# Alibaba Cloud

E-MapReduce Cluster Management

Document Version: 20220517

C-J Alibaba Cloud

### Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

- 1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloudauthorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
- 2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.
- 3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
- 4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).
- 5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud and/or its affiliates Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.
- 6. Please directly contact Alibaba Cloud for any errors of this document.

### **Document conventions**

Style	Description	Example
A Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	Danger: Resetting will result in the loss of user configuration data.
O Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	Warning: Restarting will cause business interruption. About 10 minutes are required to restart an instance.
C) Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	Notice: If the weight is set to 0, the server no longer receives new requests.
⑦ Note	A note indicates supplemental instructions, best practices, tips, and other content.	Note: You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click Settings> Network> Set network type.
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click OK.
Courier font	Courier font is used for commands	Run the cd /d C:/window command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	bae log listinstanceid Instance_ID
[] or [a b]	This format is used for an optional value, where only one item can be selected.	ipconfig [-all -t]
{} or {a b}	This format is used for a required value, where only one item can be selected.	switch {active stand}

### Table of Contents

1.Cluster types	07
2.Cluster planning	11
2.1. Select configurations	11
2.2. Node categories	12
2.3. Gateway clusters	13
2.4. ECS instances	14
2.5. Storage	15
2.6. Supported instances with local disks	16
2.7. Disaster recovery in EMR clusters	17
3.Configure clusters	19
3.1. Create a cluster	19
3.2. Create a gateway cluster	27
3.3. Manage permissions	29
3.3.1. Assign roles	29
3.3.2. EMR service role	32
3.3.3. ECS application role (used in EMR V3.32.0 and earlier	33
3.3.4. ECS application role (used in V3.X.X versions later tha	36
3.3.5. Use a custom ECS application role to access other clou	36
3.3.6. Grant permissions to RAM users	41
3.4. Manage user accounts	42
3.5. Manage and use tags	45
3.6. Use resource groups	47
3.7. Manage security groups	50
3.8. Manage Metadata	52
3.8.1. Configure DLF as a metadatabase	52
3.8.2. Configure an independent ApsaraDB RDS for MySQL d	53

3.9. Manage bootstrap actions 5	8
3.10. Cluster scripts 6	2
3.11. Manage cluster resources 6	4
3.11.1. Cluster resource overview 6	4
3.11.2. Use Capacity Scheduler 6	4
3.11.3. Use Fair Scheduler 6	6
3.12. Cluster management page 6	9
3.13. View the cluster list and cluster details 7	0
4.Manage services 7	5
4.1. View the status of services7	5
4.2. Add a service7	5
4.3. Restart a service 7	6
4.4. Manage parameters for services7	6
4.5. Customize software configurations 8	0
4.6. View the deployment information of service components 8	1
4.7. Access the Web UI 8	9
4.7.1. Create an SSH tunnel to access web UIs of open sourc	9
4.7.2. Access the web UIs of open source components9	7
5.Cluster Operations 10	1
5.1. Common file paths 10	1
5.2. Log on to a cluster 10	2
5.3. Scale out a cluster 10	6
5.4. Scale in a cluster 10	8
5.5. Auto Scaling 10	9
5.5.1. Overview 10	9
5.5.2. Create an auto scaling machine group	9
5.5.3. Manage auto scaling 11	0
5.5.4. Configure auto scaling rules by time	4

	5.5.5. View auto scaling history	116
	5.5.6. Configure auto scaling monitoring alerts	117
	5.6. Expand disk capacity	118
	5.7. Add node groups	118
	5.8. Remove abnormal ECS instances	119
	5.9. Upgrade node configurations	119
	5.10. Status list	123
	5.11. Cluster O&M guide	124
	5.12. Release a cluster	130
6.	.FAQ about cluster management	132

### 1.Cluster types

This topic describes the types of clusters that are supported by E-MapReduce (EMR) and the important operations that you can perform in the cluster of each type.

### Overview

Cluster type	Description	Important operation
Hadoop	<ul> <li>Provides Hadoop, Hive, and Spark components that serve as semi-hosted services and are used to store and compute large-scale distributed data offline.</li> <li>Provides Presto and Impala components for interactive queries.</li> <li>Provides other Hadoop ecosystem components, such as Oozie.</li> </ul>	<ul> <li>Create a cluster</li> <li>Log on to a cluster</li> <li>Release a cluster</li> </ul>
Data Science	Data Science clusters are commonly used in big data and AI scenarios. Data Science clusters support the offline extract, transform, load (ETL) of big data based on Hive and Spark, and TensorFlow model training. You can choose the CPU+ GPU heterogeneous computing framework and deep learning algorithms supported by NVIDIA GPUs to run computing jobs more efficiently.	<ul> <li>Overview</li> <li>Create a cluster</li> <li>View the status of services</li> <li>Use ds-controller</li> <li>Faiss-Server</li> <li>GKS</li> <li>Kubeflow <ul> <li>Use Kubeflow for model training</li> <li>Use Seldon Core to implement online services</li> <li>Use the EasyRec algorithm library to build a pipeline</li> </ul> </li> <li>Tutorials <ul> <li>Distributed inference solution</li> <li>XGBoost-based distributed training</li> </ul> </li> </ul>
Dataflow	Dataflow clusters provide an end-to-end (E2E) real- time computing solution. The clusters incorporate Kafka, a distributed message system with high throughput and scalability, and the commercial Flink kernel provided by Apache Flink-powered Ververica. The clusters are used to resolve various E2E real-time computing issues and are widely used in real-time data ETL, and log collection and analysis scenarios. You are free to use one of the two components or both.	<ul><li>Flink</li><li>Kafka</li></ul>

