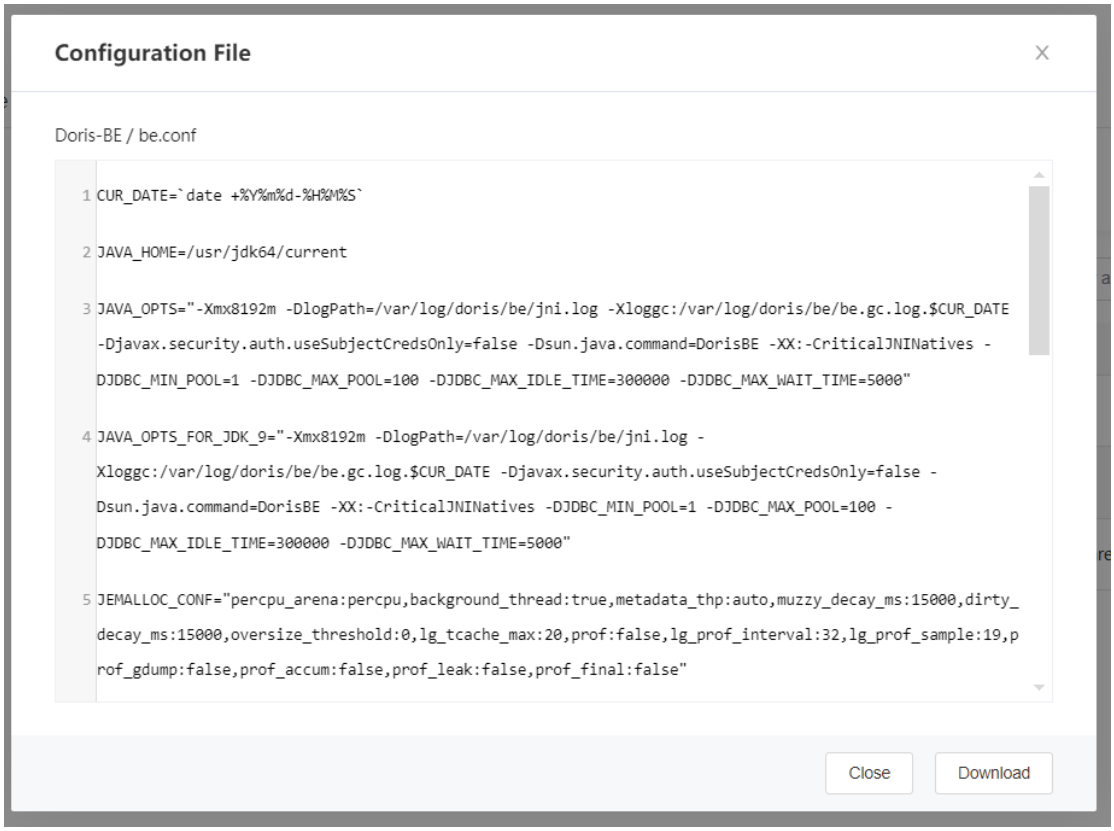

Download Configuration Files

Procedure

1. Log in to the YI-MapReduce management console.

2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Click on the configuration group, and then on the name of the configuration file. Details of this configuration will be displayed on the right side of the page.
8. Click **Download Files**, and wait for the download to complete. As shown in the figure:

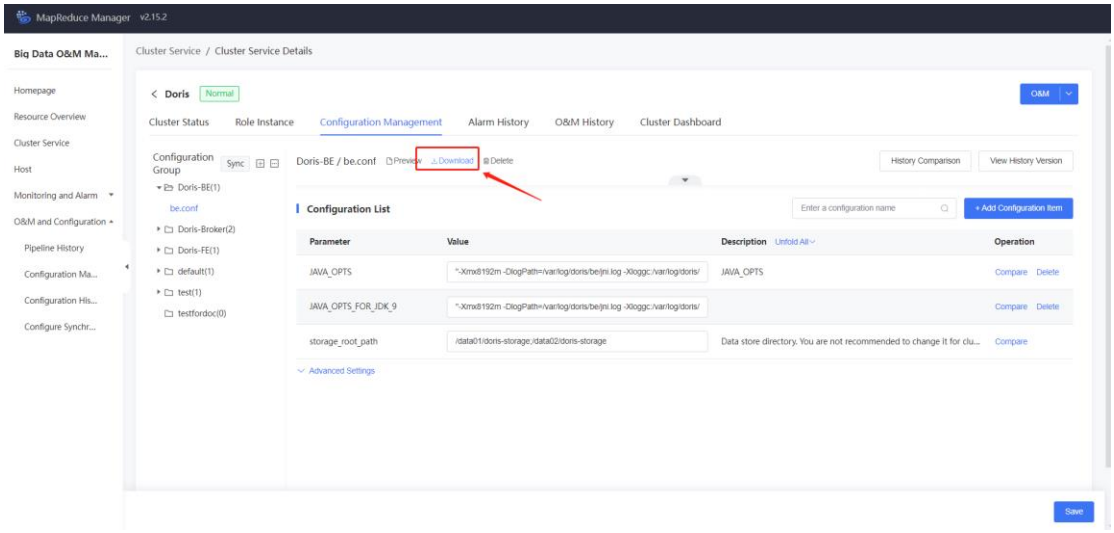

Configuration File History Comparison

Scenario

This feature allows users to compare the current configuration file with historical versions.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Click on the configuration group, and then on the name of the configuration file. Details of this configuration will be displayed on the right side of the page.
8. Click on **More > History Comparison**, and a history comparison dialog box will appear. As shown in the figure:

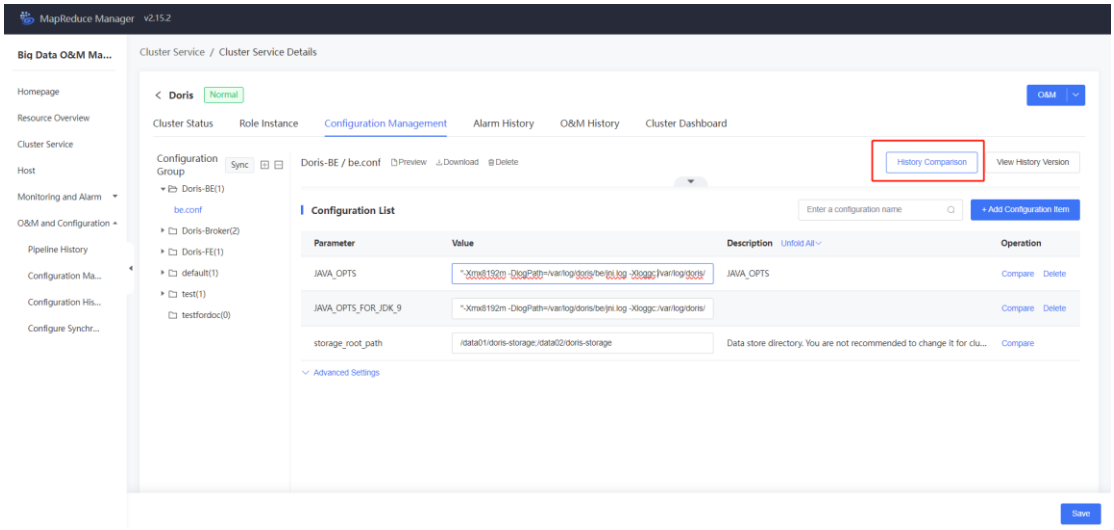

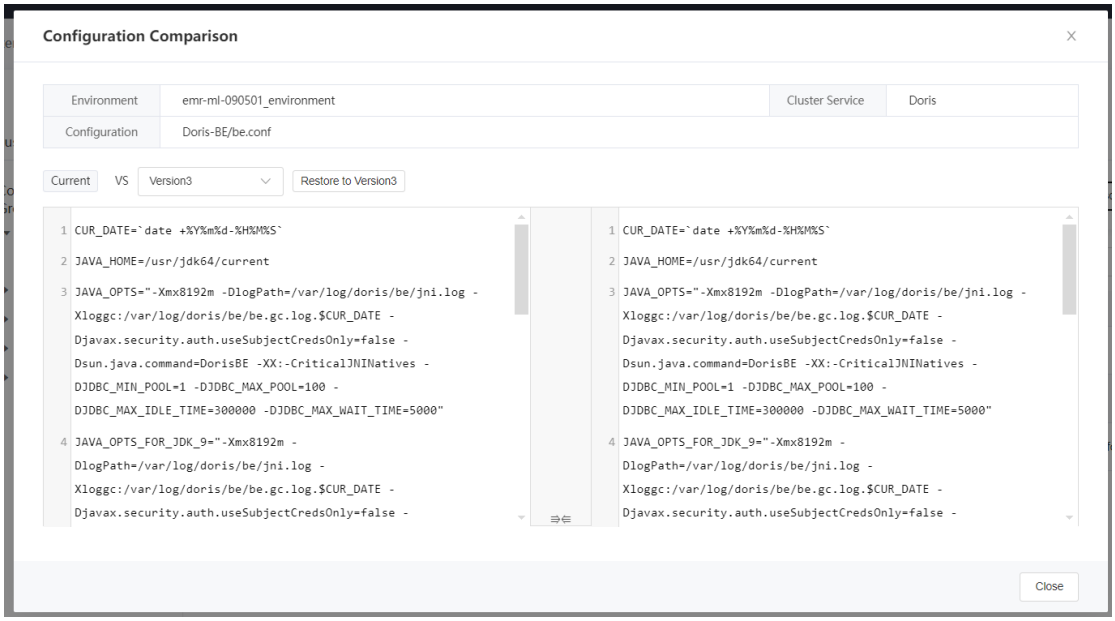

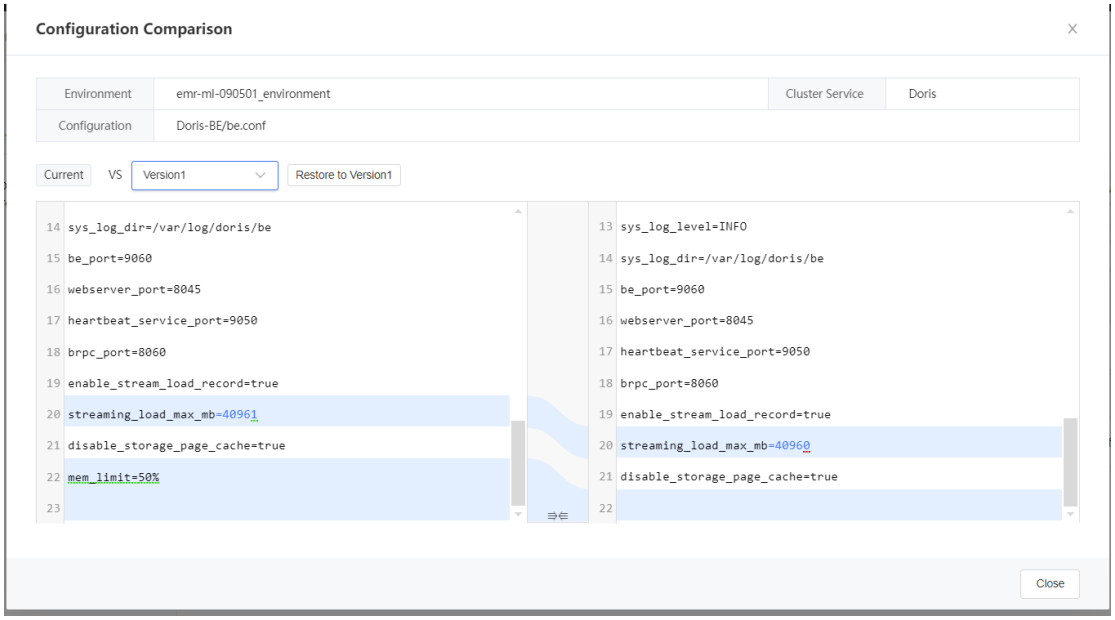

View Configuration History Versions

Scenario

This feature allows users to view the history versions of a specific configuration file, along with their contents and changes between different versions.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Click on the configuration group, and then on the name of the configuration file. Details of this configuration will be displayed on the right side of the page.
8. Click **More > View History Versions**. As shown in the figure:

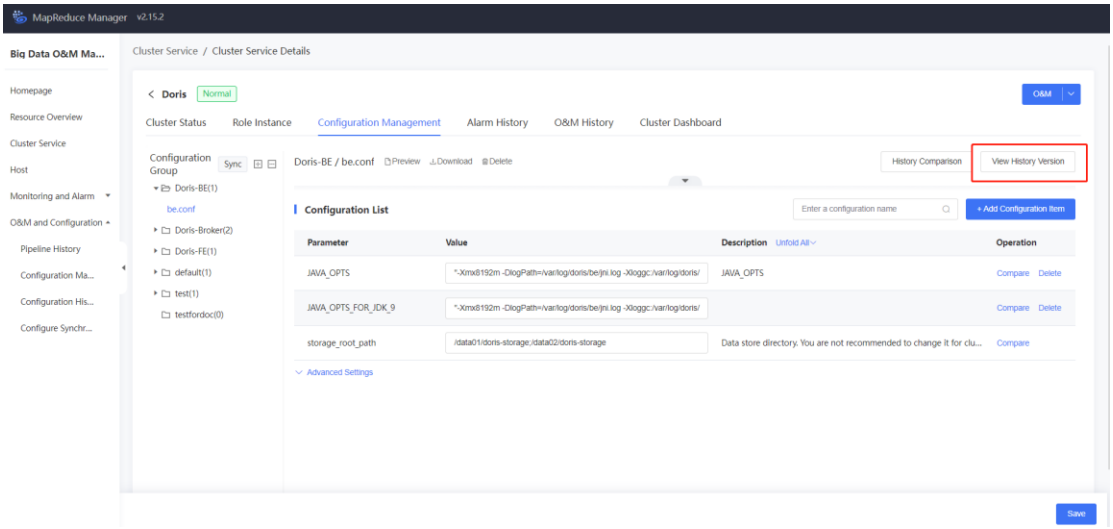

9. You will enter the configuration history page, which displays all history versions of the current configuration. As shown in the figure:

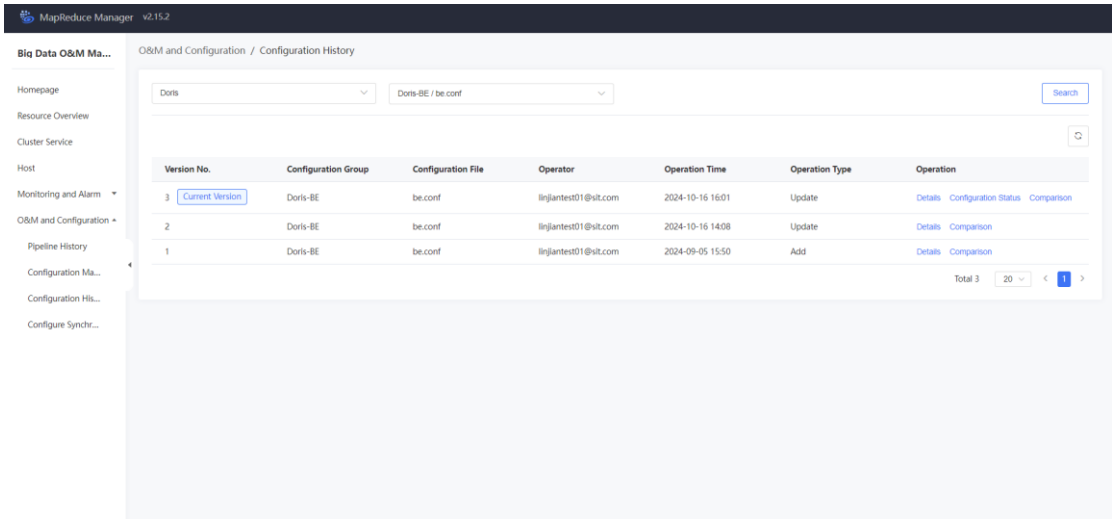

Delete Configuration

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Click on the configuration group, and then on the name of the configuration file. Details of this configuration will be displayed on the right side of the page.
8. A confirmation pop-up for deletion will appear. Click **OK**. As shown in the figure:

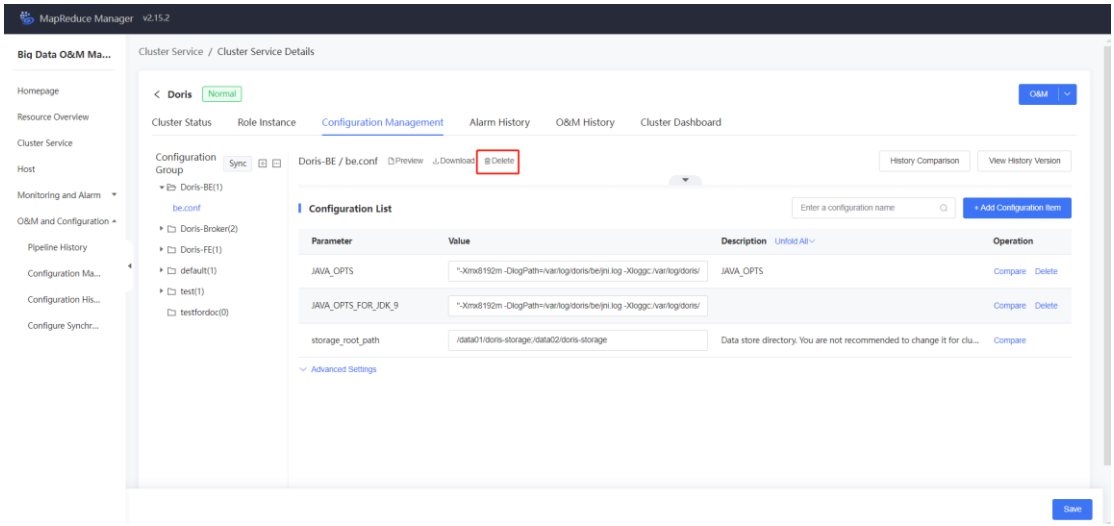

9. Click **Delete Configuration**, then click **OK** to complete the operation.

Configuration Sync

Scenario

This feature supports the synchronization of new or modified configuration files to the respective hosts.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Click the **Sync** button. As shown in the figure:

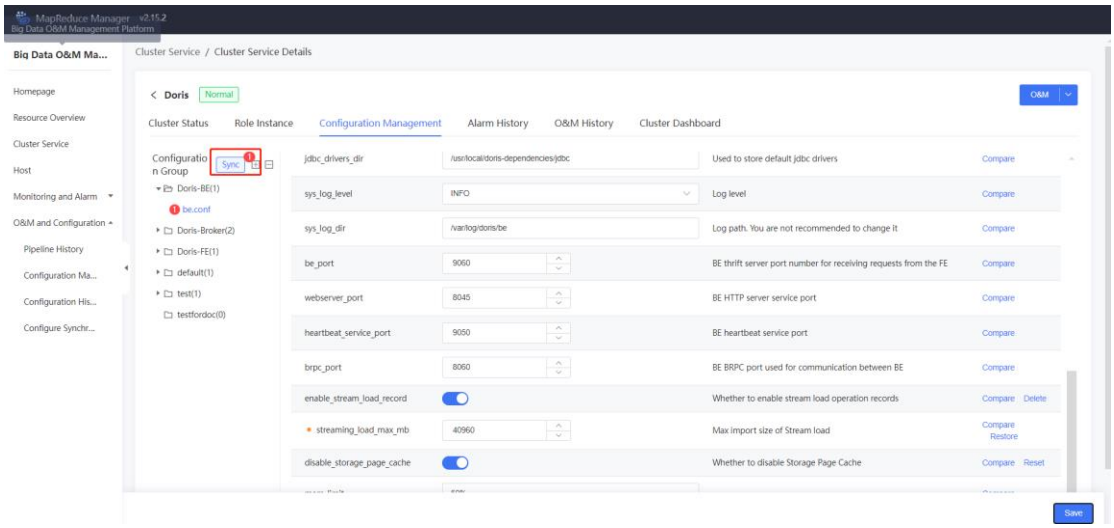

8. Enter the page of the configuration to be synced. As shown in the figure:

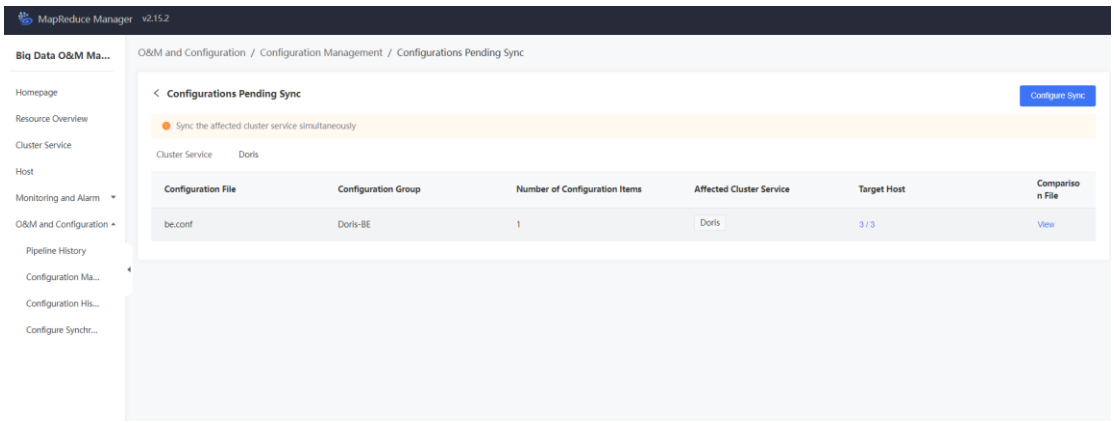

9. Click the **Configuration Sync** button to go to the configuration sync menu page. As shown in the figure:

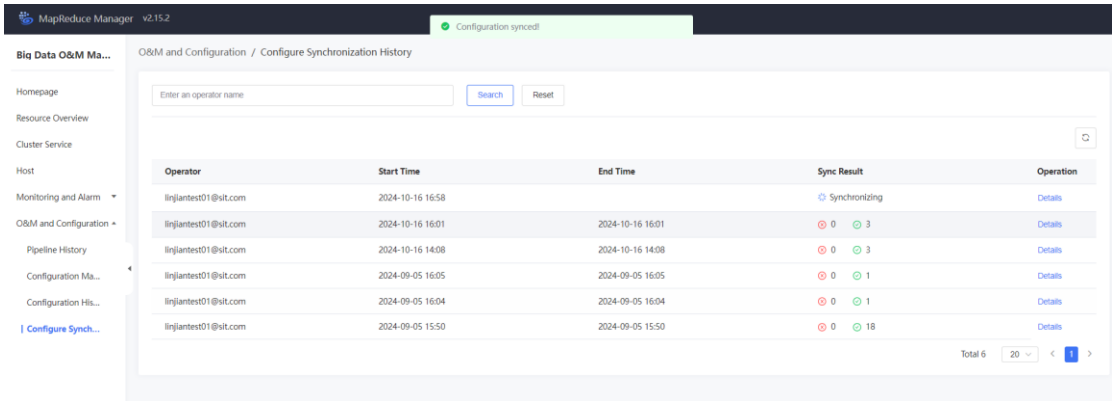

10. Click **Synchronizing** to go to the configuration sync details page, view the synchronization progress, and wait for the sync to complete. As shown in the figure:

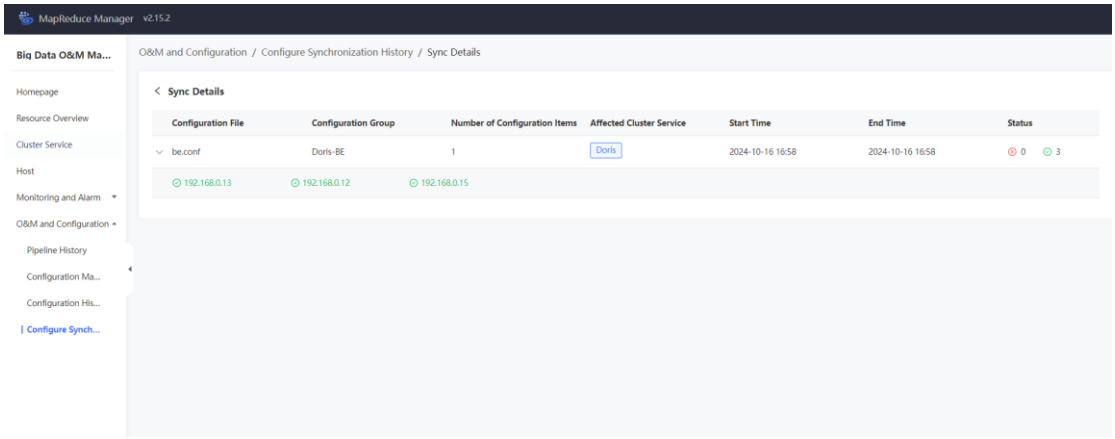

View Configuration Status

Scenario

This feature enables users to view the synchronization status of all configuration files for a chosen cluster service.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Click **More** > **View Configuration Status**. As shown in the figure:

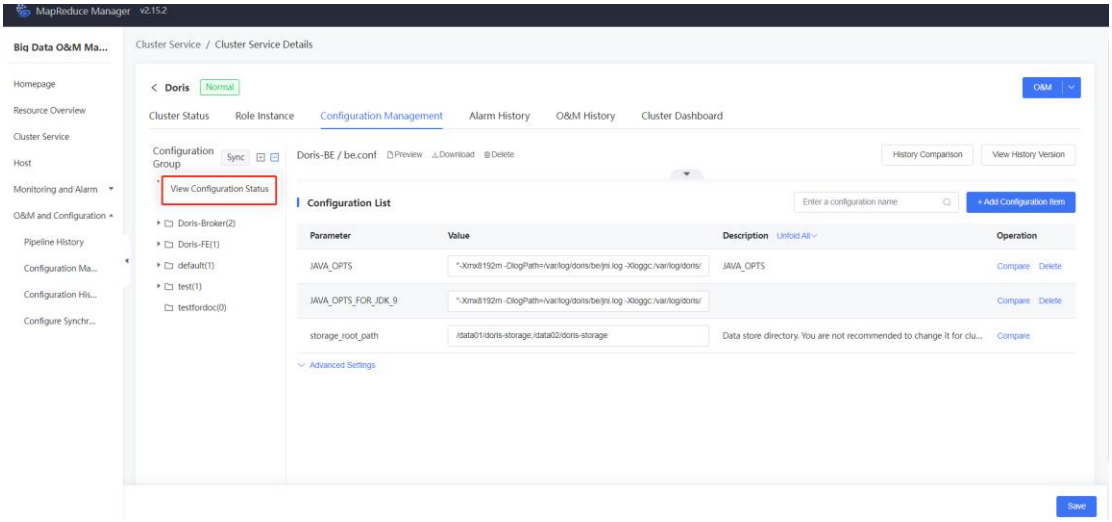

8. A dialog box displaying the synchronization status of all configuration files will appear. As shown in the figure:

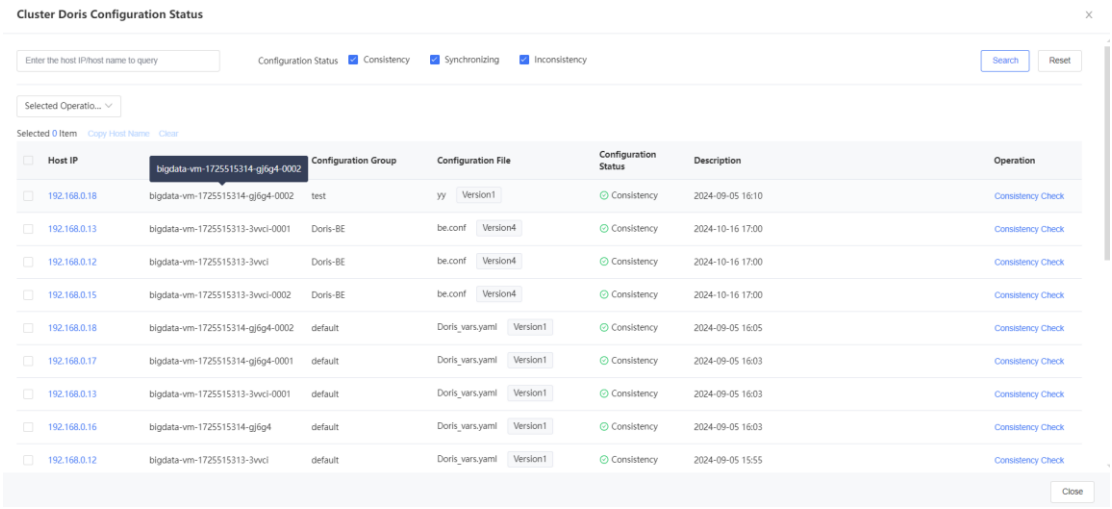

Resynchronize

Scenario

This feature supports resynchronization of configurations that have failed to sync or have inconsistent sync statuses.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.

3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Click **More > View Configuration Status**. A configuration synchronization status dialog box will appear.
8. Check the box in front of the host that requires operation.
9. Click **Operate on Selected > Resynchronize** to perform the configuration synchronization. As shown in the figure:

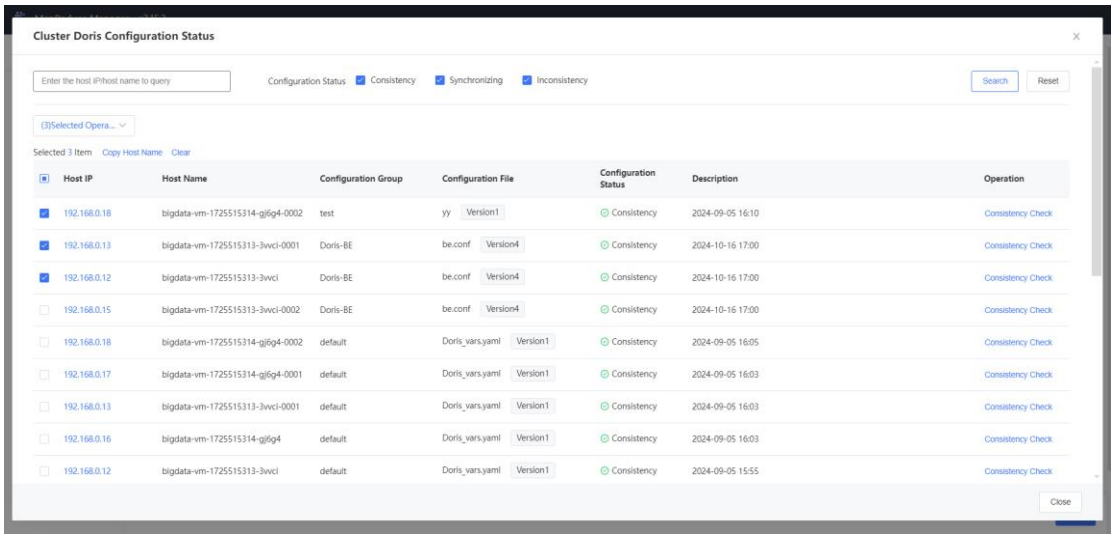

Configuration Rollback

Scenario

This feature supports the rollback of a specified configuration file to a specific historical version.

Procedure

1. Log in to the YI-MapReduce management console.

2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Configuration Management** tab.
7. Click on the configuration group, and then on the name of the configuration file. Details of this configuration will be displayed on the right side of the page.
8. Click on **More > History Comparison**, and a history comparison dialog box will appear. As shown in the figure:

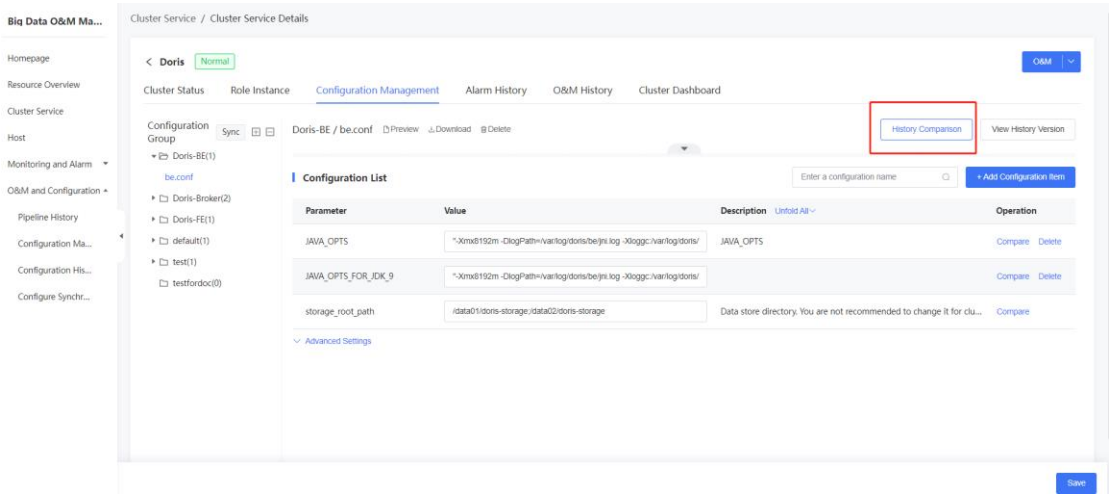

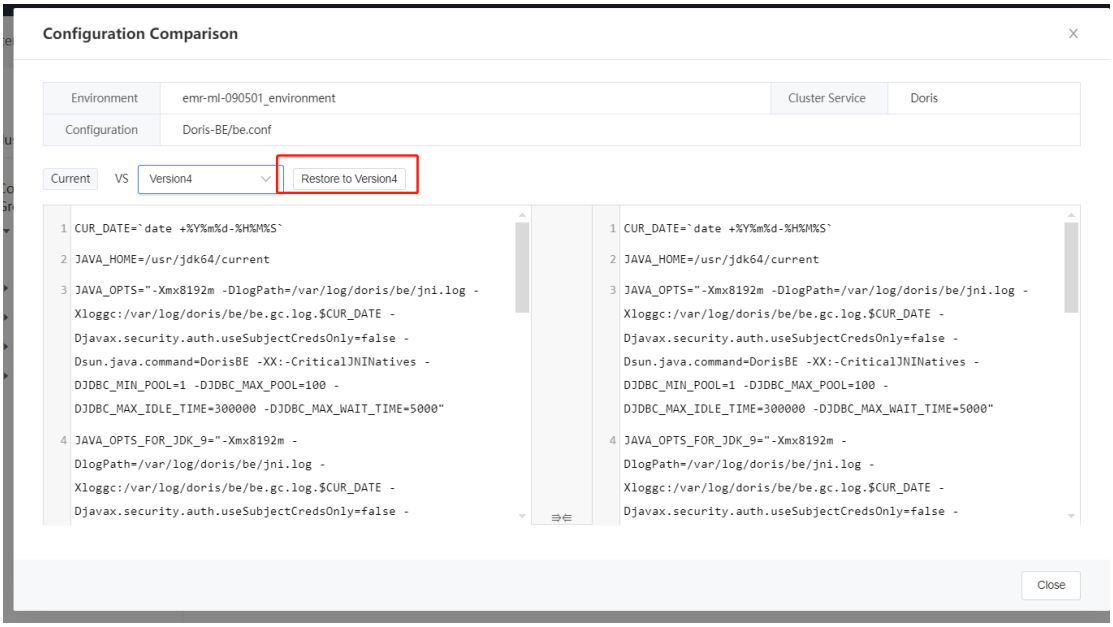

9. Click the **Restore to Version X** button. A confirmation dialog box will appear. As shown in the figure:

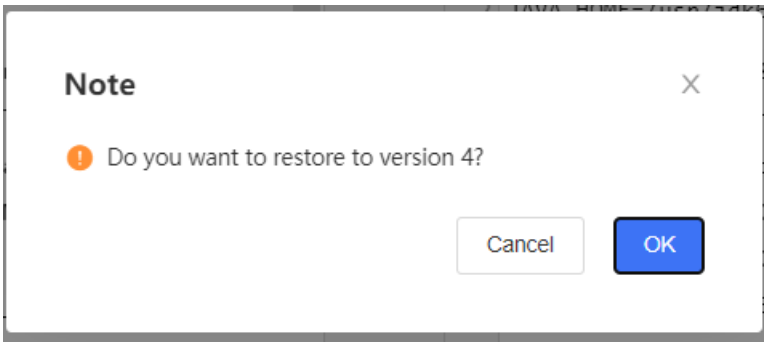

10. Click **OK**.
11. Click the **Close** button to close the configuration comparison dialog box.
12. Check the configuration file. The content of the configuration file should now match the restored version.
13. Click **Sync** to synchronize the configuration file.

View O&M History

Procedure

1. Log in to the YI-MapReduce management console.

2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **O&M History** tab to view all O&M history for this cluster service. As shown in the figure:

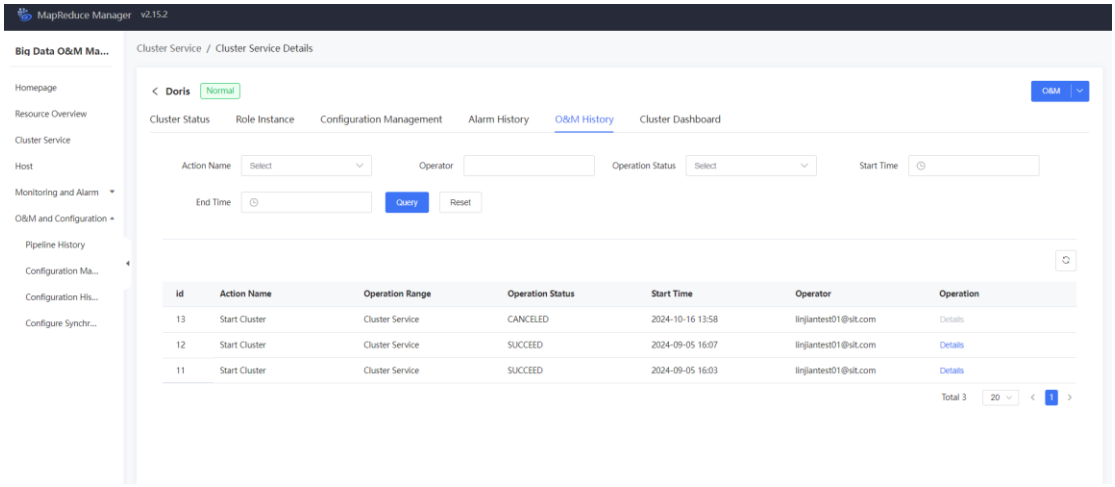

View Alarm History

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.

6. Click the **Alarm History** tab to view all alarm messages for this cluster service. As shown in the figure:

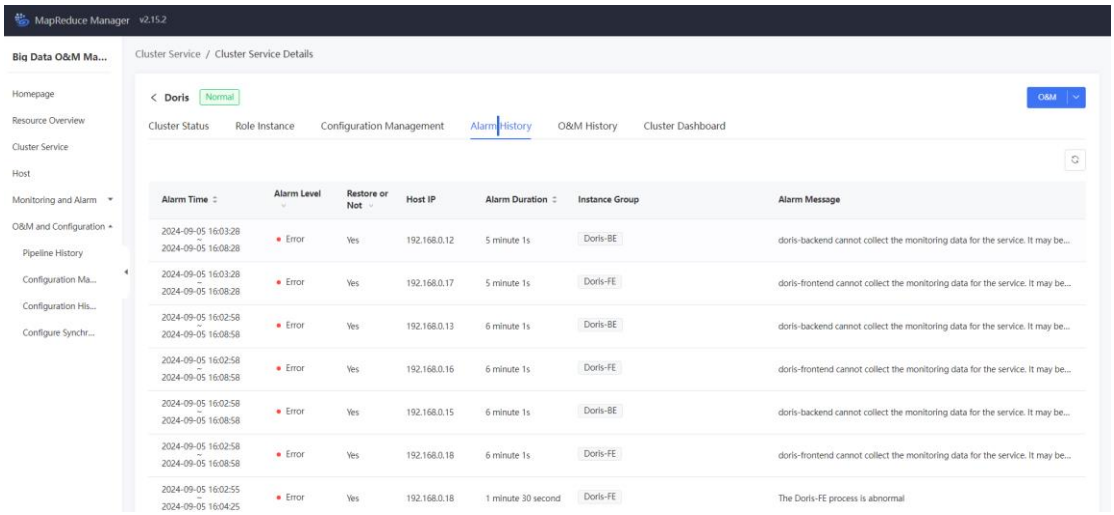

View Cluster Service Dashboard

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
5. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
6. Click the **Cluster Dashboard** tab to view the monitoring screen of the cluster service. As shown in the figure:

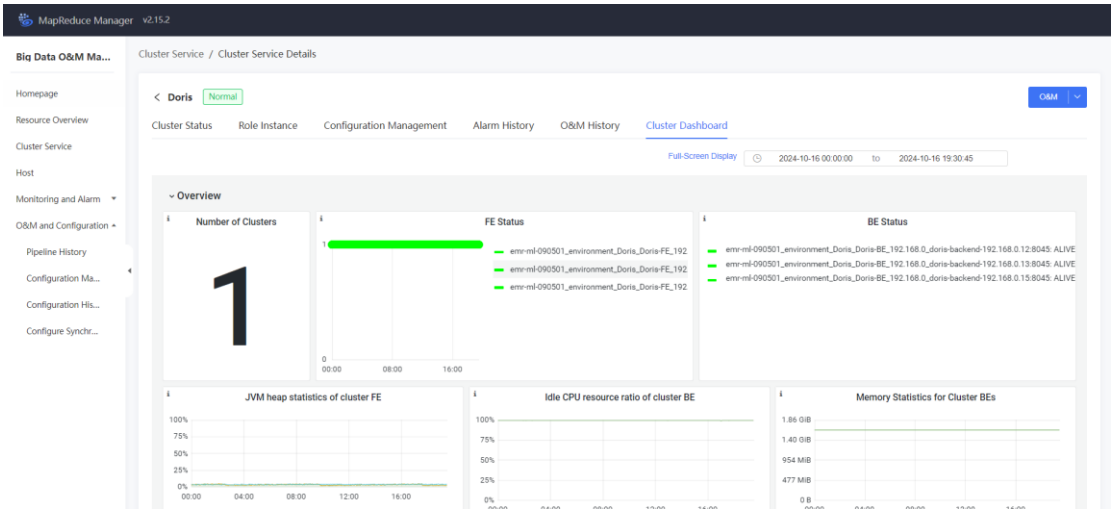

4.5.5. Host

Host Management Page Overview

Once in the YI-MapReduce Manager, click **Host** on the menu to go to the host page.

As shown in the figure:

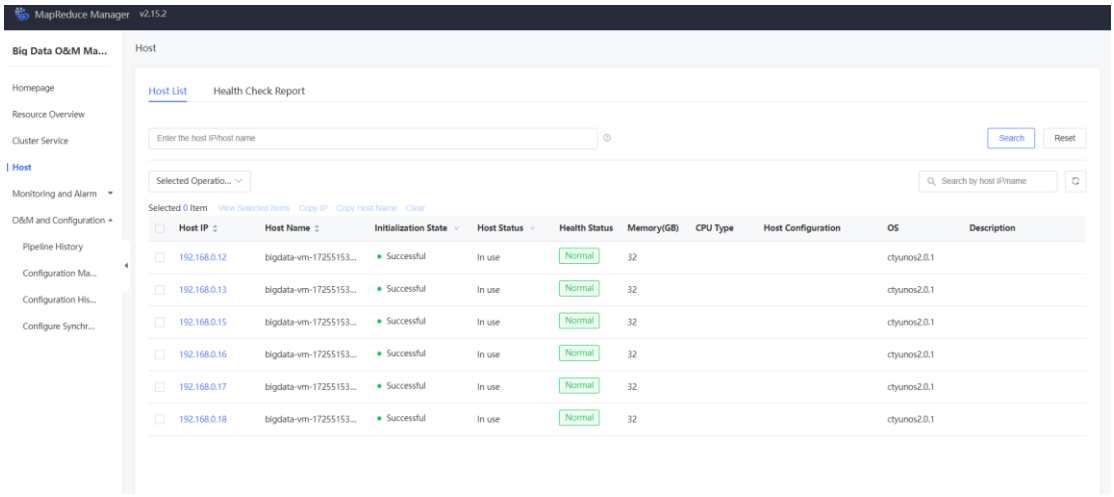

- The query function at the top of the host list allows querying by host IP/host name.
- The host list is presented in a data list format.
- The operation area located at the top left corner of the host list allows operations such as setting maintenance, canceling maintenance, and exporting data.

- Click on the host IP or the host name to go to the host details page.

Host Details Page Overview

Once in the YI-MapReduce Manager, click **Host** on the menu to go to the host list page.

Click on the host IP or the host name to go to the host details page. As shown in the figure:

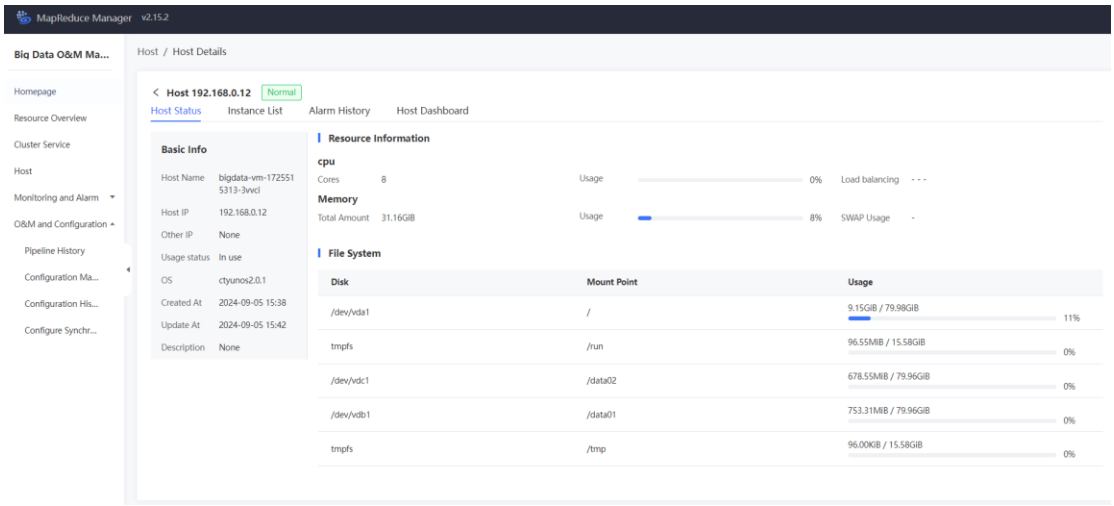

This page displays the following information in sequence: host status, instance list, and the alarm history for the host.

Host Maintenance and Cancellation

Set Maintenance

Scenario

Users may need to set a specific host or a batch of hosts into maintenance mode.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.

3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **Host** on the menu to go to the host list page.
5. Check the box in front of the host that requires operation.
6. Click on **Operate on Selected > Set Maintenance**. A set maintenance pop-up will appear. As shown in the figure:

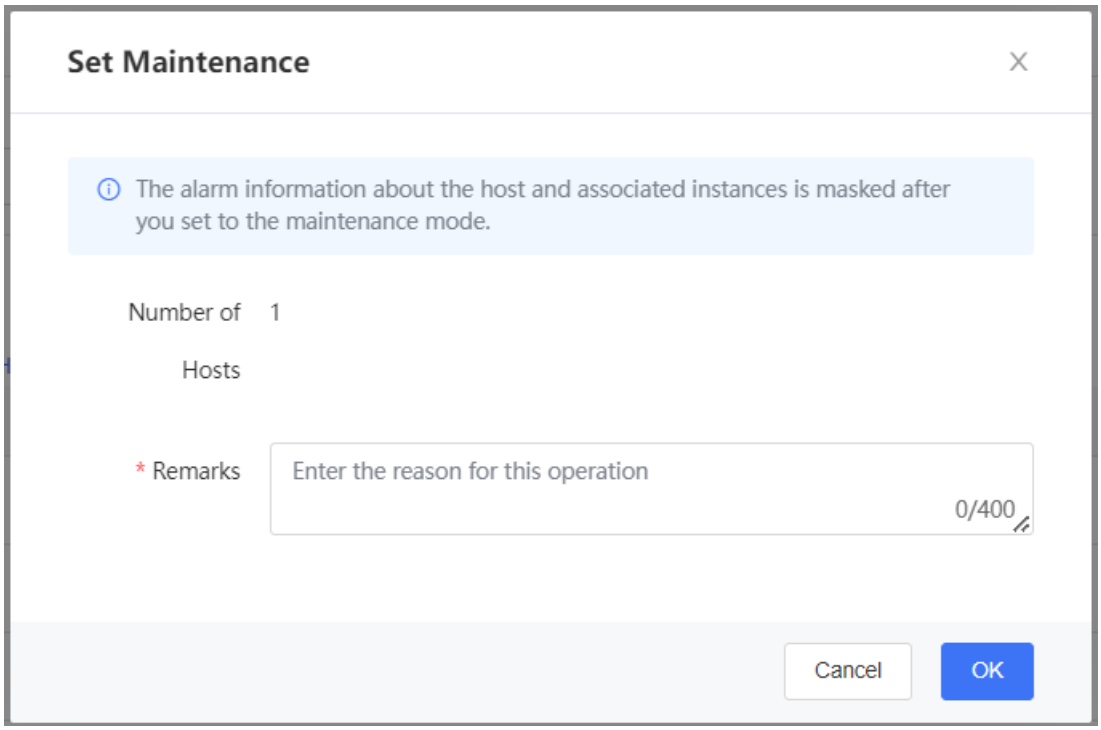

7. Enter notes and click

   OK.