Cluster type	Description	Important operation
Druid	Druid clusters provide a semi-hosted, real-time, and interactive analytic service. These clusters can query big data within milliseconds and ingest data in multiple ways. You can use Druid clusters with services such as EMR Hadoop, EMR Spark, Object Storage Service (OSS), and ApsaraDB RDS to build a flexible and stable system for real-time queries.	<ul> <li>Create a cluster</li> <li>Log on to a cluster</li> <li>Release a cluster</li> </ul>

### ☉ Hadoop Cluster ☉ Data Science Cluster ☉ Druid Cluster ☉ Dataflow Cluster

Component	References
HDFS	Overview
YARN	Overview
Hive	Overview
Spark	Overview
Knox	Overview
Tez	Overview
Sqoop	Overview
Smart Dat a	Overview
OpenLDAP	Overview
Hudi	Overview
Hue	Overview
HBase	Overview
ZooKeeper	Overview
Presto	Overview
Impala	Overview
Zeppelin	Overview
Flume	Overview
Livy	Overview
Ranger	Overview
Phoenix	Overview
ESS	Overview

Alluxio	Overview
Kudu	Overview
Oozie	Overview
Component	References
Faiss-Server	Faiss-Server
GKS	GKS
Kubeflow	Use Kubeflow for model training Use Seldon Core to implement online services Use the EasyRec algorithm library to build a pipeline

Component	References
Druid	Overview
Superset	Overview
ZooKeeper	Overview
Knox	Overview
OpenLDAP	Overview

Cluster mode	Component	References
	HDFS	Overview
	YARN	Overview
Flink	ZooKeeper	Overview
FULK	Knox	Overview
	Flink	Overview
	OpenLDAP	Overview
	ZooKeeper	Overview
	Ganglia	None
	Kafka	Overview
	Kafka-Manager	Overview
Kafka	OpenLDAP	Overview

Кпох	Overview
Ranger	Overview
Hue	Overview

### 2.Cluster planning 2.1. Select configurations

To use E-MapReduce (EMR), you must select appropriate clusters. When you select configurations for EMR clusters, consider the use of big data in your enterprise and your financial budget, and estimate the amount of data and the reliability of services that you require.

### Big data scenarios

EMR applies to the following enterprise big data scenarios:

• Batch operations

This scenario requires high disk throughput and high network throughput, involves a large amount of data, but has low requirements on the timeliness of data processing. You can use MapReduce, Pig, or Spark for this scenario. This scenario does not require a large memory capacity. You must focus on the requirements of jobs on vCPUs and memory, and the network requirements when you perform data shuffling.

• Ad hoc queries

Data scientists and data analysts use ad hoc query tools to retrieve data. This scenario requires realtime queries, high disk throughput, and high network throughput. You can use Impala or Presto for this scenario. This scenario requires a large memory capacity. You must consider the data volume and number of concurrent queries your business needs to handle.

• Stream computing, high network throughput, and compute-intensive scenarios

You can use Flink, Spark Streaming, or Storm for these scenarios.

• Message queues

In this scenario, high disk throughput and high network throughput are required, the memory usage is high, and data storage does not depend on Hadoop Distributed File System (HDFS). You can use Kafka.

• Cold backup

This scenario does not require high disk or high computing throughput but requires low cost. We recommend that you use JindoFS to store cold data to Object Storage Service (OSS) Archive or Cold Archive storage to reduce storage costs.

#### **EMR** nodes

An EMR cluster consists of three types of nodes: master, core, and task nodes. For more information, see Node categories.

You can select ultra disks, local disks, standard SSDs, or local SSDs for EMR storage. These disks are ranked in descending order of performance: local SSDs > standard SSDs > local disks > ultra disks.

EMR underlying storage supports OSS (OSS Standard storage only) and HDFS. OSS has a higher data availability than HDFS. The data availability of OSS is 99.9999999%, while the data availability of HDFS depends on the reliability of cloud disk or local disk storage. Before you use EMR to compute data, you must first restore the data that is stored to OSS Archive or Cold Archive storage, and then store the data to OSS Standard storage.

Storage prices:

• Local disk storage: USD 0.003/GB/month

- OSS Standard storage: USD 0.02/GB/month
- OSS Archive storage: USD 0.045/GB/month
- OSS Cold Archive storage: USD 0.002/GB/month
- Ultra disk storage: USD 0.05/GB/month
- Standard SSD storage: USD 0.143/GB/month

### Select configurations for EMR

- Select master node configurations.
  - Master nodes are used to deploy the master processes of Hadoop, such as NameNode and ResourceManager.
  - EMR components such as HDFS, YARN, Hive, and HBase use the high availability architecture. We recommend that you enable high availability for production clusters in the **Hardware Settings** step. If high availability is not enabled when you create an EMR cluster, it cannot be enabled later.
  - Master nodes are used to store HDFS metadata and component log files. These nodes are computeintensive with low disk I/O requirements. HDFS metadata is stored in memory. The minimum recommended memory size is 16 GB based on the number of files.
- Select core node configurations.
  - Core nodes are used to store data, run computing tasks, and run processes such as DataNode and NodeManager.
  - If the volume of data stored in HDFS exceeds 60 TB, we recommend that you use instances with local disks, such as ECS d1 and ECS d1ne instances. The local disk capacity is calculated by using the following formula: (Number of vCPUs/2) × 5.5 TB × Number of instances.

For example, if you purchase four d1 instances with eight vCPUs, the local disk capacity is 88 TB, which is calculated by using the following formula:  $8/2 \times 5.5 \times 4$ . HDFS requires three replicas. Therefore, you must purchase at least three instances that use local disks. To ensure data reliability and disaster recovery, we recommend that you purchase at least four instances.

- If the volume of data stored in HDFS is less than 60 TB, you can use ultra disks or standard SSDs.
- Select task node configurations.

Task nodes are used if the computing capabilities of vCPUs and the memory of core nodes are insufficient. Task nodes do not store data or run DataNode. You can estimate the number of task nodes based on your vCPU and memory requirements.

### EMR lifecycle

EMR supports auto scaling. You can scale out a cluster by following the instructions provided in Scale out a cluster or scale ECS instances up or down by following the instructions provided in Overview of instance upgrade and downgrade.

### Select a zone

We recommend that you deploy EMR and your business system in the same zone of the same region to ensure high efficiency. For more information, see Regions and zones.

### 2.2. Node categories

An E-MapReduce (EMR) cluster consists of three categories of nodes: master, core, and task.

The three categories of nodes run different processes to complete different tasks.

- Master node: runs the NameNode process of Hadoop HDFS and the ResourceManager process of Hadoop YARN.
- Core node: runs the DataNode process of Hadoop HDFS and the NodeManager process of Hadoop YARN.
- Task node: runs the NodeManager process of Hadoop YARN and performs only computing.

Before you create a cluster, you must determine the specifications of Elastic Compute Service (ECS) instances for each node category. Instances of the same category are in the same group. After you create a cluster, you can add ECS instances to the core node group or task node group to scale out the cluster.

**Note** In EMR V3.2.0 and later, clusters support task nodes.

#### Master node

A master node is deployed with the management components of cluster services, such as ResourceManager of Hadoop YARN.

You can access web UIs to view the running status of services in a cluster. To test or run a job in a cluster, you can connect to the master node and submit the job on the command line. For more information about how to connect to a master node, see Log on to a cluster.

#### Core node

Core nodes in a cluster are managed by the master node. Core nodes run the DataNode process of Hadoop HDFS to store all data of a cluster. They also run computing service processes such as NodeManager of Hadoop YARN to run computing tasks.

To cope with the increase of data storage and computing workloads, you can scale out core nodes at any time without affecting the running of the cluster. Core nodes can use different storage media to store data. For more information, see Local disks and Block Storage overview.

### Task node

Task nodes run only computing tasks. They cannot be used to store HDFS data. If the core nodes of a cluster offer sufficient computing capabilities, task nodes are not required. If the computing capabilities of the core nodes in a cluster become insufficient, you can add task nodes to the cluster at any time. You can run Hadoop MapReduce tasks and Spark executors on these task nodes to provide extra computing capabilities.

Task nodes can be added to or removed from a cluster at any time without any impact on the running of a cluster.

### 2.3. Gateway clusters

When you create a gateway cluster, you must associate it with an existing cluster. Gateway clusters can serve as a separate job submission node to facilitate operations on the associated cluster.

A gateway cluster is an independent cluster that consists of multiple nodes with the same configurations. Clients, such as Hadoop (HDFS+YARN), Hive, Spark, and Sqoop clients, are deployed on the cluster.

If no gateway cluster is created, jobs of a Hadoop cluster are submitted on the master node or a core node of the Hadoop cluster, which consumes the resources of this cluster. After a gateway cluster is created, you can use it to submit jobs of the cluster associated with this gateway cluster. This way, the jobs do not occupy the resources of the associated cluster, and the stability of the master and core nodes, especially the master node, in the associated cluster is improved.

Each gateway cluster can have an independent configuration environment. For example, you can create multiple gateway clusters for one cluster that is shared by multiple departments to meet their different business requirements.

### 2.4. ECS instances

This topic describes the Elastic Compute Service (ECS) instance families supported by E-MapReduce (EMR) and their use scenarios.