Cautions

1. After setting the host to maintenance, the health status of this host will be displayed as in maintenance. As shown in the figure:

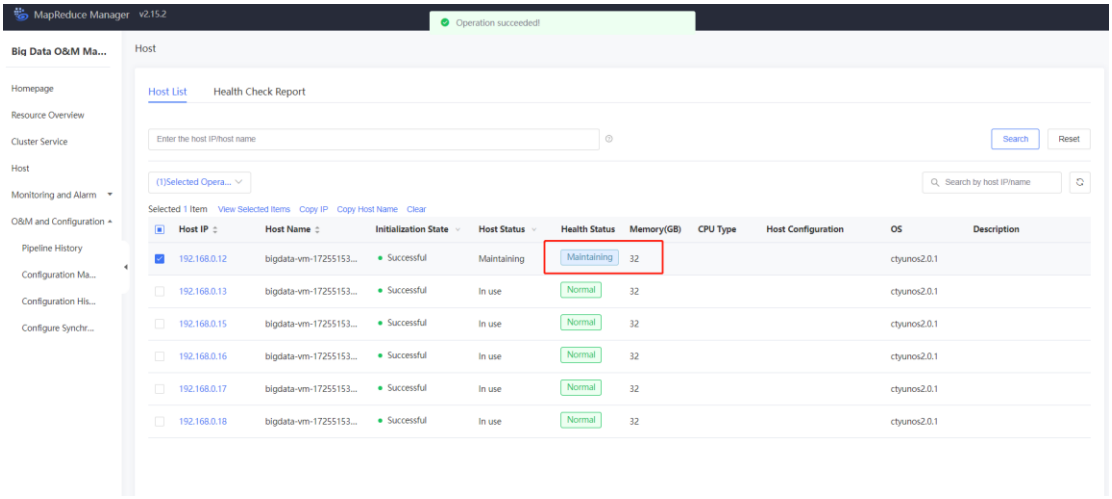

2. After setting the host to maintenance, the alarm messages of this host will be blocked.

Cancel Maintenance

Scenario

Users may need to cancel the maintenance status for a specific host or a batch of hosts.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **Host** on the menu to go to the host list page.
5. Check the box in front of the host that requires operation.
6. Click **Operate on Selected > Cancel Maintenance**. The cancel maintenance operation pop-up will appear.
7. Enter notes and click

OK.

View Host Status

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **Host** on the menu to go to the host list page.
5. Click on the IP or host name of the specific host to go to the host details page, which defaults to the host status page. As shown in the figure:

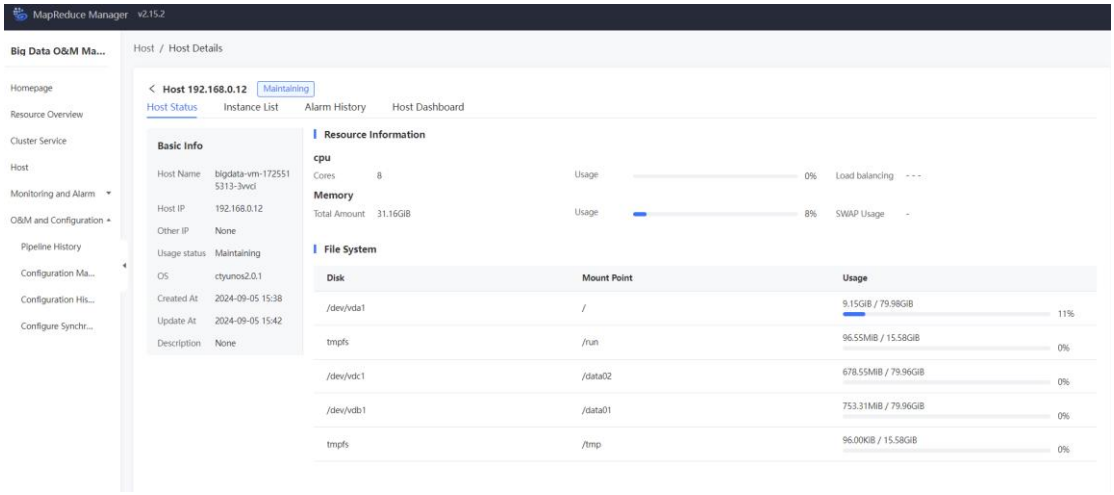

View Host Instance List

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.

3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **Host** on the menu to go to the host list page.
5. Click on the IP or the name of the specific host to go to the host details page.
6. Click the **Instance List** tab to view the instance services installed on this host. As shown in the figure:

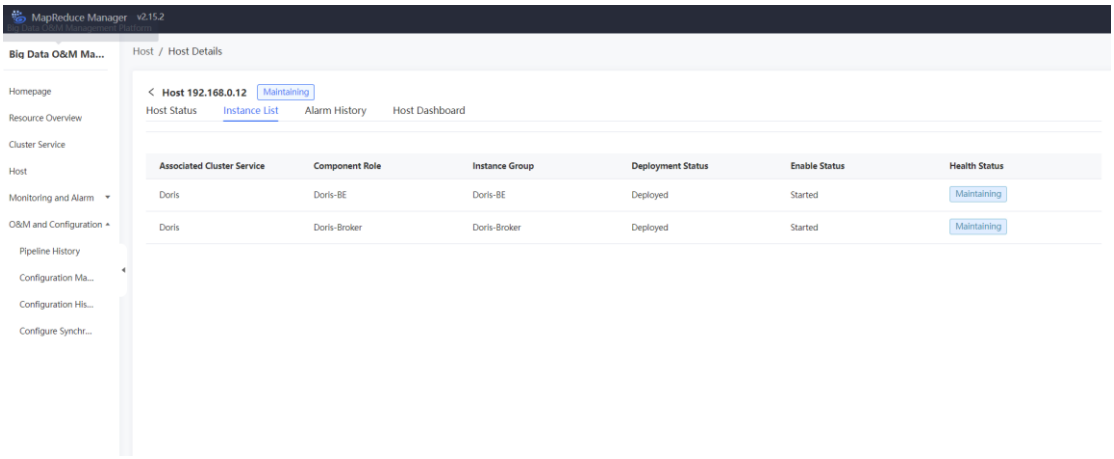

View Host Alarm History

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **Host** on the menu to go to the host list page.
5. Click on the IP or the name of the specific host to go to the host details page.
6. Click the **Alarm History** tab to view all alarm messages for this host. As shown in the figure:

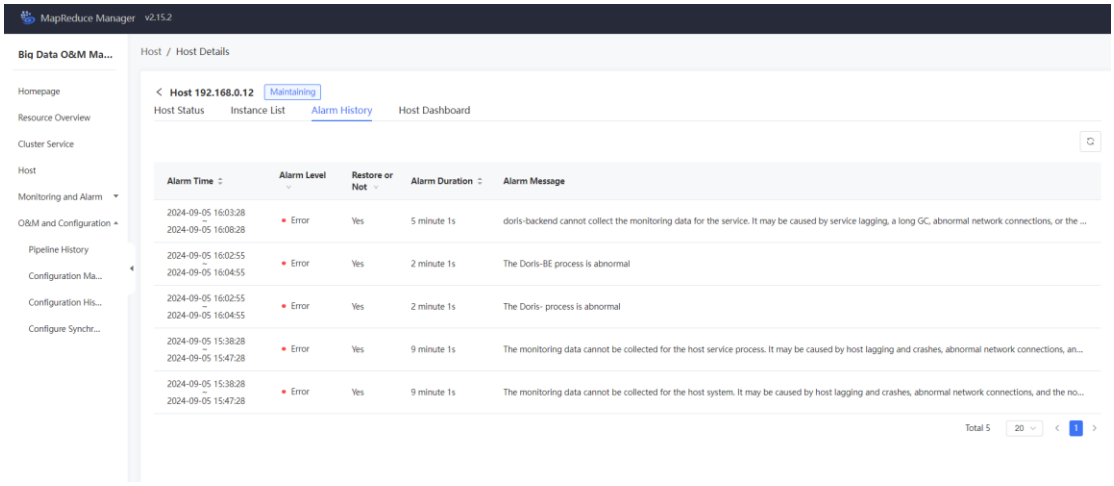

4.5.6. Monitoring and Alarm

Alarm History - Alarm History Overview

Once in the YI-MapReduce Manager, click **Monitoring and Alarm> Alarm History** on the menu to go to the alarm history page. As shown in the figure:

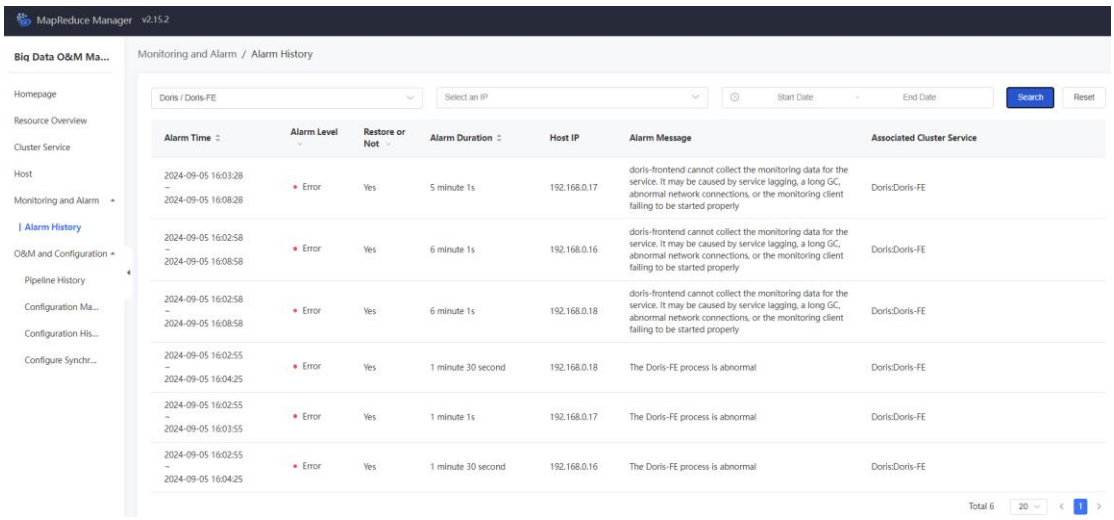

-

  The query function at the top of the page supports combined queries based on various criteria.

- Supports the querying of alarm messages at both the cluster service level, role instance level, and host level.

4.5.7. O&M and Configuration

Pipeline History - Pipeline History Overview

Once in the YI-MapReduce Manager, click **O&M > Pipeline History** on the menu to go to the pipeline history page. As shown in the figure:

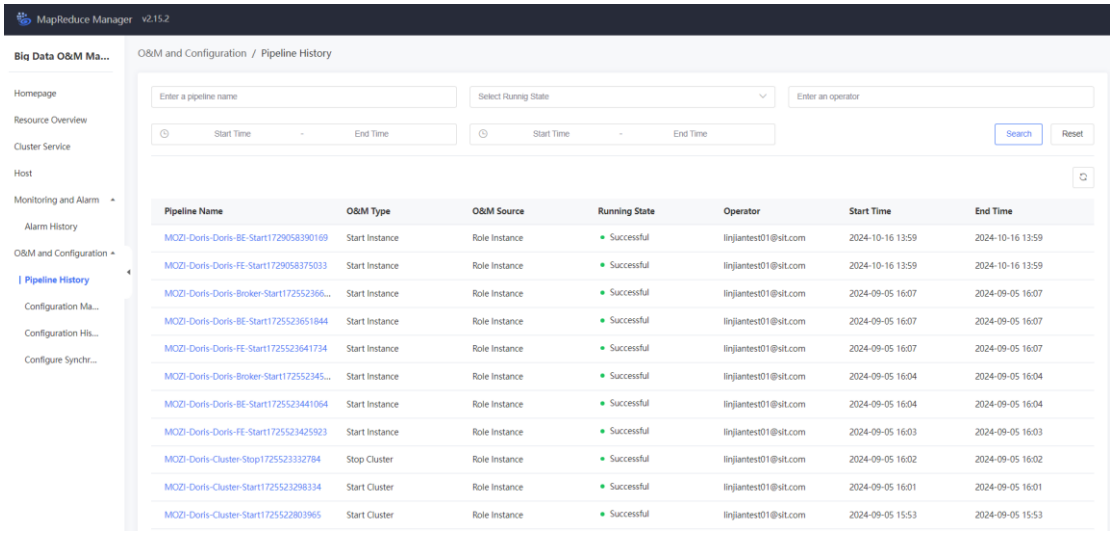

- The query function at the top of the page supports combined queries based on various criteria.
- The page displays the execution history of all pipelines, along with the operator.
- To view execution details, click on the pipeline name or the details button.

Pipeline History - View Pipeline Execution Details

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Pipeline History** on the menu.

5. To view the execution details of a specific pipeline run, click the pipeline name or the details button. As shown in the figure:

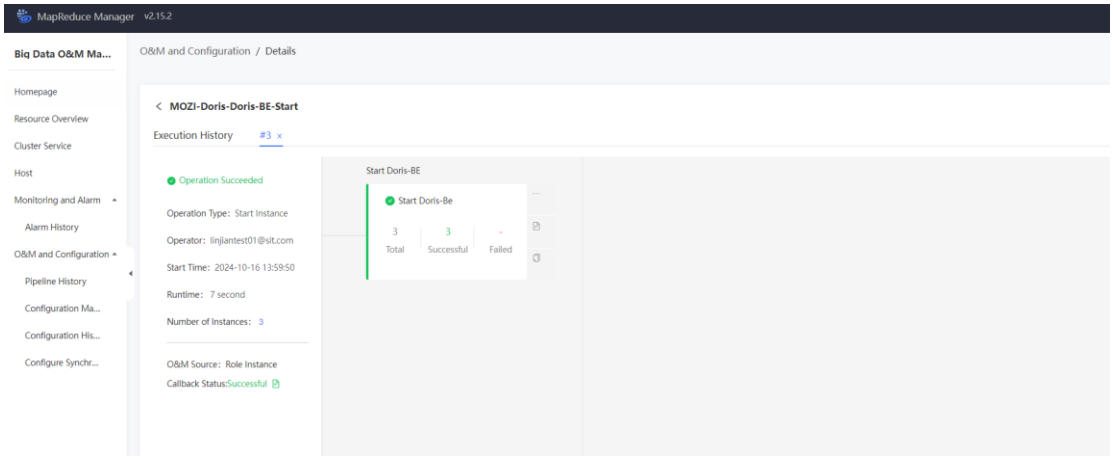

Configuration Management - Configuration Management Overview

Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Management** on the menu to go to the configuration management page. As shown in the figure:

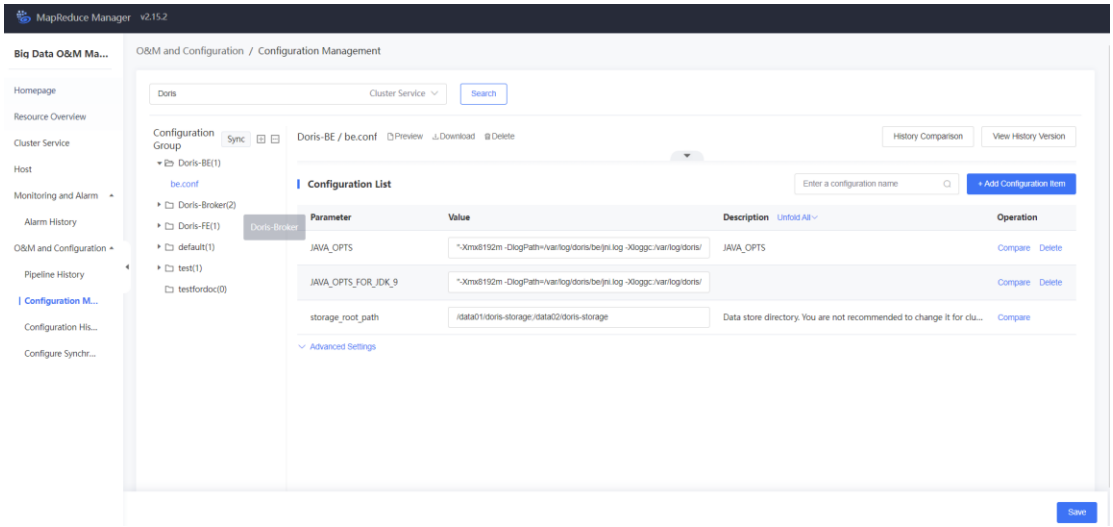

- Supports viewing configurations of different cluster services.
- Supports adding new configuration groups for different cluster services, viewing configuration group details and deleting configuration groups.

- Supports adding new configuration files, previewing and downloading configuration files, viewing configuration file details, as well as deleting configuration files.
- Supports configuration synchronization and provides a view of configuration status from the host perspective.

Configuration Management - Add Configuration Group

Scenario

Users can classify hosts within a cluster service based on different grouping criteria.

All hosts within a configuration group share a common set of configurations.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Management** on the menu.
5. Select the required cluster service for query conditions.
6. Click the **Configuration Management** tab.
7. Click on **Add Configuration Group** to go to the add configuration group page.
8. Choose the grouping method as needed, select the host, and click the **Save** button at the bottom right corner.

As shown in the figure:

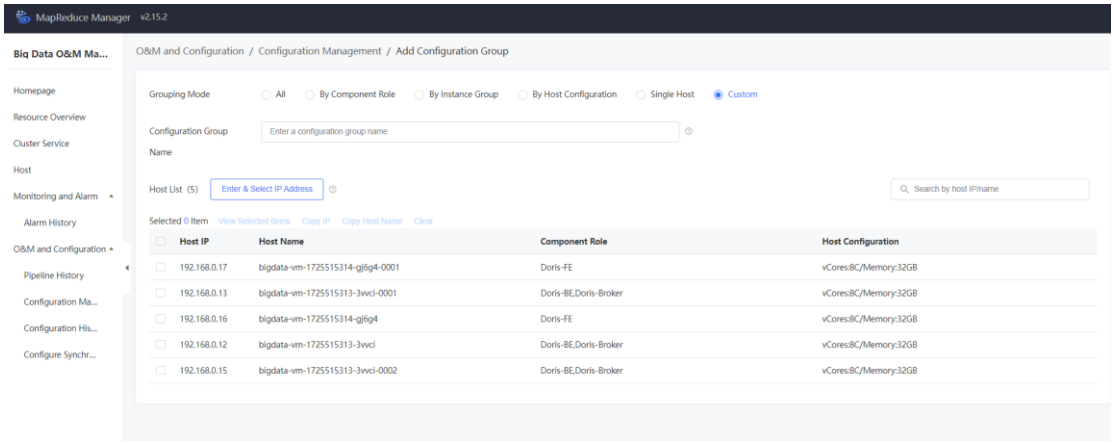

✧ **Note:**

- When the grouping method is "All", "By Component Role", "By Instance Group", or "By Host Configuration", the configuration group name is predefined and cannot be changed.

- When the grouping method is "Single Host", you need to select the host to be grouped. The configuration group name is the name of the selected host and cannot be changed.
- When the grouping method is "Custom", you need to select the host to be grouped. The configuration group name can be customized but must not be the same as other configuration group names.

Configuration Management - View Configuration Group Details

Scenario

Users can view details of different configuration groups to learn about the host information contained within them.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.

3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.

4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Management** on the menu.

5. Select the required cluster service for query conditions.

6. Choose the configuration group you want to view. Hover over the configuration group name, and click on the **Details** icon on the right to access the configuration group details page. As shown in the figure:

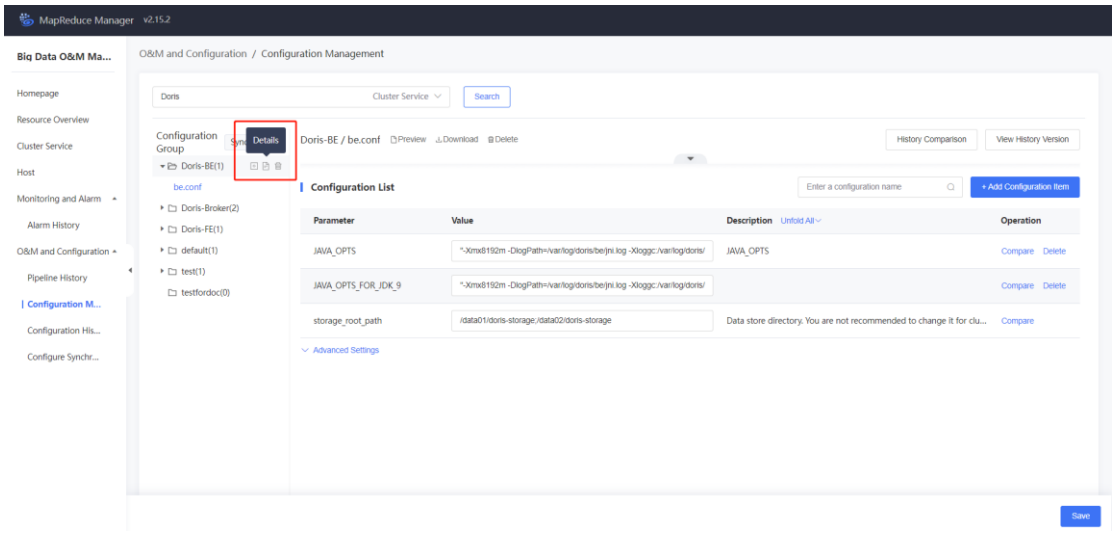

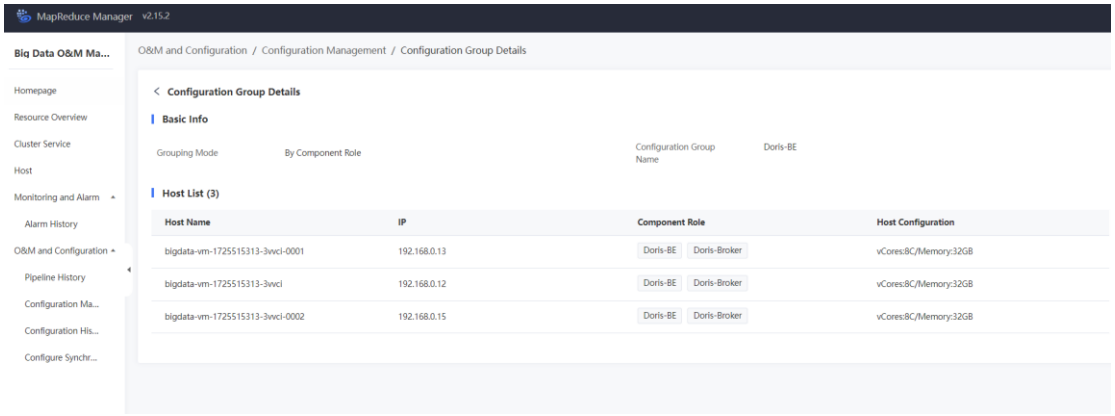

Configuration Management - Delete Configuration Group

Scenario

Users can remove configuration groups no longer needed.

Before You Begin

The number of configuration files in the configuration group must be 0 to proceed with the deletion.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Management** on the menu.
5. Select the required cluster service for query conditions.
6. Choose the configuration group you want to delete. Hover over the configuration group name, and click on the **Delete** icon on the right. A deletion confirmation pop-up will appear.
7. Click **OK** to delete the configuration group.

Configuration Management - Add Configuration

Scenario

Users can classify hosts within a cluster service based on different grouping criteria.

All hosts within a configuration group share a common set of configurations.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.

3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Management** on the menu.
5. Select the required cluster service for query conditions.
6. Click **Add** to go to the add configuration group page.
7. Choose the grouping method as needed, select the host, and click the **Save** button at the bottom right corner.

As shown in the figure:

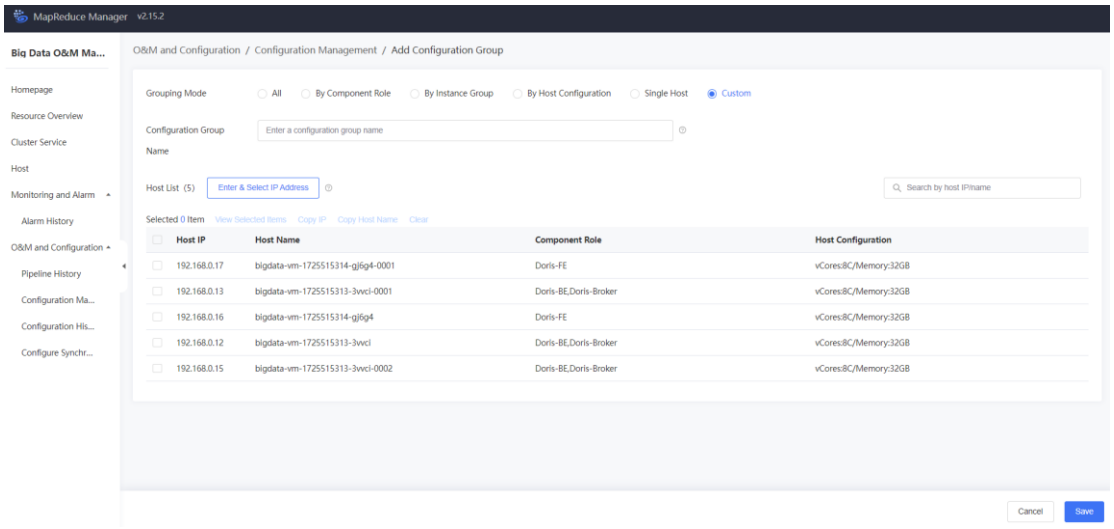

✧ Note:

- When the grouping method is "All", "By Component Role", "By Instance Group", or "By Host Configuration", the configuration group name is predefined and cannot be changed.

- When the grouping method is "Single Host", you need to select the host to be grouped. The configuration group name is the name of the selected host and cannot be changed.
- When the grouping method is "Custom", you need to select the host to be grouped. The configuration group name can be customized but must not be the same as other configuration group names.

Configuration Management - View Configuration Details

Scenario

Users can view details of different configuration groups to learn about the host information contained within them.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Management** on the menu.
5. Select the required cluster service for query conditions.
6. Choose the configuration group you want to view. Hover over the configuration group name, and click on the **Details** icon on the right to access the configuration group details page. As shown in the figure:

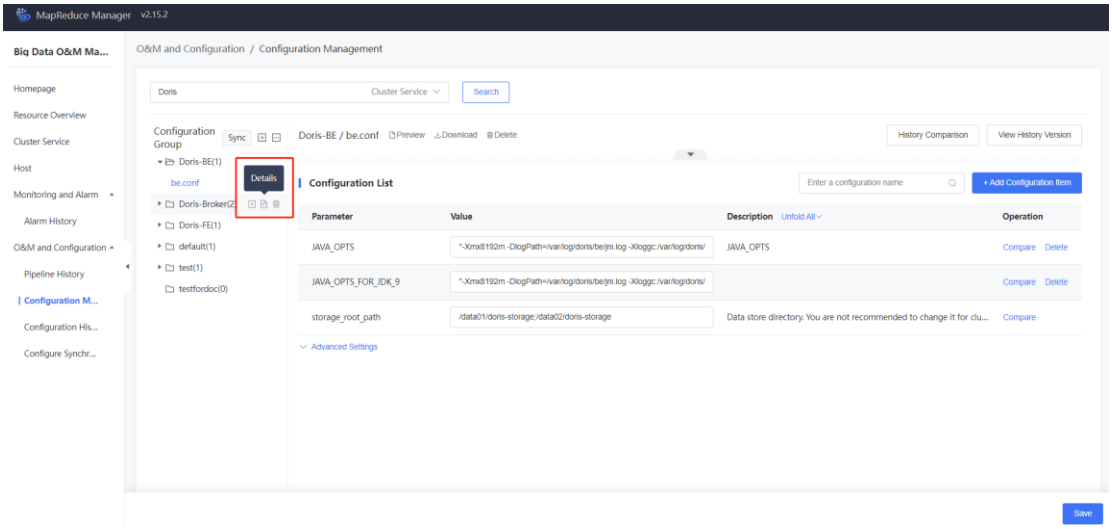

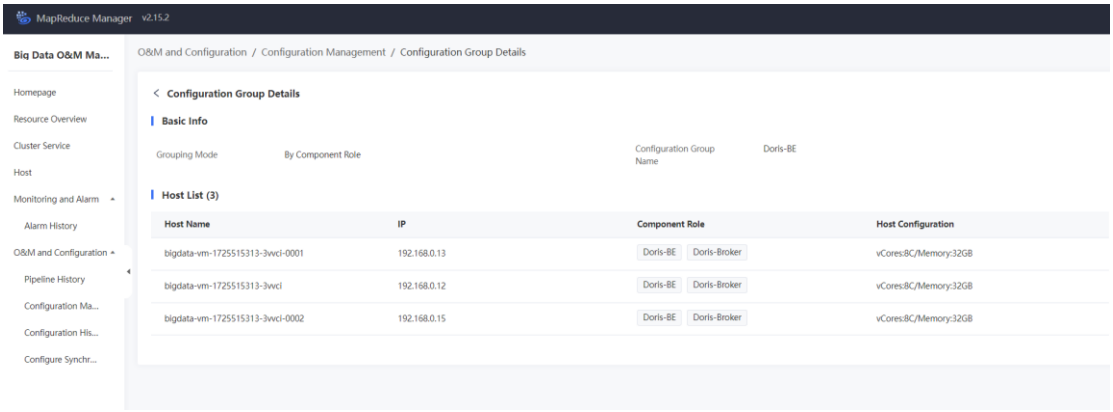

Configuration Management - Edit Configuration

Scenario

Users can modify configuration files.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Management** on the menu.
5. Select the required cluster service for query conditions.
6. Click on the configuration group, and then on the name of the configuration you want to modify. Details of this configuration will be displayed on the right side of the page. As shown in the figure:

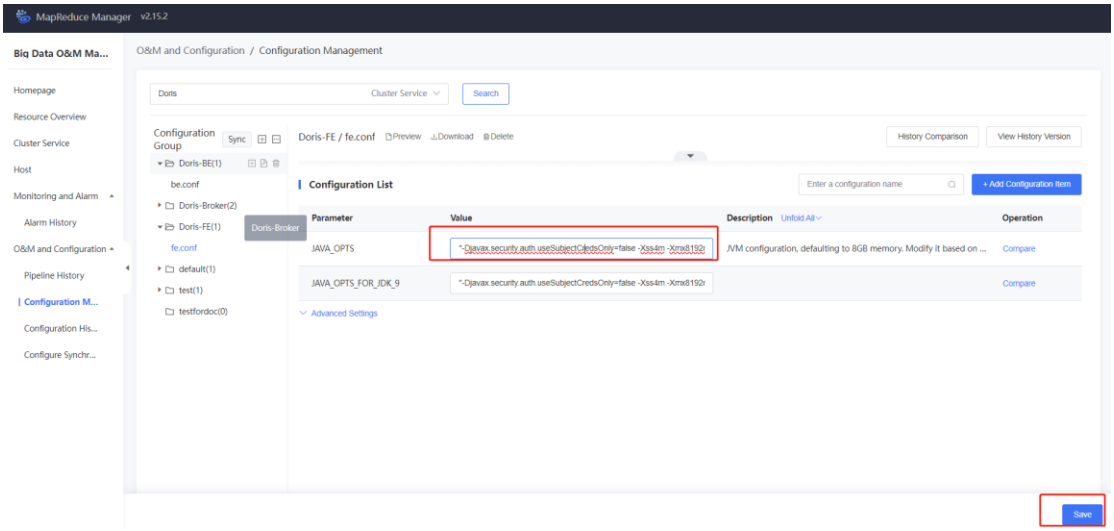

7. You can modify the configuration file directly. After completing the modifications, click **Save**.

Cautions

1. Display level: File level.

- Direct editing or modification is allowed in the text box.
- After modifications, if the current configuration file has not been synchronized, the number 1 will appear before the configuration name. As shown in the figure:

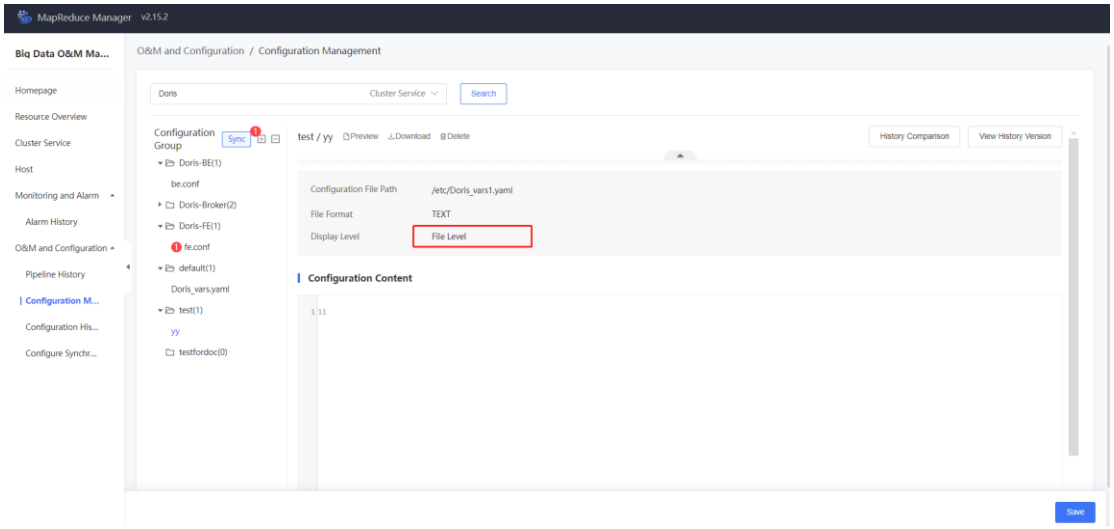

- Display level: Configuration item level.
- Configuration items need to be modified individually.
- After modifications, if the current configuration file has not been synchronized, the number of modified configuration items will be displayed before the

configuration name, and a modification mark will appear before the name of the modified configuration item. As shown in the figure:

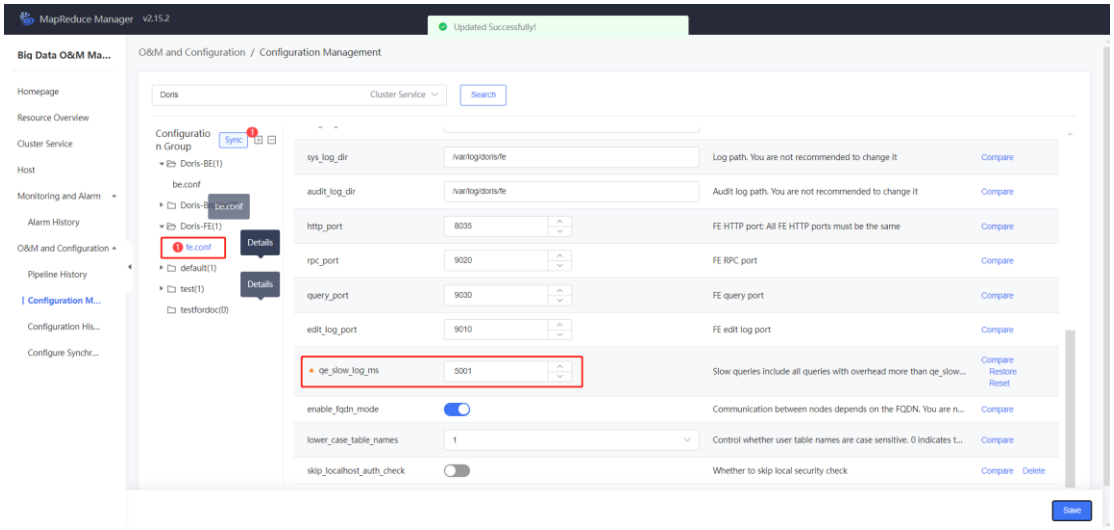

Configuration Management - Configuration File Preview

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Management** on the menu.
5. Select the required cluster service for query conditions.
6. Click on the configuration group, and then on the name of the configuration file. Details of this configuration will be displayed on the right side of the page.
7. Click **File Preview** and a configuration file preview dialog box will appear. As shown in the figure:

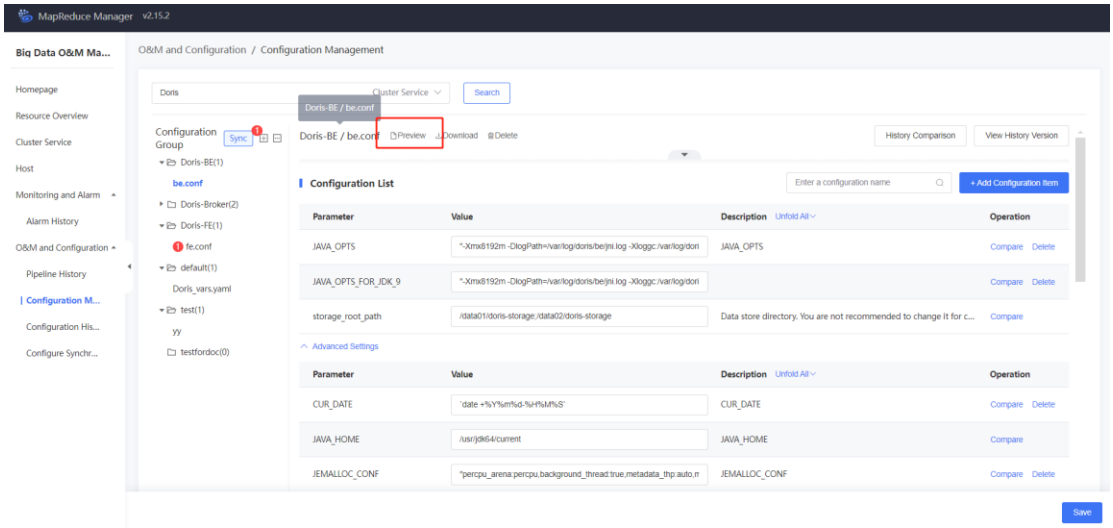
MapReduce Manager v2.15.2
Big Data O&M Ma...
O&M and Configuration / Configuration Management
Homepage
Resource Overview
Cluster Service
Host
Monitoring and Alarm
Alarm History
O&M and Configuration
Pipeline History
Configuration M...
Configuration His...
Configure Synchr...
Doris
Cluster Service
Search
Configuration Group
Sync
Doris-BE(1)
be.conf
Doris-Broker(2)
Doris-FE(1)
fe.conf
default(1)
Doris_vars.yaml
test(1)
yy
testfordoc(0)
Doris-BE / be.conf
Preview
Download
Delete
History Comparison
View History Version
Configuration List
Enter a configuration name
+ Add Configuration Item
Parameter
Value
Description
Unfold All
Operation
JAVA_OPTS
"-Xmx8192m -DlogPath=/var/log/doris/be/jni.log -Xloggc:/var/log/dori
JAVA_OPTS
Compare
Delete
JAVA_OPTS_FOR_JDK_9
"-Xmx8192m -DlogPath=/var/log/doris/be/jni.log -Xloggc:/var/log/dori
storage_root_path
/data01/doris-storage,/data02/doris-storage
Data store directory. You are not recommended to change it for c...
Advanced Settings
CUR_DATE
`date +%Y%m%d-%H%M%S`
JAVA_HOME
/usr/jdk64/current
JEMALLOC_CONF
"percpu_arena:percpu,background_thread:true,metadata_thp:auto,r
Save

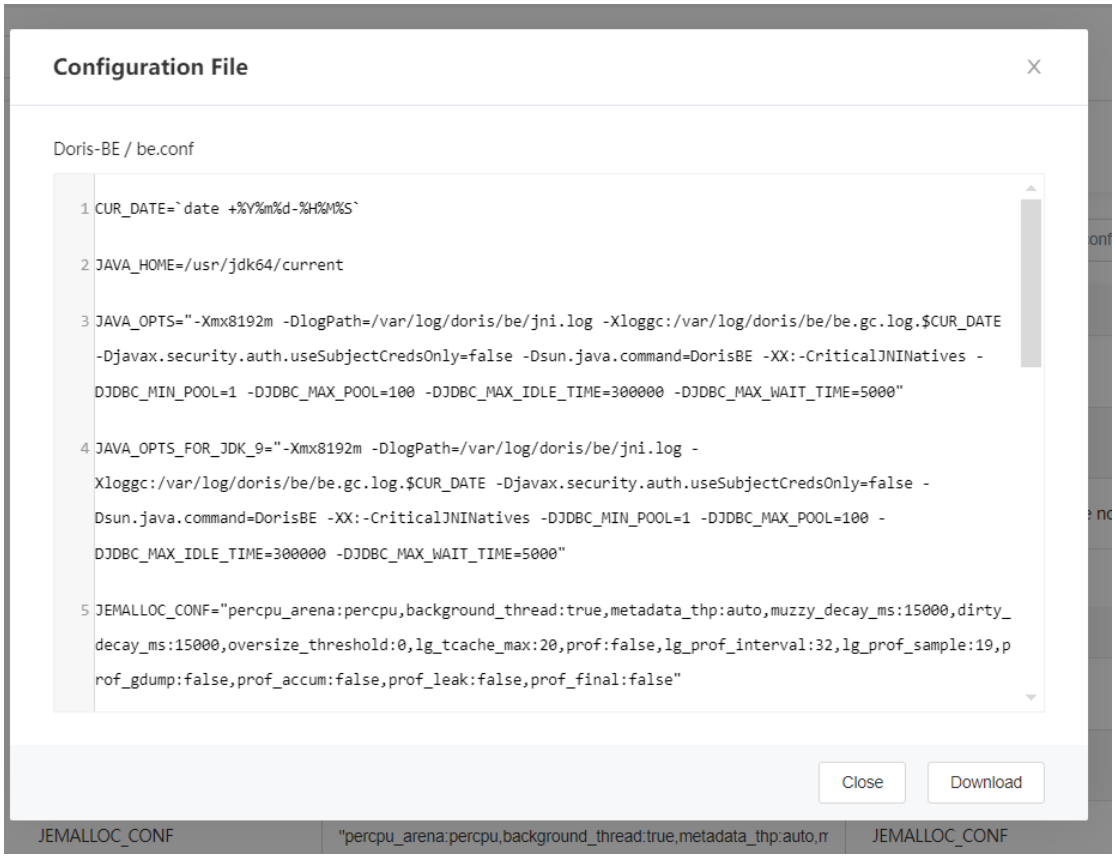
Configuration File
Doris-BE / be.conf
1 CUR_DATE=`date +%Y%m%d-%H%M%S`
2 JAVA_HOME=/usr/jdk64/current
3 JAVA_OPTS="-Xmx8192m -DlogPath=/var/log/doris/be/jni.log -Xloggc:/var/log/doris/be/be.gc.log.$CUR_DATE -Djavax.security.auth.useSubjectCredsOnly=false -Dsun.java.command=DorisBE -XX:-CriticalJNINatives -DJDBC_MIN_POOL=1 -DJDBC_MAX_POOL=100 -DJDBC_MAX_IDLE_TIME=300000 -DJDBC_MAX_WAIT_TIME=5000"
4 JAVA_OPTS_FOR_JDK_9="-Xmx8192m -DlogPath=/var/log/doris/be/jni.log -Xloggc:/var/log/doris/be/be.gc.log.$CUR_DATE -Djavax.security.auth.useSubjectCredsOnly=false -Dsun.java.command=DorisBE -XX:-CriticalJNINatives -DJDBC_MIN_POOL=1 -DJDBC_MAX_POOL=100 -DJDBC_MAX_IDLE_TIME=300000 -DJDBC_MAX_WAIT_TIME=5000"
5 JEMALLOC_CONF="percpu_arena:percpu,background_thread:true,metadata_thp:auto,muzzy_decay_ms:15000,dirty_decay_ms:15000,oversize_threshold:0,lg_tcache_max:20,prof:false,lg_prof_interval:32,lg_prof_sample:19,prof_gdump:false,prof_accum:false,prof_leak:false,prof_final:false"
Close
Download
JEMALLOC_CONF
"percpu_arena:percpu,background_thread:true,metadata_thp:auto,r
JEMALLOC_CONF

Configuration Management - Download Configuration Files

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Management** on the menu.
5. Select the required cluster service for query conditions.
6. Click on the configuration group, and then on the name of the configuration file. Details of this configuration will be displayed on the right side of the page.
7. Click **Download Files**, and wait for the download to complete. As shown in the figure:

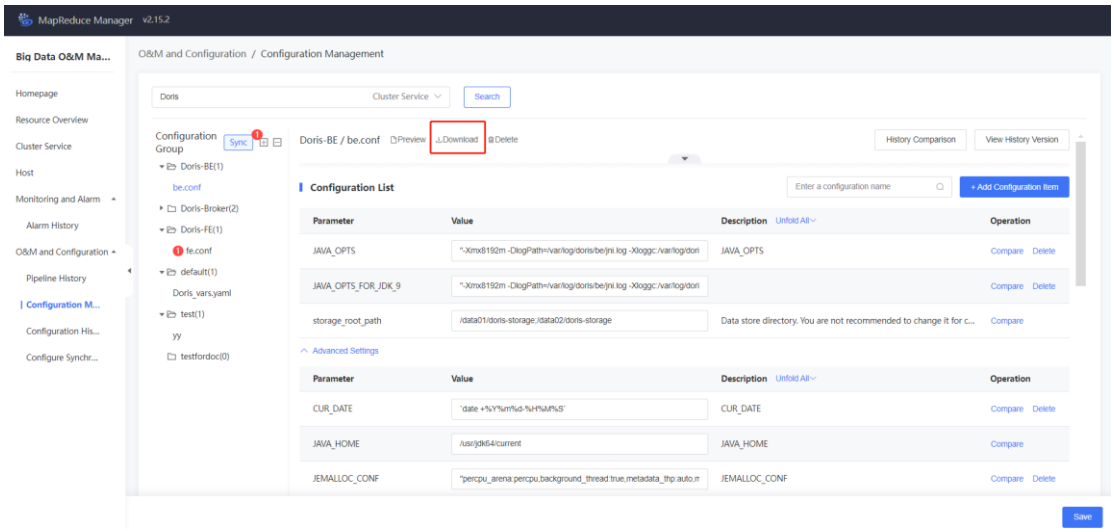

Configuration Management - Configuration File History Comparison

Scenario

This feature allows users to compare the current configuration file with historical versions.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Management** on the menu.
5. Select the required cluster service for query conditions.
6. Click on the configuration group, and then on the name of the configuration file. Details of this configuration will be displayed on the right side of the page.
7. Click on **More > History Comparison**, and a history comparison dialog box will appear. As shown in the figure:

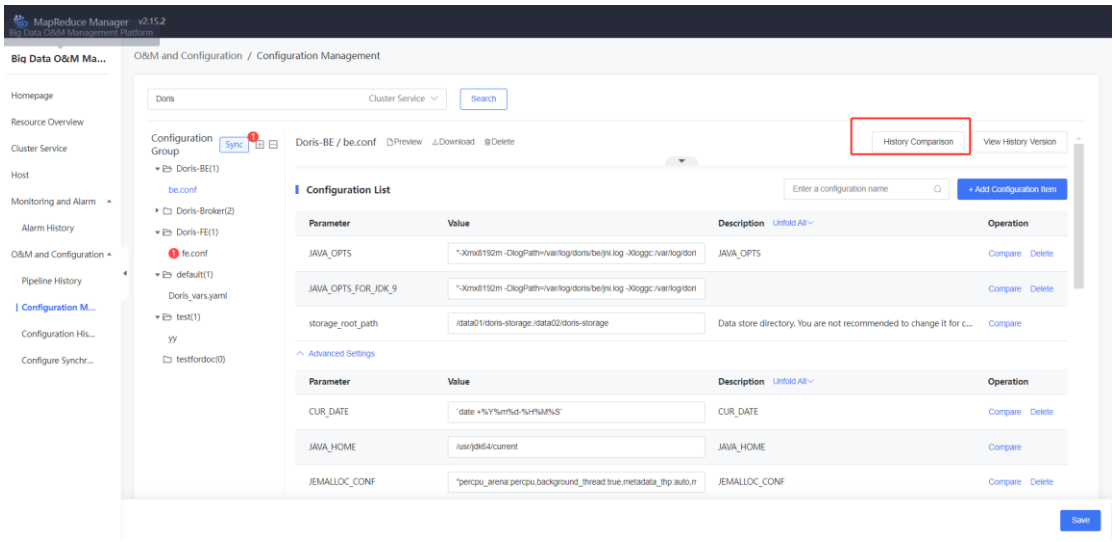

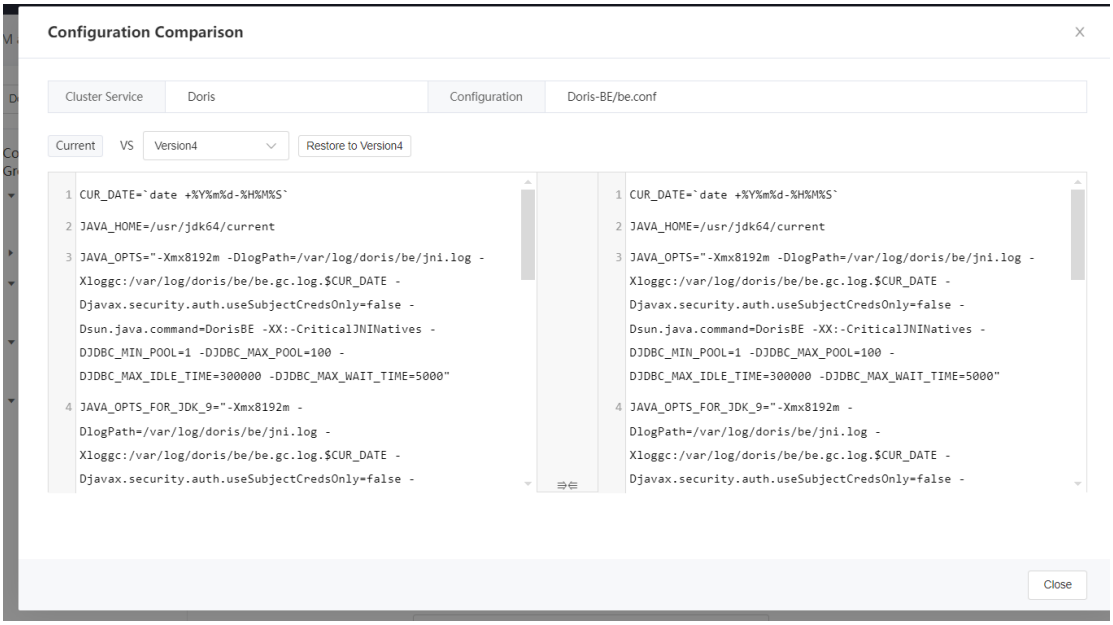

Configuration Management - View Configuration History Versions

Scenario

This feature allows users to view the history versions of a specific configuration file, along with their contents and changes between different versions.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Management** on the menu.
5. Select the required cluster service for query conditions.
6. Click on the configuration group, and then on the name of the configuration file. Details of this configuration will be displayed on the right side of the page.