### ECS instance families supported by EMR

• General purpose

This instance family uses cloud disks for storage. The ratio of vCPUs to memory is 1:4, for example, 8 vCPUs and 32 GiB of memory.

• Compute optimized

This instance family uses cloud disks for storage and provides more computing resources. The ratio of vCPUs to memory is 1:2, for example, 8 vCPUs and 16 GiB of memory.

• Memory optimized

This instance family uses cloud disks for storage and provides more memory resources. The ratio of vCPUs to memory is 1:8, for example, 8 vCPUs and 64 GiB of memory.

• Big data

This instance family uses local SATA disks for storage, which is highly cost-effective. If you want to store large volumes of data (terabytes), we recommend that you use this instance family.

⑦ Note Core nodes can be created only in Hadoop, Data Science, Dataflow, and Druid clusters.

• Local SSD type

This instance family uses local SSDs for storage, which provide high random IOPS and high throughput.

• Shared type (entry level)

Instances in this instance family share CPUs, so they are not stable in scenarios that require large volumes of computing. This instance family is suitable for entry-level users, but not enterprise customers.

• GPU

This instance family is a heterogeneous GPU-based model and applies to scenarios such as machine learning.

### Use scenarios of instance families

• Master nodes

Instances in general-purpose and memory-optimized instance families can serve as master nodes for EMR. They are suitable for scenarios in which data is stored on cloud disks to ensure high data reliability.

- Core nodes
  - Instances in general-purpose, compute-optimized, and memory-optimized instance families can serve as core nodes for EMR. They are suitable for small volumes of data (below terabytes) and scenarios in which OSS is used as primary data storage.
  - If the volume of data is 10 terabytes or more, we recommend that you use the big data type because it is more cost-effective.

Notice If local disks are used, HDFS data is stored in local disks, which cannot ensure data reliability.

• Task nodes

All instance families except big data families can be used for task nodes to improve the computing capabilities of a cluster.

### 2.5. Storage

This topic describes data storage in E-MapReduce (EMR) clusters, including the supported disk roles and disk types, and Object Storage Service (OSS).

### Background information

For more information about storage types, storage performance, and limits on storage, see Elastic Block Storage devices.

### Disk roles

Each node in an EMR cluster has two disk roles: system disk and data disk. The disks can vary in terms of configuration, type, and capacity.

Disk role	Description
System disk	A disk on which the operating system is installed. By default, the nodes of an EMR cluster use enhanced SSDs as system disks. By default, each node has one system disk.
Data disk	A disk that is used to store data. By default, the master node of an EMR cluster uses one cloud disk as data disk, and each core node of an EMR cluster uses four cloud disks as data disks.

### Cloud disks and local disks

EMR clusters allow you to use cloud disks and local disks to store data. The following table describes the different types of disks.

Disk ty	rpe D	Description	Scenario
---------	-------	-------------	----------

Disk type	Description	Scenario	
Cloud disk	Cloud disks include standard SSDs, ultra disks, and enhanced SSDs. Cloud disks are not directly attached to local compute nodes. Instead, these disks access a remote storage node over the network. Each piece of data has two real-time replicas at the backend. If the data is corrupted due to disk damage, EMR automatically uses a replica to	Cloud disks have lower IOPS and throughput than local disks. If the volume of your business data is below the terabyte level, we recommend that you use cloud disks. ?       Note       If the throughput of cloud disks is insufficient,	
	restore data. Local disks include local SATA disks and local SSDs of the big data type.	you can create a new cluster and use local disks.	
Local disk	Local disks are directly attached to compute nodes and have better performance than cloud disks. You cannot change the number of local disks. No data backup mechanism is deployed at the backend, and upper-layer software is required to ensure data reliability.	The data reliability of local disks is ensured by EMR. If the volume of your business data is at the terabyte level or higher, we recommend that you use local disks.	

When nodes in an EMR cluster are released, data on all the cloud disks and local disks is cleared. The disks cannot be kept independently and used again. Hadoop HDFS uses all data disks for data storage. Hadoop YARN uses all data disks as temporary storage for computing.

### 0 S S

OSS can be used as HDFS in an EMR cluster. You can read data from or write data to OSS by modifying the code that is originally used to access HDFS. Examples:

• Read data from HDFS:

sc.textfile("hdfs://user/path")

• Change the storage type from HDFS to OSS:

sc.textfile("oss://user/path")

• In MapReduce or Hive jobs, you can run HDFS commands to manage data in OSS. Example:

```
hadoop fs -ls oss://bucket/path
hadoop fs -cp hdfs://user/path oss://bucket/path
```

When you run the commands, you do not need to enter your AccessKey pair or the endpoint of OSS. EMR completes the information by using the data of the cluster owner. However, OSS is not suitable for scenarios that require high IOPS, such as Spark Streaming or HBase scenarios.

### 2.6. Supported instances with local disks

To meet the storage requirements in big data scenarios, Alibaba Cloud has launched d1 series instances with local disks.