7. Click **More > View History Versions**. As shown in the figure:

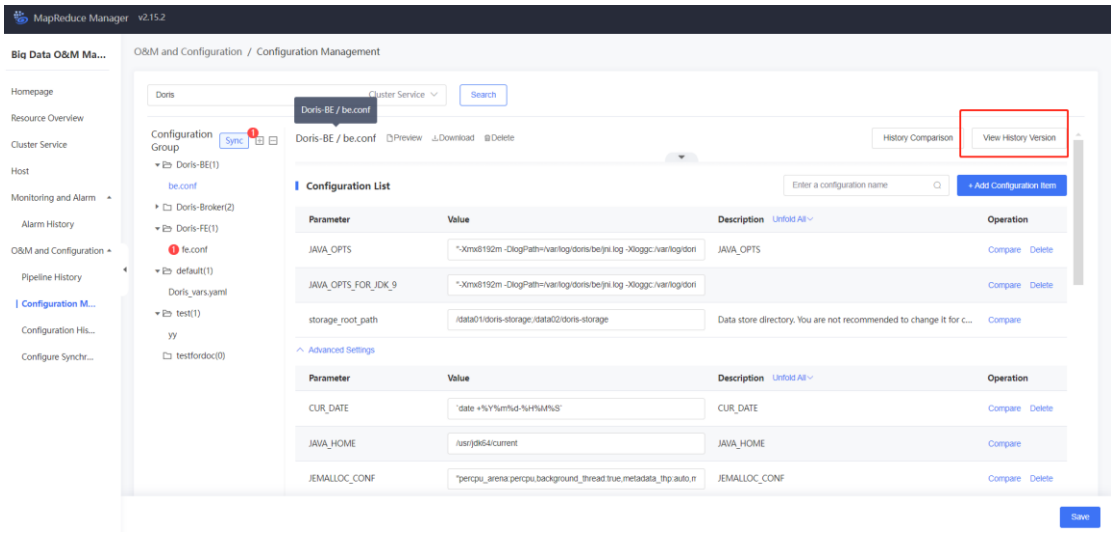

8. You will enter the configuration history page, which displays all history versions of the current configuration. As shown in the figure:

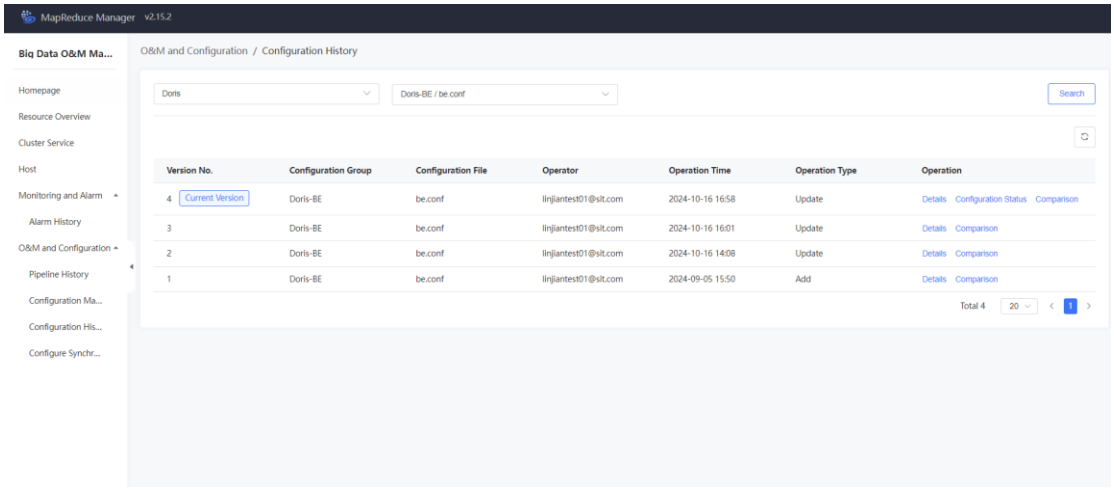

Configuration Management - Delete Configuration

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.

4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Management** on the menu.
5. Select the required cluster service for query conditions.
6. Click on the configuration group, and then on the name of the configuration file. Details of this configuration will be displayed on the right side of the page.
7. Click on **Delete Configuration**. As shown in the figure:

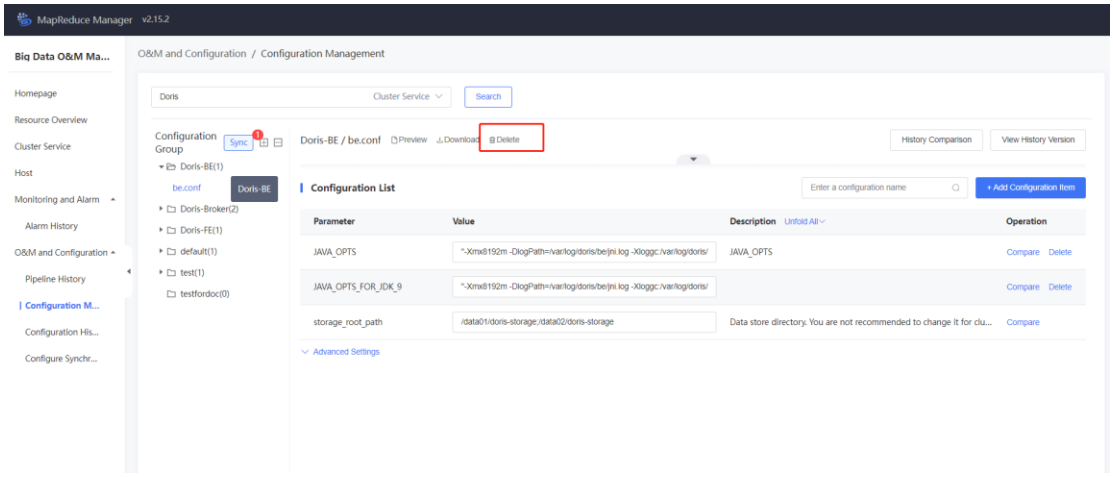

8. A confirmation pop-up for deletion will appear. Click **OK**. As shown in the figure:

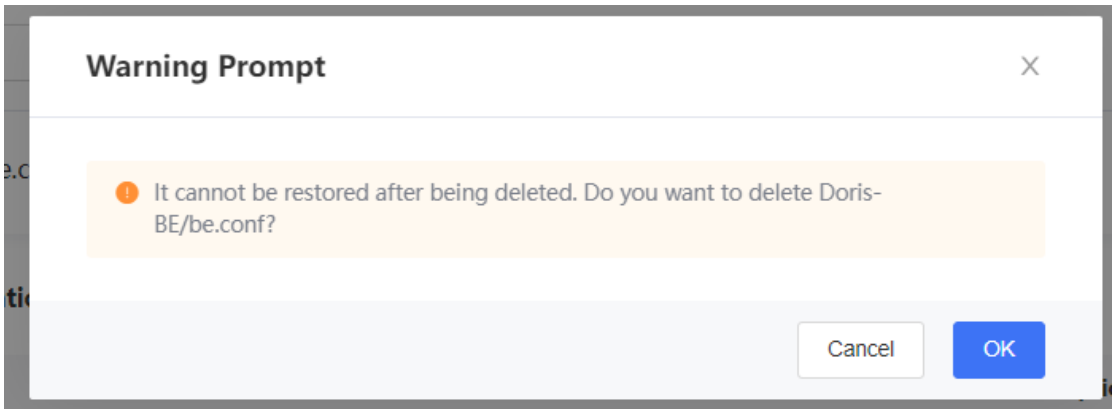

Configuration Management - Configuration Sync

Scenario

This feature supports the synchronization of new or modified configuration files to the respective hosts.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Management** on the menu.
5. Select the required cluster service for query conditions.
6. Click the **Sync** button. As shown in the figure:

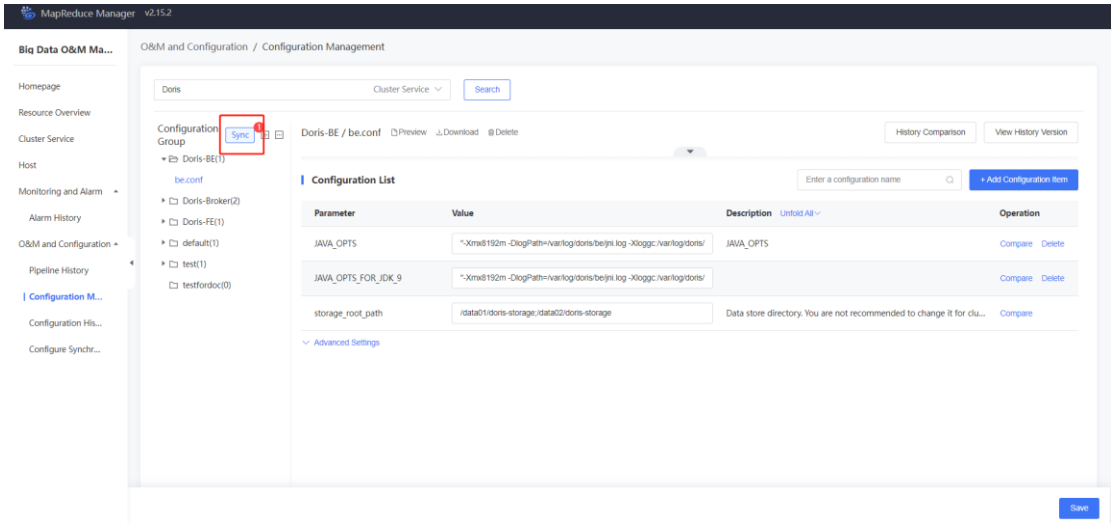

7. Enter the page of the configuration to be synced. As shown in the figure:

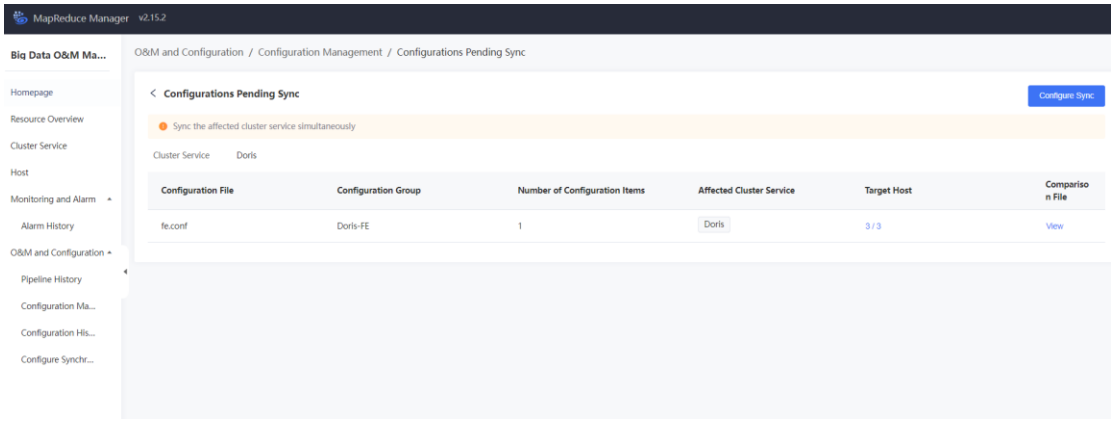

8. Click the **Configure Sync** button to go to the configuration sync menu page. As shown in the figure:

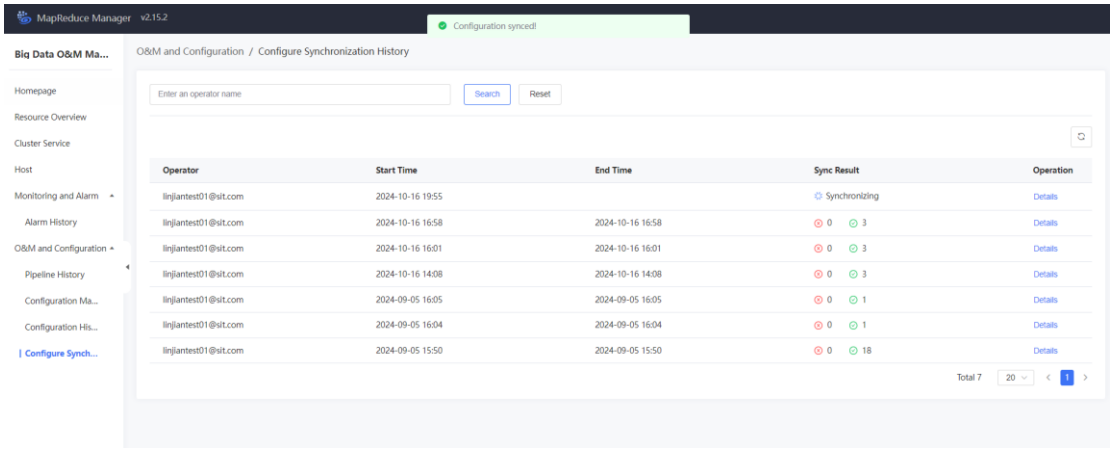

9. Click **Synchronizing** to go to the configuration sync details page, view the synchronization progress, and wait for the sync to complete. As shown in the figure:

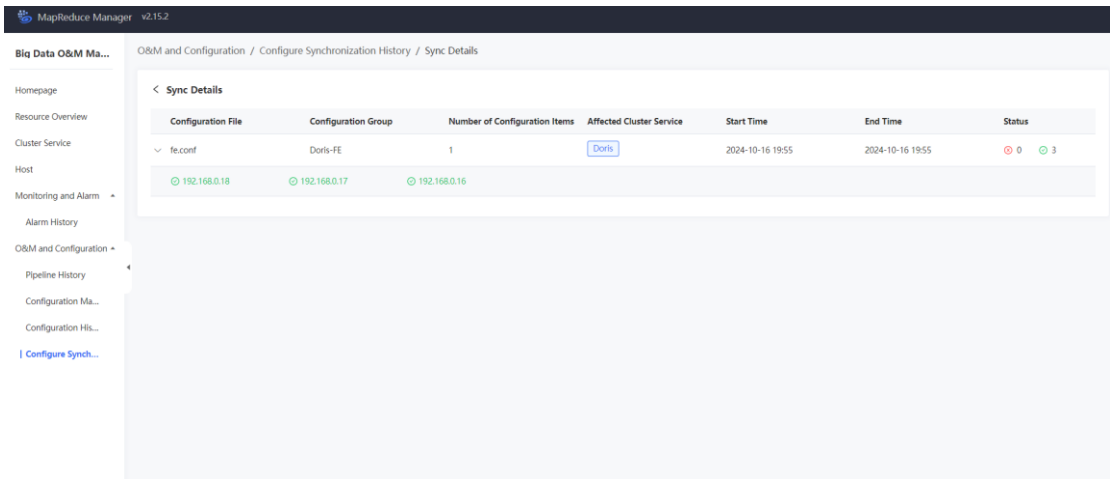

Configuration Management - View Configuration Status

Scenario

This feature enables users to view the synchronization status of all configuration files for a chosen cluster service.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Management** on the menu.
5. Select the required cluster service for query conditions.
6. Click **More > View Configuration Status**. As shown in the figure:

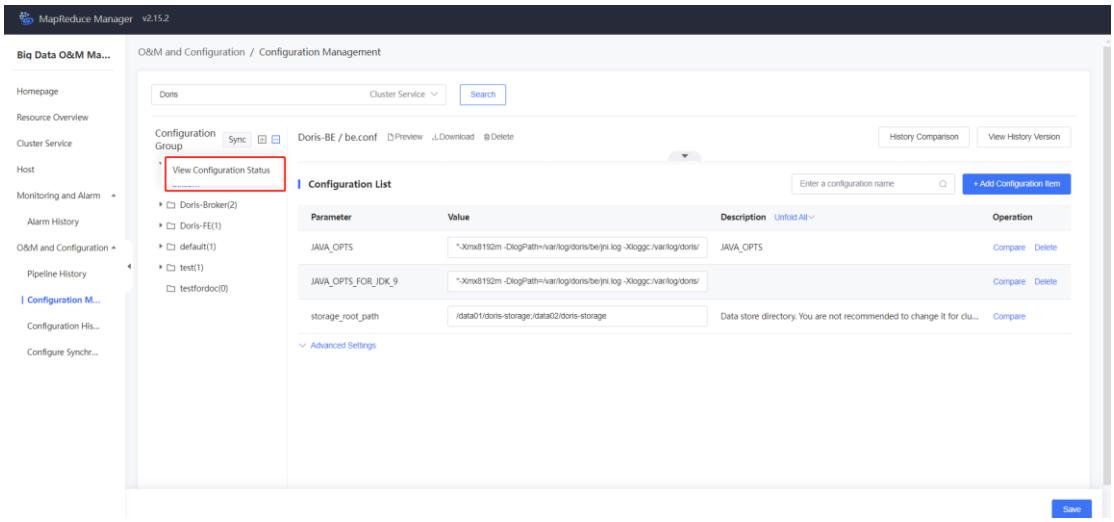

7. A dialog box displaying the synchronization status of all configuration files will appear. As shown in the figure:

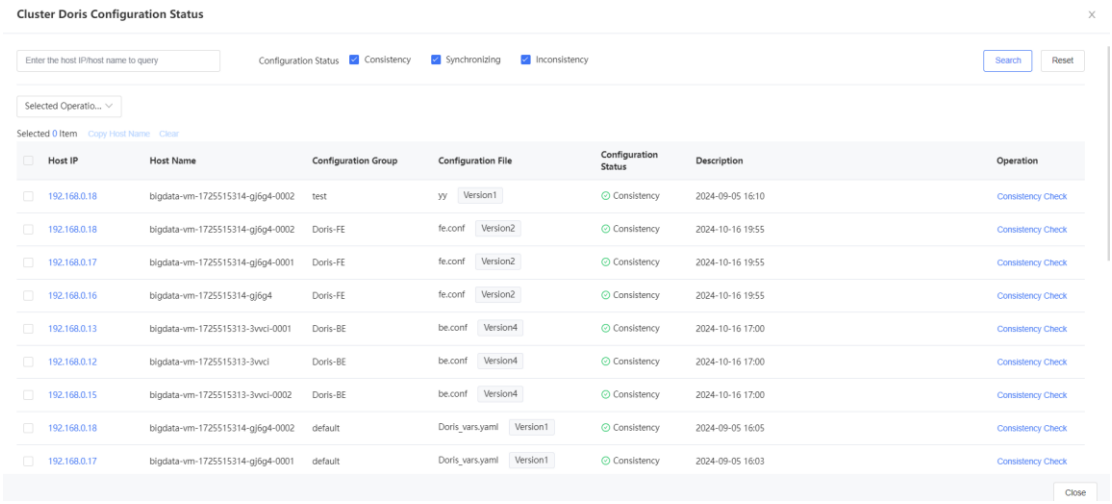

Configuration Management - Resynchronization

Scenario

This feature supports resynchronization of configurations that have failed to sync or have inconsistent sync statuses.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Management** on the menu.
5. Select the required cluster service for query conditions.
6. Click **More > View Configuration Status**. A configuration synchronization status dialog box will appear.
7. Check the box in front of the host that requires operation.
8. Click **Operate on Selected > Re-sync** to perform the configuration synchronization. As shown in the figure:

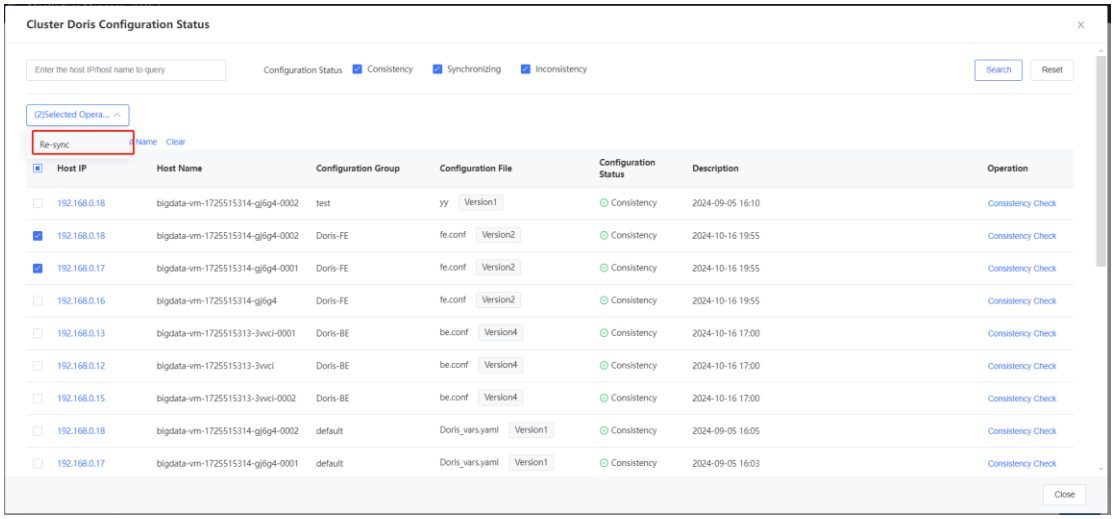

Configuration Management - Configuration Rollback

Scenario

This feature supports the rollback of a specified configuration file to a specific historical version.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Management** on the menu.
5. Select the required cluster service for query conditions.
6. Click on the configuration group, and then on the name of the configuration file. Details of this configuration will be displayed on the right side of the page.
7. Click on **More > History Comparison**, and a history comparison dialog box will appear. As shown in the figure:

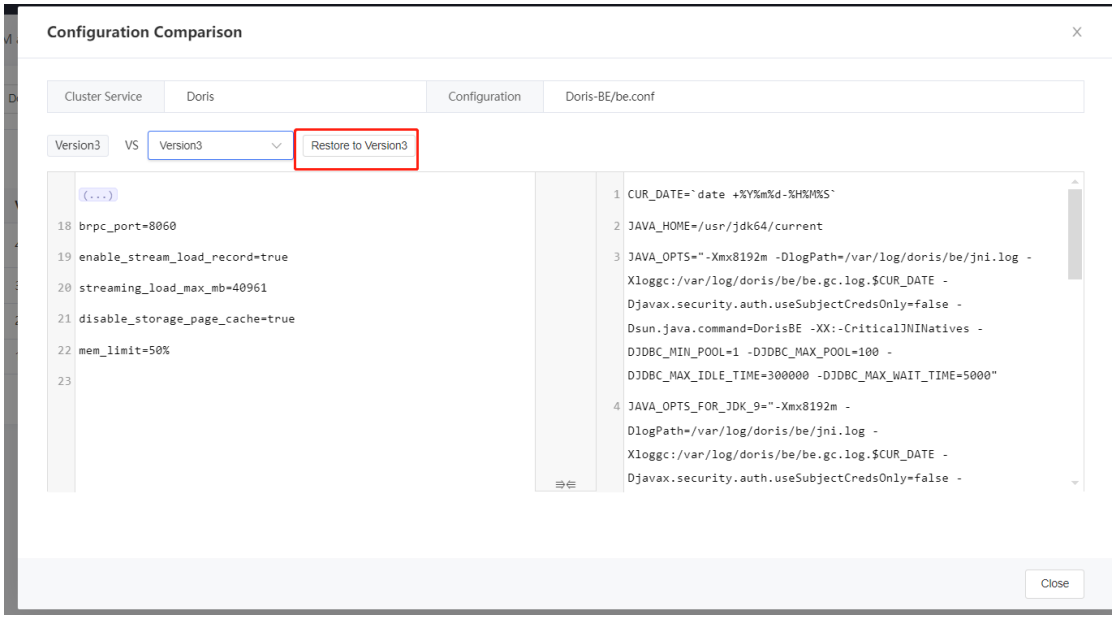

8. Click the **Restore to Version X** button. A confirmation dialog box will appear. As shown in the figure:

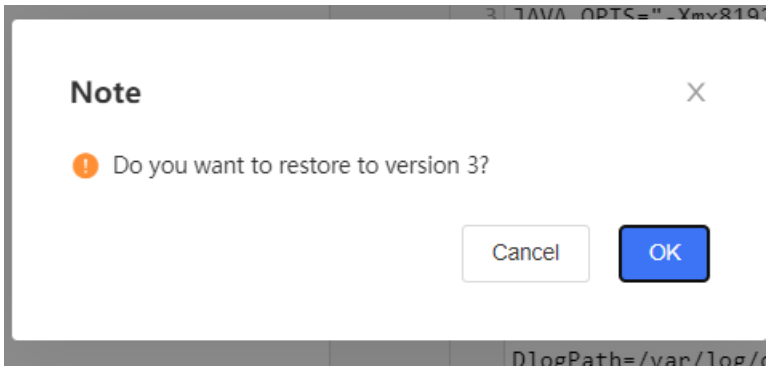

9. Click **OK**.
10. Click the **Close** button to close the configuration comparison dialog box.
11. Check the configuration file. The content of the configuration file should now match the restored version.
12. Click **Sync** to synchronize the configuration file.

Configuration History - Configuration History Overview

Once in the History, click **O&M and Configuration> Configuration History** on the menu to go to the configuration history page. As shown in the figure:

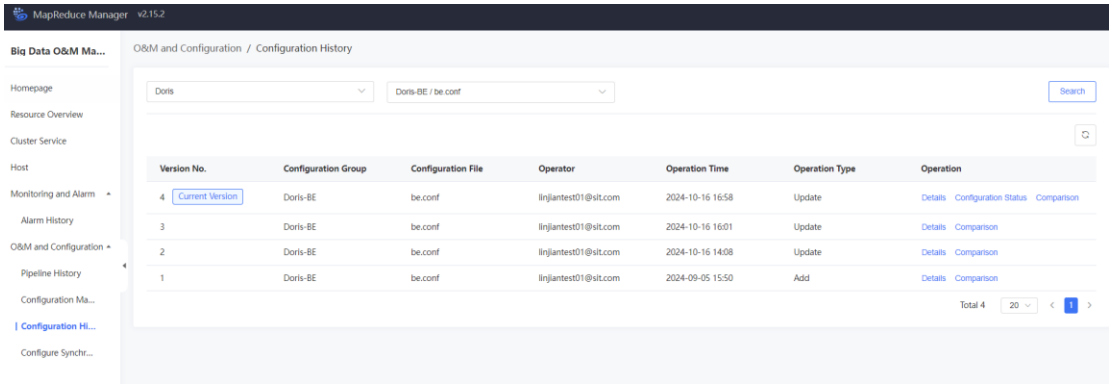

The query section is located at the top of the page.

- Supports viewing all history version information of different configuration files.
- Supports viewing the content of different versions of a configuration file and performing content comparisons between these versions.
- Supports viewing the synchronization details of the current version of a configuration file.

Configuration History - Query Configuration History

Scenario

Users can query the historical versions of a configuration file.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration History** on the menu.
5. For query conditions, select the required cluster service, configuration group, and configuration file. The page will display all historical versions of this configuration file. As shown in the figure:

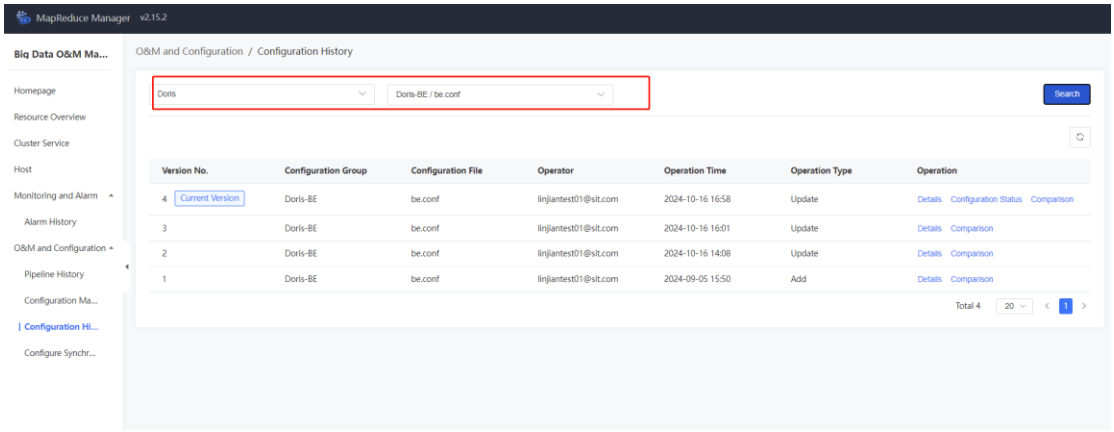

Configuration History - View Configuration Details

Scenario

After executing configuration synchronization, users may need to view the details of different configuration file versions.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration History** on the menu.
5. For query conditions, select the required cluster service, configuration group, and configuration file.
6. Choose the version you want to view and click the Details button. A dialog box will appear showing the details of the configuration file version. As shown in the figure:

Configuration History - View Configuration Status

Scenario

Viewing the configuration status of the latest version is supported.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration History** on the menu.

5. For query conditions, select the required cluster service, configuration group, and configuration file.

6. Click on **Configuration Status**. A configuration status dialog box will appear. As shown in the figure:

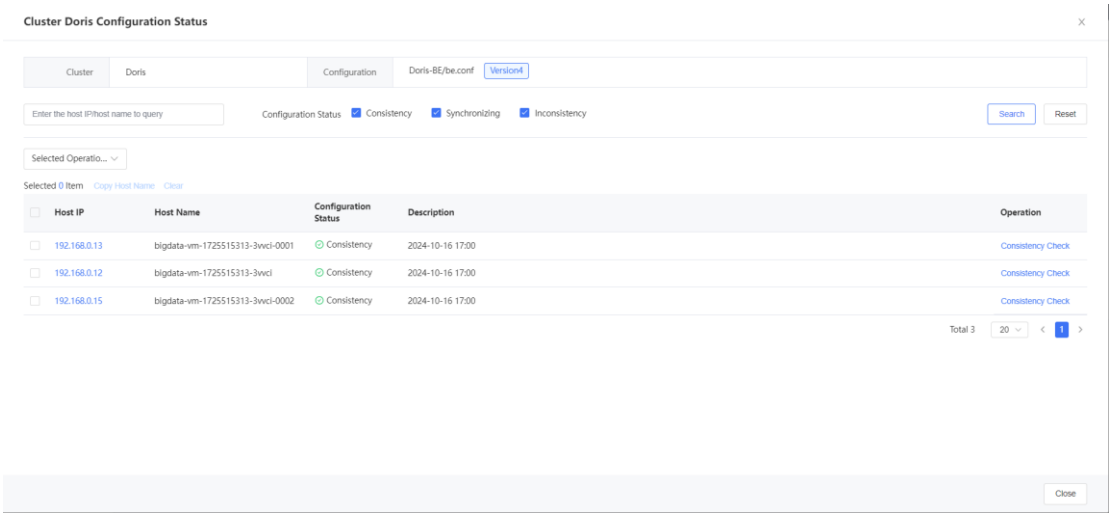

Configuration History - View Configuration Comparison

Scenario

This feature allows users to compare different versions of a configuration file.

Procedure

1. Log in to the YI-MapReduce management console.

2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.

3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.

4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration History** on the menu.

5. For query conditions, select the required cluster service, configuration group, and configuration file.

6. Click the **Compare** button. A configuration comparison dialog box will appear. As shown in the figure:

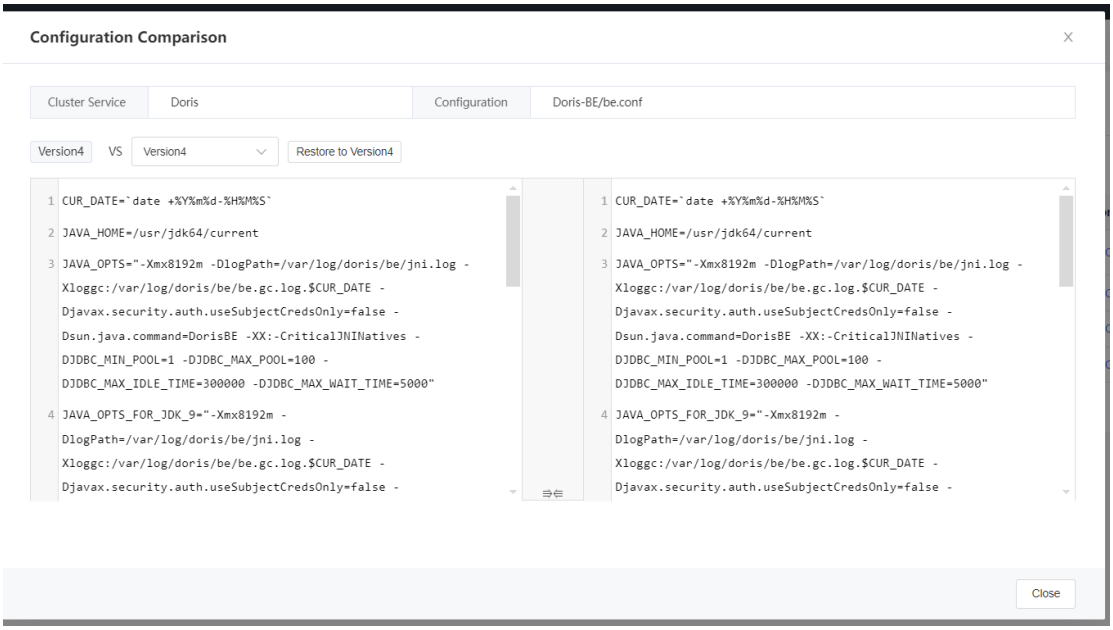

Configuration History - Resynchronization

Scenario

Viewing the configuration status of the latest version is supported.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration History** on the menu.
5. For query conditions, select the required cluster service, configuration group, and configuration file.
6. Click on **Configuration Status**. A configuration status dialog box will appear.

7. Check the box in front of the host that requires operation.
8. Click **Operate on Selected > Re-sync** to perform the configuration synchronization. As shown in the figure:

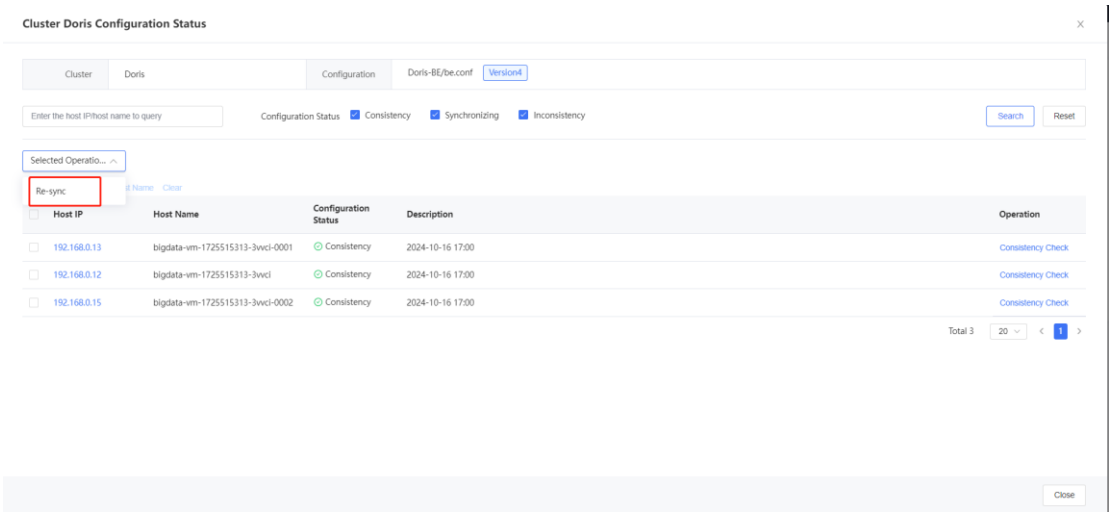

Configuration Sync History - Configuration

Synchronization History Overview

Once in the History, click **O&M and Configuration> Configuration Synchronization History** on the menu to go to the configuration sync history page.

As shown in the figure:

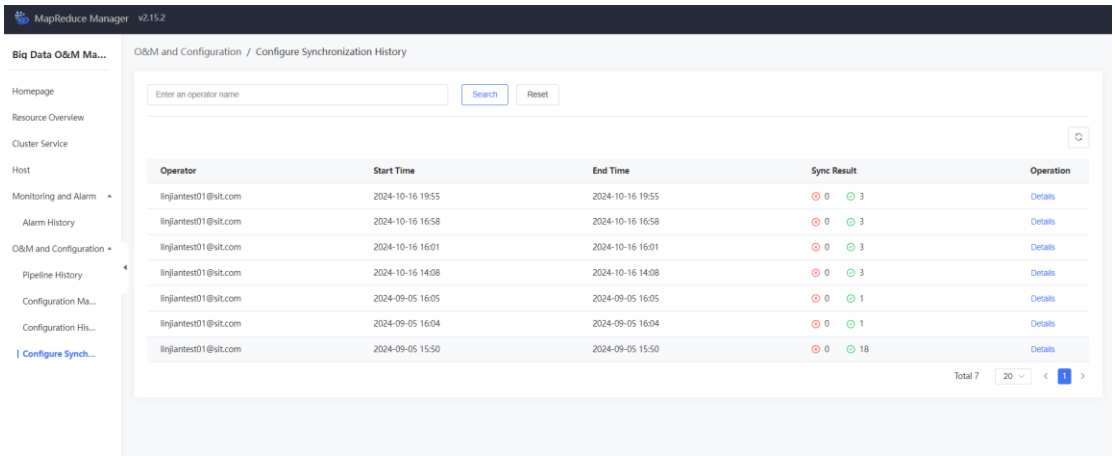

- The query section is located at the top of the page.

- The page displays all configuration synchronization operations, their operators, and details.

Configuration Sync History - View Configuration Sync Status

Scenario

After performing configuration synchronization in Configuration Management, users may need to view the synchronization status or results of this operation.

Procedure

1. Log in to the YI-MapReduce management console.
2. Click on **My Cluster** and choose a running cluster to enter the cluster information page.
3. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
4. Once in the YI-MapReduce Manager, click **O&M and Configuration> Configuration Synchronization History** on the menu. As shown in the figure:

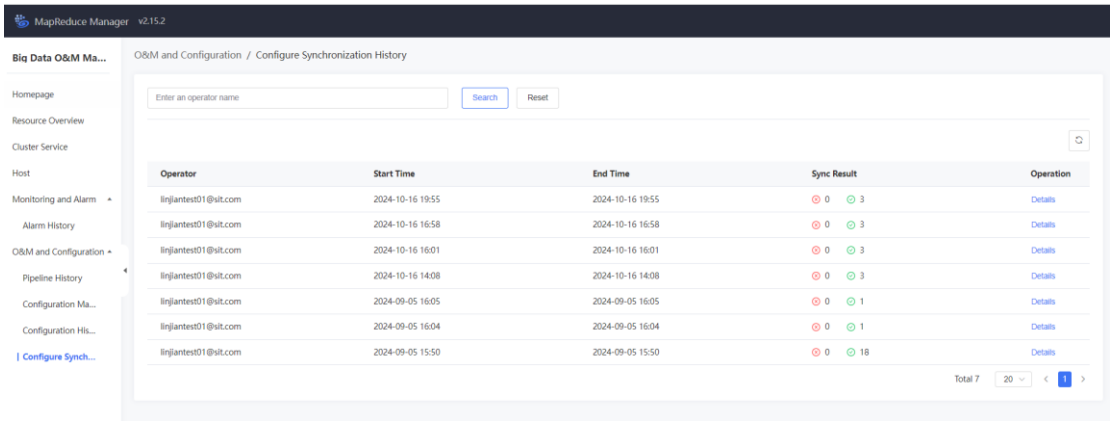

5. Select the operation record you want to view, click on **Synchronizing** or the synchronization result, to enter the configuration synchronization detail page. As shown in the figure:

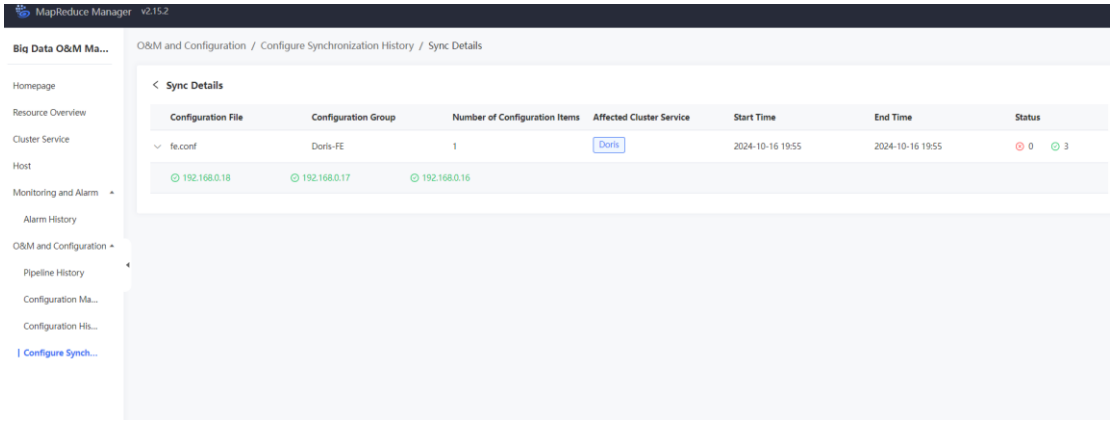

4.6. Overview of High-risk Operations

4.6.1. Prohibited Operations

The following table describes the forbidden operations that should be noted during the routine operations in the cluster operation and maintenance stage.

Type	Operation Risk
Changing the host name of a node is strictly prohibited	Changing the host name will cause the instances on that host and higher-level components to fail to provide normal service, and this cannot be fixed.
Reinstallation of node OS is prohibited	This operation will cause the YI-MapReduce cluster to enter an abnormal state, affecting the use of the YI-MapReduce cluster.
Use of private images is prohibited	This operation will cause the YI-MapReduce cluster to enter an abnormal state, affecting the use of the YI-MapReduce cluster.

4.6.2. High-risk Operations

The following tables describe high-risk operations that need to be considered during the O&M stage of each component.

Cluster High-risk Operations

Action Name	Operation Risk	Risk Level	Risk Mitigation	Major Operation Observing Items
Binding EIP	This operation exposes master nodes (such as the Doris FE) of the cluster service to the public network, increasing the risk of network attacks from the Internet.	★★★★★	Ensure that the bound EIP is a trusted public access IP. Verify if the corresponding ports have set security group rules, allowing only trusted IPs to access these ports. It is not recommended to allow 0.0.0.0 to access in the inbound direction rule.	None
Opening cluster 22 port security group rule	This operation increases the risk of users exploiting port 22 for vulnerability attacks.	★★★★★	Set security group rules for the open port 22, allowing only trusted IPs to access this port. It is not recommended to allow 0.0.0.0 to access in the inbound direction rule.	None
Deleting cluster or cluster data	This operation can lead to data loss.	★★★★★	Please confirm the necessity of this operation before deletion, and ensure that data backup has been completed.	None

Scaling down cluster	This operation can lead to data loss.	★★★★★	Please confirm the necessity of this operation before scaling down, and ensure that data backup has been completed.	None
Unmounting disk or formatting data disk	This operation can lead to data loss.	★★★★★	Please confirm the necessity of this operation before proceeding, and ensure that data backup has been completed.	None

High-risk Operations in YI-MapReduce Manager

Action Name	Operation Risk	Risk Level	Risk Mitigation	Major Operation Observing Items
Modifying log level	If modified to DEBUG, the running speed of the Manager will noticeably slow down.	★★	Confirm the necessity of the operation before modification, and revert to the default setting timely.	None
Restarting underlying service with "Also restart upper-layer	This operation will interrupt the upper layer service business, affecting the management, maintenance, and	★★★★	Confirm the necessity of the operation before execution, and ensure no	Observe for any unresolved alarms, check if cluster management and

service" checked	business of the cluster.		other maintenance operations are executing concurrently.	maintenance are normal, and if the business is operating normally.
Restarting service	The service will be interrupted during the restart. If "Also restart upper-layer service" is checked, any upper-layer service that depends on this service will be interrupted.	★★★	Confirm the necessity of the restart before execution.	Observe for any unresolved alarms, check if cluster management and maintenance are normal, and if the business is operating normally.
Modifying default SSH port of node	Modifying the default port (22) will cause the Create Cluster, Add Service/Instance, Add Host, Reinstall Host functions to be unusable, and cause the cluster health check results, including node trust, to be inaccurate.	★★★	Change the SSH port back to the default value before performing related operations.	None

4.7. Appendix

4.7.1. Scenario

If the specifications (vCPU and memory) of the ECS instances within a node group fail to meet your business needs, you can enhance these ECS instance specifications

through the feature of configuration upgrade. After upgrading the configuration, you need to manually adjust the configuration information of services such as Doris, Elasticsearch. This section guides you on how to modify configuration items on the Configuration Management page in YI-MapReduce Manager.

4.7.2. Before You Begin

You have created a cluster.

4.7.3. Procedure

1. In YI-MapReduce Manager, click on **O&M and Configuration**.
2. Click on **Configuration Management**.
3. Choose the **Selected Cluster Services**, click **Query**, and you can modify the configuration information on the current page.

4.7.4. Suggestions for Configuration Modifications

1. Doris: After the configuration of the node where Doris resides is upgraded, it is recommended that Doris FE uses half of the node's memory.
2. Elasticsearch: For a single Elasticsearch node, it is recommended that the memory does not exceed 64G. After the configuration of the node where Elasticsearch resides is upgraded, Elasticsearch will automatically set the memory value according to the node conditions. Generally, no manual modification is needed, but the cluster needs to be restarted. If you wish to manually set it, you can modify the -Xms30g and -Xmx30g parameters in the jvm.options file and restart the cluster.

5. Best Practices

5.1. Enhancing User Experience and Performance in Kibana

Kibana is a robust tool for visualizing and analyzing data in ElasticSearch. To fully leverage the capabilities of Kibana, improve user experience, and deepen data insights, here are some best practices for Kibana.

5.1.1. 1. Dashboard Design

Defined Objectives

Before creating a dashboard, it is crucial to define your objectives and requirements. Identify the insights you aim to extract from the data to tailor the design of the visualization components accordingly. For example, if you are a data analyst for an e-commerce platform, you might want to design a dashboard to monitor product sales. Your goal would be to identify the most popular product categories and buying trends.

Simplicity and Focus

Avoid overcrowding dashboards and try to ensure simplicity. Highlight key metrics and information to prevent distracting users. Refrain from cluttering a single dashboard with excessive information, ensuring that users can focus on key metrics.

5.1.2. 2. Visualization

Choosing the Appropriate Chart Type

Select the chart type that best fits the characteristics of your data, such as line charts, bar charts, or pie charts. Make sure the type of chart can effectively represent your data. For example, to show trends in product sales, a line chart may be your choice. This can clearly demonstrate how sales for various categories evolve over time.

Axis Labels and Titles

Add clear labels to the axes and titles to the charts to help users understand what the data means. Labels should be straightforward and simple, and avoid overly technical terms. For example, add labels for the horizontal and vertical axes to a line chart, denoting time and sales, respectively. Add a title to the chart such as "Product Sales Trends".

5.1.3. 3. Search and Filter

Using Filters

Use the filters of Kibana for quick selection and data filtering. Establish predefined filters for common filtering conditions, enabling users to quickly retrieve specific data. For example, create predefined filters such as "Electronics", "Clothing", and "Home Goods" to facilitate easy data filtering for specific product categories.

Using Query Language

Become familiar with the query language of Kibana to accurately extract required data during searches. Use this query language to structure complex search conditions. For example, to further filter products with sales exceeding a certain value, the query language can be used, such as sales_amount: >1000.

5.1.4. 4. Dashboard Layout

Flexible Layout

Use the flexible dashboard layout feature of Kibana, to arrange visualization components freely and adapt them according to specific requirements.

Dashboard Spacing

Ensure proper spacing between dashboard components to prevent overlap or congestion.

5.1.5. 5. Security

Access Control

Introduce access control based on user roles and permissions to ensure that specific dashboards and data are accessible only to authorized users. For example, configure the system so that only data analysts have access to the sales data dashboard, while it is not accessible to other personnel.

Processing of Sensitive Data

For dashboards that display sensitive data, consider implementing data masking or using field-level permissions to limit access to sensitive information.

5.1.6. 6. Performance and Optimization

Data Caching

Use the data caching feature of Kibana to decrease ElasticSearch query load and enhance dashboard loading speed.

Data Update Frequency

Set the dashboard refresh interval based on data update frequency to avoid excessive data refreshing.

These best practices for Kibana are designed to assist in improved dashboard design, creation, and usage. However, please note that actual practices may differ based on specific requirements and scenarios.

5.2. Elasticsearch Cluster Planning and Performance Optimization

Elasticsearch is an efficient distributed search and analytics engine, applicable to storing, retrieving, and analyzing large-scale data. To ensure optimal performance, stability, and scalability, here are some best practices for Elasticsearch.

5.2.1. 1. Cluster Planning

Number of Nodes

The number and distribution of nodes play a key role in constructing an Elasticsearch cluster. An odd number of nodes is typically preferred for better handling of master node elections and failovers. For example, a cluster with three master nodes can maintain functionality even if two nodes fail.

Single-node clusters should be avoided as they can lead to single points of failure.

Shards and Replicas

Shards are horizontal partitions of data within the cluster. When creating an index, select an appropriate number of shards based on data amount and query load, ideally no more than the number of nodes. Then, to assure data availability, set a sufficient number of replicas, but avoid setting too many to prevent resource waste.