### d1 series

d1 series instances use local disks, instead of cloud disks, for data storage. This avoids the high costs caused by multiple copies of data generated when cloud disks are used. d1 series instances do not require all data to be transmitted over the network. This improves disk throughput and leverages the advantages of Hadoop in edge computing. Local disks have better storage performance and lower storage unit price than cloud disks. Local disks are as cost-effective as physical hosts.

However, local disks cannot ensure data reliability. Alibaba Cloud provides a multi-replica data storage policy for cloud disks to ensure data reliability, and you do not need to worry about damaged disks. For local disks, data reliability is ensured by upper-layer software. Disk and node faults require manual troubleshooting.

### EMR+d1 solution

EMR provides a complete set of automated O&M solutions to help you easily and reliably use instances with local disks, such as d1 series. You do not need to worry about the O&M process, because high data reliability and high service availability are ensured.

Highlights of automated O&M solutions:

- Highly reliable distribution of required nodes
- Fault monitoring of local disks and nodes
- Automatic assessment of data migration opportunities
- Automatic data migration of faulty nodes and data balancing
- Automatic HDFS data detection
- Network topology optimization

EMR helps you make better use of local disks and develop a cost-effective big data system by using automated O&M of the entire background management and control system.

Onte If you want to create a Hadoop cluster by using d1 series instances, .

### 2.7. Disaster recovery in EMR clusters

This topic describes the disaster recovery of data and services in EMR clusters.

### Disaster recovery of data

In HDFS, each file is divided into multiple blocks, each block has three replicas by default, and the replicas are distributed to different racks. You can customize the number of replicas. In most cases, the replication factor of HDFS is 3. One replica is placed on a node of the local rack, another replica is placed on a different node of the same rack, and the last replica is placed on a node of a remote rack.

HDFS scans the replicas on a regular basis. If a replica is missing, HDFS copies the data and generates a new replica. If a node is missing, HDFS restores all data on the node. If you use disks on Alibaba Cloud, three replicas are created for each disk. If an error occurs on one of the replicas, HDFS copies data from another replica to the failed replica to ensure data reliability.

HDFS is proven to be a reliable data storage system that provides storage for large volumes of data. You can also integrate HDFS with Object Storage Service (OSS) based on cloud features to back up data. This ensures higher data reliability.

### Disaster recovery for services

The core components of HDFS, such as YARN, HDFS, Hive Server, and Hive Meta, are deployed in high availability (HA) mode. In HA mode, at least two nodes are deployed for a service to support disaster recovery. When a node is faulty, the service is switched over to another node to ensure that the service is not affected.

### 3.Configure clusters 3.1. Create a cluster

This topic describes how to create an E-MapReduce (EMR) cluster.

#### Prerequisites

Role assignment is complete in the RAM console. For more information, see Assign roles.

#### Procedure

- 1. Go to the cluster creation page.
  - i. Log on to the Alibaba Cloud EMR console.
  - ii. In the top navigation bar, select the region where you want to create a cluster and select a resource group based on your business requirements.
    - The region of a cluster cannot be changed after the cluster is created.
    - All resource groups within your account are displayed by default.
  - iii. Click Cluster Wiz ard in the Clusters section.
- 2. Configure the cluster.

To create a cluster, you must configure software parameters, hardware parameters, and basic parameters as guided by the wizard.

Notice After a cluster is created, you cannot modify its parameters except for the cluster name. Make sure that all parameters are correctly configured when you create a cluster.

#### i. Configure software parameters.

Parameter

Description

Parameter	Description	
Cluster Type	<ul> <li>Description</li> <li>The type of the cluster that you want to create. EMR supports the following types of clusters:</li> <li>Hadoop: <ul> <li>Provides Hadoop, Hive, and Spark components that serve as semi-hosted services and are used to store and compute large-scale distributed data offline.</li> <li>Provides Spark Streaming, Flink, and Storm components for stream computing.</li> <li>Provides Presto and Impala components for interactive queries.</li> <li>Provides other Hadoop ecosystem components, such as Oozie and Pig.</li> </ul> </li> <li>Data Science: Data Science clusters are commonly used in big data and Al scenarios. Data Science clusters support the offline extract, transform, load (ETL) of big data based on Hive and Spark, and TensorFlow model training. You can choose the CPU+ GPU heterogeneous computing framework and deep learning algorithms supported by NVIDIA GPUs to run computing jobs more efficiently.</li> <li>Druid: provides a semi-hosted, real-time, and interactive analytics services. Druid clusters can query big data within milliseconds and ingest data in multiple ways. You can use Druid clusters with services such as EMR Hadoop, EMR Spark, Object Storage Service (OSS), and ApsaraDB RDS to build a flexible and stable system for real-time queries.</li> <li>ClickHouse: an open source column-oriented database management system (DBMS) for online analytical processing (OLAP).</li> <li>Dataflow: provides an end-to-end (E2E) real-time computing solution. Dataflow clusters incorporate Kafka, a distributed message system with high throughput and scalability, and the commercial Flink kernel provided by Apache Flink-powered Ververica. The clusters are used to resolve various E2E real-time computing issues and are widely used in real-time data ETL, and log collection and analysis scenarios. You are free to use one of the two components or both.</li> </ul> <li>Presto: an open source interactive query engine that provides the SQL on Everything capability. Presto clusters can be used to quickly</li>	
Cloud Native Option	on ECS is selected by default.	
EMR Version	The major version of EMR. The latest version is selected by default.	
Required Services	The default components required for a specific cluster type. After a cluster is created, you can start or stop components on the cluster management page.	