Example:

```
PUT /my_index
{
"settings": {
"number_of_shards": 5,
"number_of_replicas": 1
}
}
```

5.2.2. 2. Index Design

Settings for Shards and Replicas

When creating an index, select an appropriate number of shards based on data amount and query load.

For frequently accessed indexes, consider increasing the number of replicas to enhance query performance.

Index Fields

Make wise choices for field types and avoid unnecessarily complex types. For example, use the keyword type for exact values, and the text type for full-text search.

Example:

```
PUT /my_index
{
"mappings": {"properties": {
  "product_name": {
    "type": "text"
```

```
  },
  "category": {
    "type": "keyword"
  },
  "price": {
    "type": "float"
  }
}
}
}
```

Disable fields that are not required to be indexed to reduce index size and improve performance.

5.2.3. 3. Write Operations

Bulk Write

Use the bulk API for write operations to minimize network overhead. In high-throughput write scenarios, consider moderately increasing the Refresh interval to reduce the number of index refresh operations and enhance write performance.

Example:

```
POST /my_index/_doc/_bulk
{ "index": {}}
{ "field1": "value1" }
{ "index": {}}
{ "field2": "value2" }
```

Index Refresh

Avoid frequent index refreshes. Appropriately adjust the Refresh interval to balance between write and query performance.

Example:

```
PUT /my_index/_settings
{
"refresh_interval": "30s"
}
```

5.2.4. 4. Query Optimization

Query Performance

Use the query DSL for complex searches to take full advantage of the capabilities of Elasticsearch.

Enhance query performance by using index aliases and index templates.

Example:

```
GET /my_index/_search
{
"query": {
"bool": {
  "must": [
    { "match": { "product_name": "apple" }},
    { "range": { "price": { "gte": 100 }}}
  ]
}
```

```
}

}
```

Pagination Query

For pagination queries involving a large amount of data, use Scroll to avoid performance issues associated with deep pagination.

Example:

```
POST /my_index/_search?scroll=5m

{

"query": { "match_all": {}},

"size": 100

}
```

5.2.5. 5. Hardware and Monitoring

Hardware Selection

Select high-performance hardware, particularly fast disks and large memory.

Use SSD hard drives to enhance indexing and query performance.

Monitoring and Alarm

Establish monitoring metrics such as cluster health, node status, and resource usage.

Monitor the cluster continuously with monitoring tools and set up alarms for timely issue response.

5.2.6. 6. Security

Authentication and Authorization

Enable security features and authenticate using usernames and passwords.

Restrict user access using roles and permissions.

Encrypted Communication

Enable TLS/SSL encryption to secure data during transmission.

5.2.7. 7. Maintenance

Backup and Recovery

Ensure regular backups of index data to facilitate swift data recovery in case of unforeseen events.

6. API Reference

6.1. API Overview

YI-MapReduce provides the following related API interfaces.

6.1.1. Cluster

API Title	API Description
Query Cluster Information by ID	This API provides the function to query cluster information by ID. The system takes query conditions from user input and returns the details of clusters that meet these conditions.
Page Query of Cluster Information	This API provides the function to query cluster information. The system takes query conditions from user input and returns the details of clusters that meet these conditions under the resource pool specified by the user. Information of deleted clusters will not be returned.

6.1.2. Node Group

API Title	API Description
Query Cluster Node Group Details	This API provides the function to query the details of cluster node groups. The system takes query conditions from user input and returns the details of the cluster node groups that meet these conditions.
Get Node Group Information	This API provides the function to query the details of cluster node groups. The system takes query conditions from user input and returns the details of the cluster node groups that meet these conditions.

6.2. Call APIs

6.2.1. Query Cluster Node Group Details

API Introduction

This API provides the function to query the details of cluster node groups. The system takes query conditions from user input and returns the details of the cluster node groups that meet these conditions.

API Restrictions

You can only query details of node groups from clusters that are in a running state.

Require API Approval or Not

No

URI

POST /v1/emr/openapi/cluster/clusterNodeGroup/getGroupAndHostByCondition

Content-Type

application/json

Path Parameter

None

Query Parameter

None

Request Header Parameter

Request Header Parameter

Parameter	Required or Not	Parameter Type	Description	Example	Lower-level Object
Cluster-Id	Yes	String	Cluster ID	bdbbc8cc1288a78e5851839dc26f58b9	/

Request Body Parameter

Request Body Parameter

Parameter	Required or Not	Parameter Type	Description	Example	Lower-level Object
clusterId	Yes	String	Cluster ID	6695c8e7c34af927e860ba1f043c8c91	/
nodeState	No	Integer	Host status (1: Starting, 2: Shutting down, 4: Deleted, 8: Running, 9:	1	/

			Shutdown).		
selectKey	No	String	A fuzzy query parameter that enables fuzzy searches through the node name, internal IP, or external IP.	10.0.02	/

Response Parameter

Response Parameter

Parameter	Parameter Type	Description	Example	Lower-level Object
statusCode	Integer	Status Code	200	/
message	String	Brief summary of the current API call status and necessary prompt information.	Requested	/
error	String	Error code. This field is not returned when the request is successful.	EMR_400001	/

returnObj	Object	Return result.	/	returnObj

Table returnObj

Parameter	Parameter Type	Description	Example	Lower-level Object
id	String	Node Group ID	175187849787779	/
clusterId	String	Cluster ID	da595eb1d81503b323fdc01d9bf786b7	/
payType	Integer	Payment Type	1	/
nodeGroupType	String	Code for the node type	MASTER	/
nodeGroupName	String	Code value of the node type	master	/
imageId	String	Image ID	f8d28855-ffda-426d-96e2-0f6bf46877e0	/
hostNum	Integer	Number of hosts	3	/
computeSpecificationId	Integer	ID of the host specification	101	/
iaasVmSpecCode	String	Code for the IAAS	s7.2xlarge.4	/

		VM specification		
cpuNum	Integer	List of cluster CPU core numbers	8	/
memory	Integer	Size of memory	32	/
diskSpecificationList	String	List of disk specifications	[{diskFunctionType:1,diskFunctionName:System Disk,diskType:cloud,diskTypeName:Cloud Disk,ioType:SSD-genric,ioTypeName:General Purpose SSD,volume:200,diskNum:1}]	/
createTime	Integer	Creation time	1706515356000	/
updateTime	Integer	Update time	1706515356000	/
mountPublicIp	Boolean	Whether public IP is mounted	false	/
iaasVmSpecId	String	ID of the IAAS VM specification	b307034d-cbc3-27bb-24f7-a97565814236	/

highest	Boolean	Whether the node group specification is the highest	false	/
regionId	String	Resource pool ID	bb9fdb42056f11eda1610242ac110002	/
availableZoneId	String	Availability zone ID	cn-huadong1-jsnj1A-public-ctcloud	/
vpcId	String	VPC ID	vpc-fr2xjo1gj0	/
clusterHostDtoList	Array of Objects	List of cluster node groups	/	clusterHostDtoList

Table clusterHostDtoList

Parameter	Parameter Type	Description	Example	Lower-level Object
id	String	Primary key ID	0088babec7bab47b20409df6646db856	/
nodeGroupId	String	Node Group ID	1753298494213554178	/
iaasHostId	String	Platform host ID	b2f2377d-3f0b-bc74-b9c5-67bb557b0e63	/

hostName	String	Host Name	bigdata-emr-vm-l7zb3wnk	/
manageIp	String	Management IP	10.2.3.0	/
serviceIp	String	Intranet IP	192.168.0.149	/
publicIp	String	Internet IP	212.168.0.150	/
state	Integer	Enumeration of host states at the point of creation, not the real-time state	8	/
vpcId	String	VPC ID	vpc-fr2xjo1gj0	/
subnetId	String	Subnet ID	subnet-vkixrw8xw7	/
regionId	String	Resource pool ID	bb9fdb42056f11eda1610242ac110002	/
availableZoneId	String	Availability zone ID	cn-huadong1-jsnj1A-public-ctcloud	/
eipId	String	EIP ID	eip_ssww	/
createTime	Integer	Creation time	1706515356000	/

updateTime	Integer	Update time	1706515356000	/
hostStateValue	String	Host state	Running	/
deployRoleInstance	Array of Strings	Deployed role instance	["NodeMaster"]	/

Enumeration Parameters

None

Request Header Example

```
{
    "Cluster-Id":"bdbbc8cc1288a78e5851839dc26f58b9"
}
```

Request URL Example

```
https://emr-global.ctapi.ctyun.cn/v1/emr/cluster/openapi/clusterNodeGroup/getGroupAndHostByCondition
```

Request Body Example

```
{
    "clusterId":"bdbbc8cc1288a78e5851839dc26f58b9",
    "selectKey":"192.168.0.208"
}
```

Response Example

Successful request return value example


```
{
    "statusCode": 200,
    "message": "success",
    "returnObj": [
        {
            "id": "1759528882838106113",
            "clusterId": "bdbbc8cc1288a78e5851839dc26f58b9",
            "payType": 1,
            "nodeGroupType": "CORE",
            "nodeGroupName": "core",
            "imageId": "f8d28855-ffda-426d-96e2-0f6bf46877e0",
            "hostNum": 3,
            "computeSpecificationId": 101,
            "iaasVmSpecCode": "s7.2xlarge.4",
            "cpuNum": 8,
            "memory": 32,
            "diskSpecificationList":
"[{\"diskFunctionType\":1,\"diskFunctionName\":\"System
Disk\",\"diskType\":\"cloud\",\"diskTypeNamfe\":\"Cloud
Disk\",\"ioType\":\"SATA\",\"ioTypeName\":\"Common
IO\",\"volume\":200,\"diskNum\":1},{\"diskFunctionType\":2,\"diskFunctionNa
me\":\"Data Disk\",\"diskType\":\"cloud\",\"diskTypeName\":\"Cloud
```

Disk\",\"ioType\":\"SATA\",\"ioTypeName\":\"Common IO\",\"volume\":200,\"diskNum\":2}]",
"createTime": 1708339350000,
"updateTime": 1708339350000,
"mountPublicIp": false,
"iaasVmSpecId": "b307034d-cbc3-27bb-24f7-a97565814236",
"highest": null,
"regionId": "bb9fdb42056f11eda1610242ac110002",
"availableZoneId": "cn-huadong1-jsnj1A-public-ctcloud",
"vpcId": "vpc-fr2xjo1gj0",
"clusterHostDtoList": [
{
"id": "af84f22fdefd94bb30f7af4c24368337",
"nodeGroupId": "1759528882838106113",
"iaasHostId": "3d328644-8521-89f7-a114-74b4de602c50",
"hostName": "bigdata-emr-vm-iwxkgpil",
"manageIp": "-",
"serviceIp": "192.168.0.208",
"publicIp": "-",
"state": 8,
"vpcId": "vpc-fr2xjo1gj0",
"subnetId": "subnet-vkixrw8xw7",
"regionId": "bb9fdb42056f11eda1610242ac110002",
"availableZoneId": "cn-huadong1-jsnj1A-public-ctcloud",

```
                "eipId": "",
                "createTime": 1708339353000,
                "updateTime": 1708339353000,
                "hostStateValue": "Running",
                "deployRoleInstance": []
            }
        ]
    }
  ]
}
```


Failed request return value example


```
{
    "statusCode": 200,
    "error": "EMR_400010",
    "message": "Error with clusterId parameter type",
    "returnObj": []
}
```


Status Code

Please refer to status codes

Error Codes

Please refer to error codes

6.2.2. Query Cluster Information by ID

API Introduction

This API provides the function to query cluster information by ID. The system takes query conditions from user input and returns the details of clusters that meet these conditions.

API Restrictions

You can only query information of clusters that are in a running state.

Require API Approval or Not

No

URI

GET /v1/emr/openapi/cluster/clusterDetail/getById

Content-Type

application/json

Path Parameter

None

Query Parameter

Query Parameter

Parameter	Required or Not	Parameter Type	Description	Example

id	Yes	String	Cluster ID	00c3a04292996955752f073c995a1cc6

Request Header Parameter

Request Header Parameter

Parameter	Required or Not	Parameter Type	Description	Example	Lower-level Object
Cluster-Id	Yes	String	Cluster ID	00c3a04292996955752f073c995a1cc6	/

Request Body Parameter

None

Response Parameter

Response Parameter

Parameter	Parameter Type	Description	Example	Lower-level Object
statusCode	Integer	Status Code	200	/
message	String	Brief summary of the current API call status and necessary prompt information.	Requested	/

error	String	Error code. This field is not returned when the request is successful.	EMR_400001	/
returnObj	Object	Return result.	/	returnObj

Table returnObj

Parameter	Parameter Type	Description	Example	Lower-level Object
id	String	Cluster ID	00c3a04292996955752f073c995a1cc6	/
managerClusterId	Integer	Cluster ID as defined by the manager	1	
iaasType	String	Platform	Public Cloud	/
regionId	String	Resource pool ID	bb9fdb42056f11eda1610242ac110002	/
regionName	String	Resource pool name	East China 1	/
availableZoneId	String	Availability zone ID	cn-huadong1-jsnj1A-public-ctcloud	/
availableZoneName	String	Availability zone name	AZ 1	/

clusterName	String	Cluster Name	test_vpc_24	/
payType	String	Payment Type	Annual and Monthly Subscription	/
clusterType	String	Cluster Type	Elastic Search	/
clusterTypeVersion	String	Product Version	YI-MapReduce-2.12.0	/
clusterPlanCode	String	Code of the cluster plan	cloud-search	/
componentNameList	String	List of component names	[{componentTitle:ElasticSearch,version:7.10.2}, {componentTitle:Kibana,version:7.10.2}]]	/
datasourceConfigs	String	Information about the data source	[]	/
vpcId	String	VPC ID	vpc-0k5xl6w5	/
subnetId	String	Subnet ID	subnet-i2ys8sp	/
securityGroupId	String	Security Group ID	[sg-4h7w9cl1]	/
clusterDueTime	Integer	Expiry time of the cluster	1709193751000	/

userId	String	User ID	ba14c8e729e447d69698f81ac7d55555	/
accountId	String	Account ID	ed24e4b414a048b0a9cb995f59cc85jj	/
clusterCreateTime	Integer	Creation time of the cluster	1706515357000	/
clusterState	String	Cluster Status	Running	/
createTime	Integer	Creation time	1706515356000	/
updateTime	Integer	Update time	1706515356000	/
pathMap	Map of String	Component Name: Component Link	{ "Kibana": "1,1", "ElasticSearch": "1,1" }	/

Enumeration Parameters

None

Request Header Example

```
{
    "Cluster-Id":"00c3a04292996955752f073c995a1cc6"
}
```

Request URL Example

https://emr-global.ctapi.ctyun.cn/v1/emr/openapi/cluster/clusterDetail/getById?id=00c3a04292996955752f073c995a1cc6

Request Body Example

None

Response Example

Successful request return value example


```
{
    "statusCode": 200,
    "message": "success",
    "returnObj": {
        "id": "00c3a04292996955752f073c995a1cc6",
        "managerClusterId": 1,
        "iaasType": "Public Cloud",
        "regionId": "bb9fdb42056f11eda1610242ac110002",
        "regionName": "East China 1",
        "availableZoneId": "cn-huadong1-jsnj1A-public-ctcloud",
        "availableZoneName": "Availability Zone 1",
        "clusterName": "test1218",
        "payType": "Annual or Monthly",
        "clusterType": "Elastic Search",
        "clusterTypeVersion": "YI-MapReduce-2.12.0",
```

```
        "clusterPlanCode": "cloud-search",
        "componentNameList":
"[{\"componentTitle\":\"ElasticSearch\",\"version\":\"7.10.2\"},{\"componentTitl
e\":\"Kibana\",\"version\":\"7.10.2\"}]",
        "datasourceConfigs": "[]",
        "vpcId": "vpc-vgo9wazk0m",
        "subnetId": "subnet-6hd5rpedsp",
        "securityGroupId": "[\"sg-mz2suubydc\"]",
        "clusterDueTime": 1705543555000,
        "userId": "ba14c8e729e447d69698f81ac7d506b5",
        "accountId": "ed24e4b414a048b0a9cb995f59cc85dd",
        "clusterCreateTime": 1702865155000,
        "clusterState": "Running",
        "createTime": 1702865154000,
        "updateTime": 1702979599000,
        "pathMap": {
          "Kibana": "1,1",
          "ElasticSearch": "1,1"
        }
    }
}
```


Failed request return value example


```
{
```

```
    "statusCode": 200,
    "error": "EMR_400019",
    "message": "Cluster ID is required",
    "returnObj": []
}
```


Status Code

Please refer to status codes

Error Codes

Please refer to error codes

6.2.3. Get Node Group Information

API Introduction

This API provides the function to query the details of cluster node groups. The system takes query conditions from user input and returns the details of the cluster node groups that meet these conditions.

API Restrictions

You can only query node group information of clusters that are in a running state.

Require API Approval or Not

No

URI

GET /v1/emr/openapi/cluster/clusterNodeGroup/getByClusterId

Content-Type

application/json

Path Parameter

None

Query Parameter

Query Parameter

Parameter	Required or Not	Parameter Type	Description	Example
clusterId	Yes	String	Cluster ID	da595eb1d81503b323fdc01d9bf786b7

Request Header Parameter

Request Header Parameter

Parameter	Required or Not	Parameter Type	Description	Example	Lower-level Object
Cluster-Id	Yes	String	Cluster ID	da595eb1d81503b323fdc01d9bf786b7	/

Request Body Parameter

None

Response Parameter

Response Parameter

Parameter	Parameter Type	Description	Example	Lower-level Object
statusCode	Integer	Status Code	200	/
message	String	Brief summary of the current API call status and necessary prompt information.	Requested	/
error	String	Error code. This field is not returned when the request is successful.	EMR_400001	/
returnObj	Array of Objects	Return result.	/	returnObj

Table returnObj

Parameter	Parameter Type	Description	Example	Lower-level Object
id	String	Node Group ID	175187849787779	/
clusterId	String	Cluster ID	da595eb1d81503b323fdc01d9bf786b7	/

payType	Integer	Payment Type	1	/
nodeGroupType	String	Code for the node type	MASTER	/
nodeGroupName	String	Code value of the node type	master	/
imageId	String	Image ID	f8d28855-ffda-426d-96e2-0f6bf46877e0	/
hostNum	Integer	Number of hosts	3	/
computeSpecificationId	Integer	ID of the host specification	101	/
iaasVmSpecCode	String	Code for the IAAS VM specification	s7.2xlarge.4	/
cpuNum	Integer	Number of CPU cores	8	/
memory	Integer	Size of memory	32	/
diskSpecificationList	String	List of disk	[{diskFunctionType:1,diskFunctionName:System	/

		specifications	Disk,diskType:cloud,diskTypeName:Cloud Disk,ioType:SSD-genric,ioTypeName:General Purpose SSD,volume:200,diskNum:1}]	
createTime	Integer	Creation time	1706515356000	/
updateTime	Integer	Update time	1706515356000	/
mountPublicIp	Boolean	Whether public IP is mounted	false	/
iaasVmSpecId	String	ID of the IAAS VM specification	b307034d-cbc3-27bb-24f7-a97565814236	/
highest	Boolean	Whether the node group specification is the highest	false	/

Enumeration Parameters

None

Request Header Example

```
{
    "Cluster-Id":"da595eb1d81503b323fdc01d9bf786b7"
}
```

Request URL Example

```
https://emr-
global.ctapi.ctyun.cn/v1/emr/openapi/cluster/clusterNodeGroup/getByClusterI
d?clusterId=da595eb1d81503b323fdc01d9bf786b7
```

Request Body Example

None

Response Example

Successful request return value example


```
{
    "statusCode": 200,
    "message": "success",
    "returnObj": [
        {

            "id": "1751878497877790722",
            "clusterId": "da595eb1d81503b323fdc01d9bf786b7",
            "payType": 1,
            "nodeGroupType": "MASTER",
```

```
            "nodeGroupName": "master",
            "imageId": "f8d28855-ffda-426d-96e2-0f6bf46877e0",
            "hostNum": 3,
            "computeSpecificationId": 101,
            "iaasVmSpecCode": "s7.2xlarge.4",
            "cpuNum": 8,
            "memory": 32,
            "diskSpecificationList":
"[{\"diskFunctionType\":1,\"diskFunctionName\":\"System
Disk\",\"diskType\":\"cloud\",\"diskTypeName\":\"Cloud
Disk\",\"ioType\":\"SSD-genric\",\"ioTypeName\":\"General Purpose
SSD\",\"volume\":200,\"diskNum\":1},{\"diskFunctionType\":2,\"diskFunctionNa
me\":\"Data Disk\",\"diskType\":\"cloud\",\"diskTypeName\":\"Cloud
Disk\",\"ioType\":\"SSD-genric\",\"ioTypeName\":\"General Purpose
SSD\",\"volume\":200,\"diskNum\":2}]",
            "createTime": 1706515356000,
            "updateTime": 1706515356000,
            "mountPublicIp": false,
            "iaasVmSpecId": "b307034d-cbc3-27bb-24f7-a97565814236",
            "highest": null
        },
        {
            "id": "1751878497890373634",
            "clusterId": "da595eb1d81503b323fdc01d9bf786b7",
```

```
            "payType": 1,
            "nodeGroupType": "CORE",
            "nodeGroupName": "core",
            "imageId": "f8d28855-ffda-426d-96e2-0f6bf46877e0",
            "hostNum": 1,
            "computeSpecificationId": 101,
            "iaasVmSpecCode": "s7.2xlarge.4",
            "cpuNum": 8,
            "memory": 32,
            "diskSpecificationList":
"[{\"diskFunctionType\":1,\"diskFunctionName\":\"System
Disk\",\"diskType\":\"cloud\",\"diskTypeName\":\"Cloud
Disk\",\"ioType\":\"SSD-genric\",\"ioTypeName\":\"General Purpose
SSD\",\"volume\":200,\"diskNum\":1},{\"diskFunctionType\":2,\"diskFunctionNa
me\":\"Data Disk\",\"diskType\":\"cloud\",\"diskTypeName\":\"Cloud
Disk\",\"ioType\":\"SSD-genric\",\"ioTypeName\":\"General Purpose
SSD\",\"volume\":200,\"diskNum\":2}]",
            "createTime": 1706515356000,
            "updateTime": 1706515356000,
            "mountPublicIp": false,
            "iaasVmSpecId": "b307034d-cbc3-27bb-24f7-a97565814236",
            "highest": null
        }
    ]
```

```
}
```

Failed request return value example

```
{
    "statusCode": 200,
    "error": "EMR_400021",
    "message": "Invalid cluster ID",
    "returnObj": []
}
```

Status Code

Please refer to status codes

Error Codes

Please refer to error codes

6.2.4. Page Query of Cluster Information

API Introduction

This API provides the function to query cluster information. The system takes query conditions from user input and returns the details of clusters that meet these conditions under the resource pool specified by the user. Information of deleted clusters will not be returned.