Parameter	Description		
	The other components that you can specify based on your business requirements. By default, the relevant service processes for the components you specify are started.		
Optional Services	<b>Note</b> The more components you specify, the higher instance specifications a cluster needs to handle the components. You must select the instance type that matches the number of components you specified when you configure the hardware. Otherwise, the cluster may have insufficient resources to run the components.		
Advanced Settings	<ul> <li>Kerberos Mode: specifies whether to enable Kerberos authentication for clusters. This feature is disabled by default. It is not required by clusters created for common users.</li> <li>Custom Software Settings: customizes software settings. You can use a JSON file to customize the parameters of the basic components required for a cluster, such as Hadoop, Spark, and Hive. For more information, see Customize software configurations. This feature is disabled by default.</li> </ul>		

#### ii. Configure hardware parameters.

|--|

Description	
pports the following billing allows you to pay for an system charges you for a actually used. You are and that you use pay-as- or dynamically scheduled ws you to use an instance e a pay-as-you-go ster passes the test, cluster for production. Iling Method, you must d and Auto Renewal. By d is one month and the ned on. If you turn on ystem renews your th seven days before nformation, see	Billing d Method
er. Zones are different ion. They are ost cases, you can use the	Zone
etwork type is selected by	Network Type
ant to deploy the cluster. . If no VPC is available in reate a VPC.	VPC
in the specified zone. If vSwitch.	vSwitch
reate a VPC. in the specified	

Section	Parameter	Description		
	Security Group Name	The security group of the cluster. An existing security group is selected by default. For more information about security groups, see Overview. You can click Create Security Group and enter a security group name to create a security group.		
High Availabilit Y	High Availabilit Y	<ul> <li>This feature is disabled by default. For a Hadoop cluster, if High Availability is enabled, two or three master nodes are created in the cluster to ensure the availability of the ResourceManager and NameNode processes.</li> <li>HBase clusters always work in high availability mode. If you do not enable high availability, only one master node is created, but a core node is used to support high availability. If you enable high availability, two master nodes are created to ensure higher security and reliability.</li> </ul>		

Section	Parameter	Description		
		<ul> <li>Master Instance: runs control processes, such as ResourceManager and NameNode.</li> </ul>		
		You can select an instance type based on your business requirements. For more information, see Instance families.		
		<ul> <li>System Disk Type: You can select an SSD, ESSD, or ultra disk based on your business requirements.</li> </ul>		
		<ul> <li>Disk Size: You can resize a disk based on your business requirements. The recommended minimum disk size is 120 GB. Valid values: 40 to 2048. Unit: GB.</li> </ul>		
		<ul> <li>Data Disk Type: You can select an SSD, ESSD, or ultra disk based on your business requirements.</li> </ul>		
		<ul> <li>Disk Size: You can resize a disk based on your business requirements. The recommended minimum disk size is 80 GB. Valid values: 40 to 32768. Unit: GB.</li> </ul>		
		<ul> <li>Master Nodes: One master node is configured by default. If high availability is enabled, two or three master nodes are configured.</li> </ul>		
		<ul> <li>Core Instance: stores all the data of a cluster. You can add core nodes as needed after a cluster is created.</li> </ul>		
Instance	Learn More	<ul> <li>System Disk Type: You can select an SSD, ESSD, or ultra disk based on your business requirements.</li> </ul>		
		<ul> <li>Disk Size: You can resize a disk based on your business requirements. The recommended minimum disk size is 120 GB.</li> </ul>		
		<ul> <li>Data Disk Type: You can select an SSD, ESSD, or ultra disk based on your business requirements.</li> </ul>		
		<ul> <li>Disk Size: You can resize a disk based on your business requirements. The recommended minimum disk size is 80 GB.</li> </ul>		
		<ul> <li>Core Nodes: Two core nodes are configured by default. You can change the number of core nodes based on your business requirements.</li> </ul>		
		• <b>Task Instance</b> : stores no data. It is used to adjust the computing capabilities of clusters. No task node is configured by default. You can add task nodes based on your business requirements.		

iii. Configure basic parameters.

Section	Parameter	Description		
	Cluster Name	The name of the cluster. The name must be 1 to 64 characters in length and can contain only letters, digits, hyphens (-), and underscores (_).		
Basic Informati on	Туре	<ul> <li>DLF Unified Metadata (recommended): Metadata is stored in a data lake.</li> <li>Data Lake Formation (DLF) provides a fully managed, maintenancefree, unified metadata service that provides high availability and performance. The metadata service is compatible with multiple versions of Hive and facilitate metadata migration between a Hive metastore and DLF. For more information, see Overview.</li> <li>Self-managed RDS: Metadata is stored in an ApsaraDB RDS database. For more information, see Configure an independent ApsaraDB RDS for MySQL database.</li> <li>Built-in MySQL (not recommended): Metadata is stored in the local MySQL database of a cluster.</li> <li>Note You can select this option only in test scenarios. The local MySQL database is deployed on a single node of an EMR cluster. This cannot ensure high availability for services and may cause stability risks. We recommend that you select DLF Unified Metadata or Self-managed RDS in production scenarios.</li> </ul>		
	Assign Public IP Address	Specifies whether an elastic IP address (EIP) is associated with the cluster. This feature is disabled by default.  Once To access the cluster over the Internet, apply for a public IP address on ECS. For information about how to apply for an EIP address, see Elastic IP addresses.		
	Key Pair	For information about how to use a key pair, see SSH key pair overview.		
	Password	The password used to log on to a master node. The password mu be 8 to 30 characters in length and contain uppercase letters, lowercase letters, digits, and special characters. The following special characters are supported: ! @ # \$ % ^ & *		
	Add User	The user added to access the web UIs of open source big data software.		

Section	Parameter	Description		
Advanced Settings	Permissio n Settings	<ul> <li>The RAM roles that allow applications running in a cluster to access other Alibaba Cloud services. You can use the default RAM roles.</li> <li>EMR Role: The value is fixed as AliyunEMRDefaultRole and cannot be changed. This RAM role authorizes a cluster to access other Alibaba Cloud services, such as ECS and OSS.</li> <li>ECS Role: You can also assign an application role to a cluster. Then, EMR applies for a temporary AccessKey pair when applications running on the compute nodes of that cluster access other Alibaba Cloud services, such as OSS. This way, you do not need to manually enter an AccessKey pair. You can grant the access permissions of the application role on specific Alibaba Cloud services based on your business requirements.</li> </ul>		
	Data Disk Encryption	This feature is disabled by default. If you turn on <b>Enable Encryption</b> , data in all cloud disks that serve as the data disks of the ECS instances in the cluster is encrypted. By default, a service-managed key is used to encrypt your data. You can also use a user-managed key to encrypt your data.		
	Bootstrap Actions	Optional. You can configure bootstrap actions to run custom scripts before a cluster starts Hadoop. For more information, see Manage bootstrap actions.		
	Tag	Optional. You can add a tag pair when you create a cluster or add a tag pair on the cluster details page after a cluster is created. For more information, see Manage and use tags.		
	Resource Group	Optional. For more information, see Use resource groups.		

**Note** The cluster configurations appear on the right side of the page when you configure parameters. After you complete the configurations, click Next: Confirm. You are directed to the Confirm step, in which you can confirm the configurations and the fee for the creation of your cluster. The fee varies based on the billing method.

### 3. Verify that the configuration is correct, read and select E-MapReduce Service Terms, and then click **Create**.

#### ♥ Notice

- Pay-as-you-go clusters: Creation immediately starts after you click Create.
  - After the cluster is created, its status changes to Idle.
- Subscription clusters: An order is generated after you click Create. The cluster is created after you pay the fee.

### 3.2. Create a gateway cluster

You can use a gateway cluster to balance loads and isolate clusters in a secure manner. You can also use the gateway cluster to submit jobs to an E-MapReduce (EMR) cluster. This topic describes how to create a gateway cluster.

### Prerequisites

A Hadoop or Kafka cluster is created in EMR. For more information, see Create a cluster.

### Limits

A gateway cluster can be associated with only a Hadoop or Kafka cluster in EMR.

### Procedure

- 1. Go to the Cluster Management page.
  - i.
  - ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
  - iii. Click the **Cluster Management** tab.
- 2. In the upper-right corner of the **Cluster Management** page, click **CreateGateway**.
- 3. On the **Create Gateway** page, configure the parameters.

Section	Parameter	Description
	Cluster Name	The name of the gateway cluster. The name must be 1 to 64 characters in length and can contain only letters, digits, hyphens (-), and underscores (_).
	Assign Public IP Address	Specifies whether to assign an elastic IP address to the gateway cluster.
		• <b>Password</b> : the password that is used to log on to the gateway cluster. The password must be 8 to 30 characters in length and must contain uppercase letters, lowercase letters, digits, and special characters.
		The password can contain the following special characters:
		! @ # \$ % ^ & *
Basic Information	Password and Key Pair	<ul> <li>Key Pair: the name of the key pair that is used to log on to the gateway cluster. If no key pair is created, click</li> <li>Create Key Pair next to this field to go to the SSH Key Pairs page of the Elastic Compute Service (ECS) console and create a key pair.</li> </ul>
inionnation		Keep the <i>.pem</i> private key file secure. After you create a gateway cluster, the public key is automatically bound to the ECS instance. When you log on to the gateway cluster by using SSH, you must enter the private key in the private key file.