API Restrictions

None

Require API Approval or Not

No

URI

POST /v1/emr/openapi/cluster/clusterDetail/selectPage

Content-Type

application/json

Path Parameter

None

Query Parameter

None

Request Header Parameter

None

Request Body Parameter

Request Body Parameter

Parameter	Required or Not	Parameter Type	Description	Example	Lower-level Object
regionId	Yes	String	Resource pool ID	bb9fdb42056f11eda1610242ac110002	/
pageIndex	Yes	Integer	Current page,	1	/

			parameter value must be a positive integer		
pageSize	Yes	Integer	Size of each page, parameter value must be a positive integer	10	/
clusterName	No	String	Cluster Name	cluster_name	/
clusterStateCode	No	Integer	Code for the cluster state (1: Starting, 2: Running, 3: Abnormal Termination, 4: Releasing, 5: Terminated, 10: Frozen)	1	/
clusterTypeCode	No	Integer	Code for the cluster	1	/

			type (1: Data Lake, 2: Data Service, 4: Data Analysis, 5: Elastic Search)		

Response Parameter

Response Parameter

Parameter	Parameter Type	Description	Example	Lower-level Object
statusCode	Integer	Status Code	200	/
message	String	Brief summary of the current API call status and necessary prompt information.	Requested	/
error	String	Error code. This field is not returned when the request is successful.	EMR_400001	/
returnObj	Object	Return result.	/	returnObj

Table returnObj

Parameter	Parameter Type	Description	Example	Lower-level Object

total	Integer	Total number of clusters	23	/
size	Integer	Number of clusters returned in the current query	1	/
current	Integer	Current page	1	/
orders	Array of Objects	Sort field and sort method	/	orderItem
optimizeCountSql	Boolean	Automatic optimization of count SQL query	true	/
searchCount	Boolean	Whether to perform count query, default is true	true	/
countId	String	ID for performing count query	1	/
maxLimit	Integer	Maximum limit on the number of entries per page	20	/
pages	Integer	Total number of pages	5	/
records	Array of Objects	Cluster List	/	records

Table orderItem

Parameter	Parameter Type	Description	Example	Lower-level Object

column	String	The field that is to be sorted	regionId	/
acs	Boolean	Specifies whether the sorting is in ascending order, the default is true	true	/

Table records

Parameter	Parameter Type	Description	Example	Lower-level Object
id	String	Cluster ID	8571aa313aaf77d45bf6765805a9ce55	/
managerClusterId	Integer	Cluster ID as defined by the manager	1	/
iaasType	String	Platform	Public Cloud	/
regionId	String	Resource pool ID	bb9fdb42056f11eda1610242ac110002	/
regionName	String	Resource pool name	East China 1	/
availableZoneId	String	Availability zone ID	cn-huadong1-jsnj1A-public-ctcloud	/
availableZoneName	String	Availability zone name	AZ 1	/

clusterName	String	Cluster Name	test_vpc_24	/
payType	String	Payment Type	Annual and Monthly Subscription	/
clusterType	String	Cluster Type	Elastic Search	/
clusterTypeVersion	String	Product Version	YI-MapReduce-2.12.0	/
clusterPlanCode	String	Code of the cluster plan	cloud-search	/
componentNameList	String	List of component names	[{componentTitle:ElasticSearch,version:7.10.2}, {componentTitle:Kibana,version:7.10.2}]]	/
datasourceConfigs	String	Information about the data source	[]	/
vpcId	String	VPC ID	vpc-0k5xl6w5	/
subnetId	String	Subnet ID	subnet-i2ys8sp	/
securityGroupId	String	Security Group ID	[sg-4h7w9cl1]	/
clusterDueTime	Integer	Expiry time of the cluster	1709193751000	/

userId	String	User ID	ba14c8e729e447d69698f81ac7d55555	/
accountId	String	Account ID	ed24e4b414a048b0a9cb995f59cc85jj	/
clusterCreateTime	Integer	Creation time of the cluster	1706515357000	/
clusterState	String	Cluster Status	Running	/
createTime	Integer	Creation time	1706515356000	/
updateTime	Integer	Update time	1706515356000	/

Enumeration Parameters

None

Request Header Example

None

Request URL Example

```
https://emr-
global.ctapi.ctyun.cn/v1/emr/openapi/cluster/clusterDetail/selectPage
```

Request Body Example

```
{
    "clusterName":"",
```

```
    "clusterStateCode": 2,
    "clusterTypeCode": 2,
    "pageIndex": 1,
    "pageSize": 10,
    "regionId":"bb9fdb42056f11eda1619242ac110002"
}
```

Response Example

Successful request return value example

```
{
    "statusCode": 200,
    "message": "success",
    "returnObj": {
        "records": [
            {
                "id": "2cd01a834b8e0c14d4a03c01d6052924",
                "managerClusterId": 1,
                "iaasType": "Public Cloud",
                "regionId": "bb9fdb42056f11eda1610242ac110002",
                "regionName": "East China 1",
                "availableZoneId": "cn-huadong1-jsnj1A-public-ctcloud",
                "availableZoneName": "Availability Zone 1",
                "clusterName": "sjfw-sxy-0207-3",
                "payType": "Annual or Monthly",
                "clusterType": "Data Service",
```

```
        "clusterTypeVersion": "YI-MapReduce-2.12.1",
        "clusterPlanCode": "data-service",
        "componentNameList":
"[{\"componentTitle\":\"OpenLDAP\",\"version\":\"2.4.50\"},{\"componentTitle\
":\"Kerberos\",\"version\":\"1.18.2\"},{\"componentTitle\":\"ZooKeeper\",\"versi
on\":\"3.7.1\"},{\"componentTitle\":\"HDFS\",\"version\":\"3.3.3\"},{\"componen
tTitle\":\"YARN\",\"version\":\"3.3.3\"},{\"componentTitle\":\"HBase\",\"version\
":\"2.4.12\"},{\"componentTitle\":\"Ranger\",\"version\":\"2.2.0\"}]",
        "datasourceConfigs":
"[{\"compType\":\"Ranger\",\"properties\":[{\"propName\":\"DB_Host\",\"prop
Value\":\"192.168.0.156\"},{\"propName\":\"DB_Port\",\"propValue\":\"13049\"}
,{\"propName\":\"DB_Name\",\"propValue\":\"ranger0207_4_sxy\"},{\"propNa
me\":\"DB_User\",\"propValue\":\"root\"},{\"propName\":\"DB_Password\",\"pr
opValue\":\"t3O15dEAnNxQbKytUnCpi+ir5lbdkt45Zxm9BSO1YoIlowerqD0dCA
22hMVLDwhff5na3fz4KAnNW7fQpP84ZaKgX6XRovjEsWJ0Gyn9ss0NZ3nRI7jjO
W08CnNOmGRirY3OebmYO9PsSucjVMGCPAplyhnjMefhGJHOFnvw4xU=\"}]}]",
        "vpcId": "vpc-fr2xjo1gj0",
        "subnetId": "subnet-vkixrw8xw7",
        "securityGroupId": "[\"sg-72kksosxfn\"]",
        "clusterDueTime": 1709795918000,
        "userId": "ba14c8e729e447d69698f81ac7d506b5",
        "accountId": "ed24e4b414a048b0a9cb995f59cc85dd",
        "clusterCreateTime": 1707290310000,
        "clusterState": "Running",
```

```
            "createTime": 1707290309000,
            "updateTime": 1707291860000
        }
    ],
    "total": 1,
    "size": 10,
    "current": 1,
    "orders": [],
    "optimizeCountSql": true,
    "searchCount": true,
    "countId": "",
    "maxLimit": null,
    "pages": 1
  }
}
```


Failed request return value example


```
{
  "statusCode": 200,
  "error": "EMR_400010",
  "message": "pageIndex, regionId are required, not provided",
  "returnObj": []
}
```

Status Code

Please refer to status codes

Error Codes

Please refer to error codes

6.3. API

API

7. FAQs

7.1. Product

7.1.1. What Types of Distributed Storage Does YI-MapReduce Support?

When using YI-MapReduce, customers can choose to include the Doris component. Doris provides distributed data storage capabilities, with version 2.0.4 being the default offering.

For more details on the components of the YI-MapReduce product, please consult the Version Overview.

7.1.2. Does YI-MapReduce Support Modification of the CIDR Block?

Currently, YI-MapReduce does not support modification of the CIDR block for nodes within a cluster. Therefore, we advise customers to partition the CIDR block based on business needs before activating the YI-MapReduce cluster. Also, ensure the availability of suitable resources in other eSurfing Cloud products, such as the Virtual Private Cloud (CT-VPC).

7.1.3. Can I Downgrade the Configuration Specifications of Nodes Within a YI-MapReduce clusters?

Currently, YI-MapReduce does not support the downgrade of configuration specifications for nodes within a cluster. Customers can upgrade the configuration

specifications for Master and Core node groups within the YI-MapReduce cluster to meet business needs.

7.1.4. Can I Clean up the Log Files in the /var/log directory on nodes in the YI-MapReduce cluster?

The /var/log directory on nodes in the YI-MapReduce cluster contains essential log data related to operations of the operating system, security, and application and audit logs of instances of various components deployed in the YI-MapReduce cluster. This information can assist customers in identifying issues with the cloud host system and provide security auditing for the YI-MapReduce component cluster.

- It is not recommended to clean up the system log files mentioned above. This is a high-risk operation that may lead to the failure or unavailability of the Elastic Cloud Server.
- For the application and audit logs of the YI-MapReduce cluster mentioned above, you should consider the need for cleanup. If cleanup is necessary, you should prioritize deleting older historical logs.

7.1.5. Can I Modify the IP Address of the Nodes in the YI-MapReduce cluster?

Currently, modifying the IP address of nodes within the YI-MapReduce cluster is not possible. Forcing a modification of the IP could result in communication abnormalities among the YI-MapReduce cluster services and could lead to a total business shutdown of the YI-MapReduce cluster. We recommend planning the availability zones and VPC in advance before activating the YI-MapReduce cluster.

7.1.6. What Are the Differences and Connections Between the YI-MapReduce console and the YI-MapReduce Manager?

Users can access the YI-MapReduce Manager page at the YI-MapReduce Manager tab on the YI-MapReduce management console. The differences and connections between the management console and the Manager page are outlined in the table below:

Common Operations	YI-MapReduce Console	YI-MapReduce Manager
View basic cluster information, IAM synchronization	Support	Not supported
Configuration upgrades, remote connections, binding/unbinding Elastic IPs	Support	Not supported
Operation Logs	Support	Support
User Management	Support	Not supported
Access Link and Port	Support	Not supported
Start service, stop service, roll-restart service	Not supported	Support
Start, stop, restart, roll-stop, roll-restart instances	Not supported	Support
View monitoring information and alarm messages	Not supported	Support
O&M and configuration management	Not supported	Support

7.2. Billing

7.2.1. What billing methods does YI-MapReduce support?

Currently, YI-MapReduce supports a subscription-based billing mode.

7.2.2. Is there a low balance reminder for YI-MapReduce accounts?

Users can set a balance threshold in the User Center. An SMS will be sent when the balance drops below this threshold.

7.2.3. Are there any additional charges for purchasing YI-MapReduce?

When creating an YI-MapReduce cluster, you will be charged separately for the vCPU, memory, storage, and bandwidth resources you select based on actual requirements, as well as the corresponding software service fees. The system will automatically calculate the price based on your selection.

7.2.4. How to handle exceptions when unsubscribing from Elastic Cloud Servers in ECS after unsubscribing from YI-MapReduce service?

Please collect the IDs of the Elastic Cloud Servers and contact your account manager or submit a service ticket. Our technical support team will assist you in resolving the issue.

7.2.5. Why isn't the price displayed when creating a YI-MapReduce clusters?

If you only select the number of disks and do not choose the number of virtual machine instances for the cluster, no charges will be incurred for the product quotation.

Once you select the number of instances, the product quotation will be displayed in real time at the bottom of the browser.

7.3. Purchase

7.3.1. In which resource pools of eSurfing Cloud can I subscribe to YI-MapReduce?

Currently, YI-MapReduce is available in the China(Hongkong2) resource pool. We will continuously add more resource pools in the future.

7.3.2. What types of distributed storage does YI-MapReduce support?

YI-MapReduce provides mainstream Hadoop support, currently offering Hadoop version 3.3.3, and will continue to update in line with community developments.

7.3.3. What big data cluster business scenarios does YI-MapReduce support for purchase?

YI-MapReduce mainly encompasses four business scenarios: Data Analysis and Elastic Search.

Scenario	Description

Data Analysis	OLAP: Data Analysis Apache Doris: An open-source MPP architecture OLAP analysis engine that supports sub-second data queries and multi-table joins.
Elastic Search	ES: Elastic search It delivers low-cost, high-performance, and reliable search and analysis service capabilities for structured/unstructured data.

7.3.4. What billing items are included in the purchase of the YI-MapReduce product?

The cost of the YI-MapReduce cluster consists of the YI-MapReduce service management fee and ECS product fees. The ECS product fees include the fees for Elastic Cloud Server and Elastic Volume Service.

7.4. Account and Password

7.4.1. What is the default user for logging into the YI-MapReduce Manager after activating a YI-MapReduce clusters?

Upon activating a YI-MapReduce clusters, the system defaults to the "op_admin" account for login into the YI-MapReduce Manager. Users can access the YI-MapReduce Manager directly through password-free login from the YI-MapReduce console.

7.4.2. How can I change and apply the passwords for the related big data components after activating a YI-MapReduce clusters?

1. You can go to the YI-MapReduce console and select a YI-MapReduce clusters from **My Cluster** to adjust the account password.
2. After selecting the relevant cluster, click on the **YI-MapReduce Manager** tab of the cluster to access the YI-MapReduce Manager control platform through password-free login.
3. In the **O&M and Configuration - Configuration Management** function menu, switch between different component services such as HDFS, ZooKeeper, etc. Then, identify the passwords to be adjusted, complete the configuration changes, and click **Save Changes** at the bottom right.
4. Click the **Configuration Sync** button. Follow the instructions to complete the distribution of the new configuration.
5. Go to the **Cluster Service** function menu. Locate the corresponding component service and go to its details page. From the **O&M Operations** dropdown menu, execute the cluster restart at the top-right to apply the cluster configuration changes.

7.4.3. How to change the password for the root account used to log in to a node after activating the YI-MapReduce cluster?

1. You can go to the YI-MapReduce console and select a YI-MapReduce clusters from **My Cluster** to adjust the account password used to log in to a node.
2. After selecting the relevant cluster, go to **Node Management** of the cluster, click **Remote Connection** on the right side of the specific node, and access the server system by entering the account and password required in the VNC console.
3. In the node command line, type "passwd root" to complete the password change for the root account.

7.5. Cluster Creation

7.5.1. How to Create an YI-MapReduce Cluster Using a Custom Security Group?

1. Currently, only the China(Hongkong2) resource pool supports first-class nodes for the YI-MapReduce product. Future updates will be based on actual product rollouts.
2. Before creating a YI-MapReduce clusters, you need to establish a Virtual Private Cloud (VPC) and add the related security group rules in the first-class node resource pools such as China(Hongkong2) .
3. Refer to the Create a VPC section in the eSurfing Cloud - VPC product documentation. Create related VPC networks in the first-class node resource pool. By default, the CIDR block is "192.168.0.0/16".
4. After the above steps, you can find a default security group automatically added to the VPC created in the Control Center - Network Console - Access Control - Security Group menu. In addition, you can create custom security groups.
5. Before creating a YI-MapReduce clusters, refer to the Adding a Security Group Rule section in the eSurfing Cloud - Virtual Private Cloud product documentation. Add the following two inbound rules:
    a. Allow remote 198.19.128.0/20 with Any protocol, and port range is 1-65535:
        i. Authorization Policy: Choose Allow.
        ii. Priority: Enter any number between 1-99.
        iii. Protocol: Choose Any.
        iv. Source Address: Choose IP, and enter 198.19.128.0/20.
        v. Description: Fill in as needed.

b. Allow remote 192.168.0.0/24 (**actual CIDR block address, depending on the VPC subnet CIDR block address created by the user**) with Any protocol, and port range is Any:

i. Priority: Enter any number between 1-99.

ii. Authorization Policy: Choose Allow.

iii. Protocol: Choose Any.

iv. Source Address: Choose IP, enter 192.168.0.0/24 (**actual CIDR block address, depending on the VPC subnet CIDR block address created by the user**).

v. Description: Fill in as needed.

Description

After creating the VPC, the default security group automatically contains Rule a mentioned above, and the user needs to add Rule b. If the user uses a custom security group, ensure both Rule a and Rule b have been added.

7.5.2. How to change the login method for nodes within the YI-MapReduce cluster?

Changing is not available. The cluster login method chosen during the creation of the cluster cannot be altered later.

7.6. Cluster Management

7.6.1. FAQ About Using

How can I view the configuration information of the cluster?

- Once a cluster is created, click on the cluster name in the YI-MapReduce console to access the basic information page. Here, you can see the basic configuration details of the cluster, including the cluster name, creation time, payment type, cluster version, component information, VPC and security group, node instance specifications and storage, etc. The specifications and storage of the node instance determine the data analysis and processing capabilities of the cluster. The higher the node instance specifications and the more storage, the faster the cluster runs, the stronger its analysis and processing capabilities, and the higher the associated cost.
- Click to go to **YI-MapReduce Manager** and proceed to the **O&M and Configuration** page under **Configuration Management**. By selecting different cluster services, you can view the full configuration list information for the corresponding cluster service.

Is it necessary to shut down the Master node for an upgrade?

When upgrading the MASTER/CORE node group of the YI-MapReduce service cluster, an automatic shutdown operation will be triggered. Once the upgrade is complete, an automatic boot operation will be triggered, eliminating the need for manual intervention. However, please note that all cluster services must be stopped before the upgrade, and all cluster services should be manually restarted after the upgrade is complete.

Can installed components in the YI-MapReduce cluster be removed?

Components within an existing YI-MapReduce-2.14.1 version cluster cannot be removed. If not in use, you can log in to the cluster service page of the Manager page, locate the corresponding component, and stop it.

How do I find the path of each component configuration file?

After the cluster creation, click on the cluster name in the YI-MapReduce console to view the details. Click on **YI-MapReduce Manager** to enter the big data O&M management platform. Click on **O&M and Configuration**. Enter the **Configuration Management** page. Then, click on the Unfold button in the upper-right corner to view the configuration file paths for the relevant cluster services.

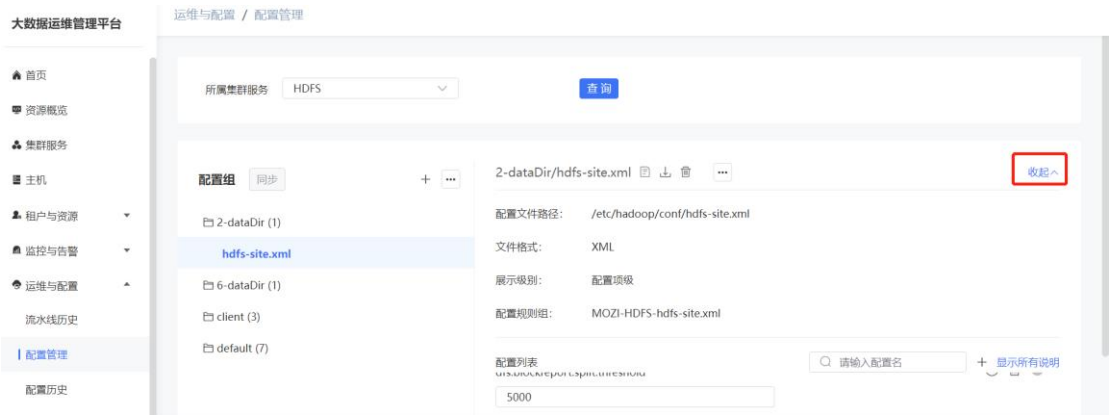

How do I start/stop all cluster services?

- Log in to the YI-MapReduce management console, click on **My Cluster**, click on the name of the specified cluster, and enter the cluster information page.
- Click on the **YI-MapReduce Manager** tab, click on the **Cluster Service** menu, and enter the cluster service list page. Click on **More Operations > Start All Cluster Services** or **Stop All Cluster Services**. After confirming the operation, start/stop all cluster services and wait for completion.

How to perform a rolling restart of cluster services?

- Log in to the YI-MapReduce management console, click on **My Cluster**, click on the name of the specified cluster, and enter the cluster information page.
- Click on the **YI-MapReduce Manager** tab, click on the **Cluster Service** menu, and enter the cluster service list page. Select the specific cluster service to roll-restart. Click on the cluster service name to enter the cluster service details page. Click on **Operations > Roll-restart Cluster Service** to access the page for the rolling restart of the cluster service. Follow the instructions for the next steps and wait for the rolling restart of the cluster service to complete.

How to view the cluster service monitoring dashboard?

- Log in to the YI-MapReduce management console, click on **My Cluster**, click on the name of the specified cluster, and enter the cluster information page.
- Click the **YI-MapReduce Manager** tab, select the **Cluster Service** menu, and select the cluster service you want to view. On the details page of the cluster service, click on the **Cluster Dashboard** tab to view the monitoring dashboard for that cluster service.

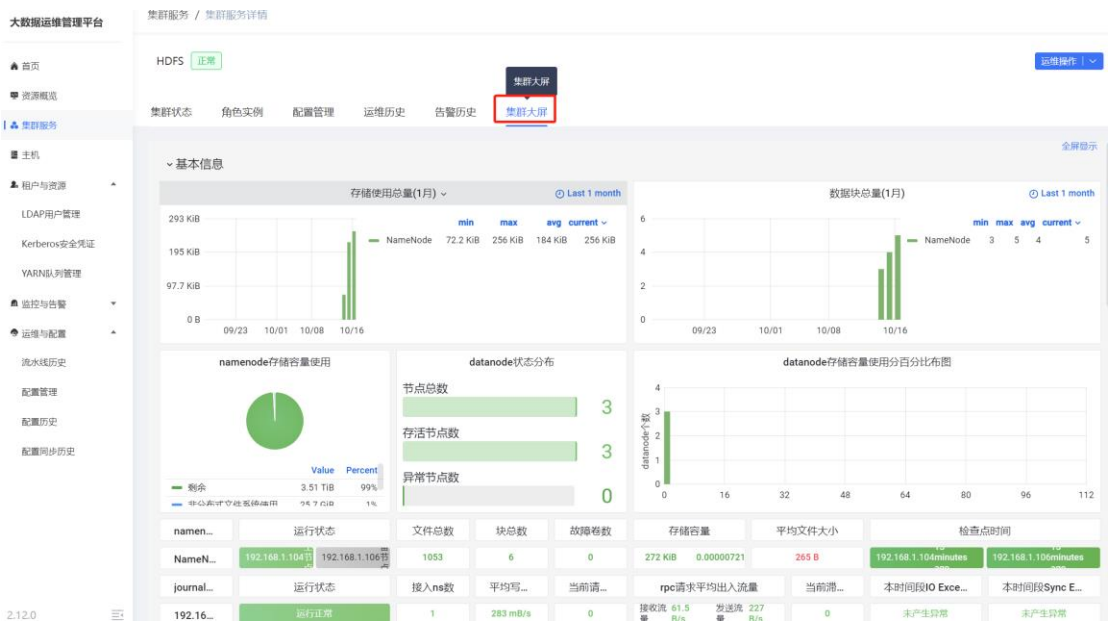

7.6.2. How to view all clusters?

All clusters in YI-MapReduce are displayed on the My Cluster page. This page provides a comprehensive view of all cluster information. When the number of clusters is high, they can be paginated. Clusters of any status are visible. Highlighted cluster names signify clusters operating as expected, whereas non-highlighted names denote clusters operating abnormally. The possible cluster states include:

- Starting: The deployment of the cluster is in progress, with its specific process visible.
- Running: The cluster has successfully completed deployment or renewal, and is operating as expected.
- Abnormally Terminated: The cluster has either failed to start or has been terminated due to unexpected reasons. The reasons for failure can be viewed.
- Releasing: Clusters that are in the process of being unsubscribed display a "Releasing" status.
- Terminated: Clusters that have expired for more than 15 days or where users have completed the unsubscribe process display a "Terminated" status and will be removed.
- Renewing: When a cluster is in the midst of renewal, its status displays as "Renewing".
- Frozen: Clusters that are expired but not over 15 days are kept in a "Frozen" state and will automatically unfreeze after renewal.

7.6.3. How to view host and component logs?

Procedure

1. Log in to the YI-MapReduce management console.
2. Select **My Cluster**, choose a running cluster and click on its name to enter the cluster information page.
3. Under the **Node Management** tab, click on the master node and select the master node to access.

4. Click on **Remote Connection** on the right side of the node.
5. As prompted, enter the username and password for the master node. The username is "root". The password is the one set when creating the cluster.
6. Go to the log directory of the corresponding component to view the related logs.

7.6.4. How to optimize cluster service configuration?

Procedure

Step 1: Modify configuration

1. Log in to the YI-MapReduce management console, click on **My Cluster**, click on the name of the specified cluster, and enter the cluster information page.
2. Click the **YI-MapReduce Manager** tab, and then select **Go to YI-MapReduce Manager**.
3. Upon entering the YI-MapReduce Manager, click **Cluster Service** on the menu to go to the cluster service list.
4. Choose the specific cluster service. Click on the cluster service name to enter the cluster service details page.
5. Click on the **Configuration Management** tab, select the configuration group, and choose the configuration name you want to modify. The right side of the page will display the detailed information of this configuration. Make the necessary modifications and click **Save Changes** when complete.

Step 2: Synchronize configuration

1. Click on the **Sync** button to enter the page for pending configuration synchronization.
2. Click on the **Configuration Sync** button to execute the configuration synchronization and access the configuration sync details page, view the synchronization progress, and wait for the sync to complete.

Step 3: Configuration takes effect

1. After the configuration synchronization is complete, return to the cluster service details page.
2. Click on **O&M Operations > Roll-restart Cluster**. The modified configurations take effect after the cluster service restarts.

7.7. Performance Optimization

7.7.1. Does the YI-MapReduce Cluster Support System Reinstallation?

the YI-MapReduce cluster does not support system reinstallation. And we advise against clients remotely accessing CT-ECS to perform high-risk operations such as system reinstallation. These actions may cause faults in the YI-MapReduce cluster, leading to service unavailability.

7.7.2. Does the YI-MapReduce Cluster Support System Switching?

the YI-MapReduce cluster does not support operating system switching on its nodes. If customers are experiencing issues with the YI-MapReduce cluster, we recommend reaching out to the professional customer service of eSurfing Cloud by submitting a work ticket for consultation and communication. We appreciate your understanding.

7.7.3. How to Shut Down the Firewall Service?

1. Access each node of the cluster as the root user.
2. Observe if the firewall service is running.
   On ECS nodes, execute the command systemctl status firewalld.service.

3. Turn off the firewall service.
   On ECS nodes, execute the command systemctl stop firewalld.service.

8. Download

8.1. YI-MapReduce Service User Guide

Click to download: YI-MapReduce Service User Guide.pdf

9. Related Agreements

9.1. YI-MapReduce Service Agreement

Click to preview: YI-MapReduce Service Agreement.ghkvbn