Section	Parameter	Description
	Billing Method	<ul> <li>Subscription: You are charged only once for each subscription period. The unit price of a subscription cluster is lower than the unit price of a pay-as-you-go cluster of the same specifications. The amount of savings increases with the subscription period.</li> <li>Pay-As-You-Go: You are charged for the hours during which a cluster is running.</li> </ul>
	Associated Cluster	The cluster that you want to associate with the gateway cluster. The gateway cluster submits jobs to this cluster.
	Zone	The zone where the associated cluster resides.
	Network Type	The network type of the associated cluster.
Cluster Configuration	VPC	The virtual private cloud (VPC) to which the associated cluster belongs.
	VSwitch	The vSwitch that you want the gateway cluster to use. Select a vSwitch in the zone and the VPC in which the cluster resides.
	Security Group Name	The name of the security group to which the associated cluster belongs.
		<ul> <li>The available instance types in the current region. For more information, see Instance families.</li> <li>System Disk Type: the type of the system disk that you want the gateway cluster to use. System disks are classified into ultra disks, standard SSDs, and ESSDs. The types of system disks that you can use to create a gateway cluster vary based on the region and instance type that you select. By default, system disks are released after the relevant cluster is released.</li> </ul>
Instance	Gateway Instance	<ul> <li>Disk Size: the size of the system disk. Unit: GB. Valid values: 40 to 500. Default value: 300.</li> <li>Data Disk Type: the type of data disks you want the gateway cluster to use. Data disks are classified into ultra disks, standard SSDs, and ESSDs. The types of data disks that you can use to create a gateway cluster vary based on the region and instance type that you select. By default, the data disks are released after the relevant cluster is released.</li> <li>Disk Size: the size of a data disk. Unit: GB. Valid values: 200 to 4000. Default value: 300.</li> <li>Count: the number of data disks. Valid values: 1 to 10.</li> </ul>

Section	Parameter	Description
Advanced Settings	Permission Settings	<ul> <li>The RAM roles that allow applications running in a cluster to access other Alibaba Cloud services. You can use the default RAM roles.</li> <li>EMR Role: This parameter has a fixed value of AliyunEMRDefaultRole and cannot be modified. This RAM role allows a cluster to access other Alibaba Cloud services, such as ECS and Object Storage Service (OSS).</li> <li>ECS Role: You can also assign an application role to a cluster. Then, EMR applies for a temporary AccessKey pair when applications running on the compute nodes of that cluster access other Alibaba Cloud services, such as OSS. This way, you do not need to manually enter an AccessKey pair. You can grant the access permissions of the application role on specific Alibaba Cloud services based on your business requirements.</li> </ul>
	Bootstrap Actions	Optional. You can configure bootstrap actions to run custom scripts before a cluster starts. For more information, see Bootstrap actions.
	Data Disk Encryption	This feature is disabled by default. If you turn on <b>Enable Encryption</b> , data in all cloud disks that serve as the data disks of the ECS instances in the cluster is encrypted. By default, a service-managed key is used to encrypt your data. You can also use a user-managed key to encrypt your data.

 Read the terms of service, select E-MapReduce Service Terms, and then click Create. The gateway cluster that you created appears in the cluster list and its state changes from Initializing to Idle a few minutes later.

## 3.3. Manage permissions

### 3.3.1. Assign roles

When you run components such as Hadoop and Spark in your E-MapReduce (EMR) cluster, you must grant the components the permissions to access other Alibaba Cloud services and perform related operations. Each EMR cluster must be configured with service roles and ECS application roles. This topic describes how to assign roles to EMR and also describes the roles that are associated with EMR.

### Context

EMR provides default system roles and default system policies. System policies are created and maintained by Alibaba Cloud. If service requirements change, the system policies are accordingly updated.

When you use EMR for the first time, you must use your Alibaba Cloud account to assign the AliyunEMRDefaultRole and AliyunECSInstanceForEMRRole or AliyunEmrEcsDefaultRole roles to EMR. After the roles are assigned, you can view the roles in the RAM console and attach policies to the roles. For more information about the roles, see RAM role overview.

#### ♥ Notice

- Roles that are required for EMR vary based on the EMR version.
  - In EMR V3.32.0 or an earlier minor version, or EMR V4.5.0 or an earlier minor version: AliyunEmrEcsDef ault Role
  - In a minor version later than EMR V3.32.0 or EMR V4.5.0: AliyunECSInstanceForEMRRole
- When you use EMR for the first time, you must use your Alibaba Cloud account to assign default system roles to EMR. Otherwise, your Alibaba Cloud account and the RAM users within your Alibaba Cloud account cannot use EMR.
- If you want to delete a service role, make sure that the resources that use the role are released. Otherwise, the use of the resources is affected.
- If only some roles are assigned, the EMR console sends you a notification. You can create a cluster only after all roles are assigned.

#### Procedure

EMR V3.30 is used in this example.

1. In the Alibaba Cloud EMR console, click Authorize in RAM.

Onte When you use EMR for the first time, you must assign default system roles to EMR. Then, you do not need to repeat the assignment operation when you use EMR again.

When you create an EMR cluster or create an execution plan as required with a new cluster, if default roles are not assigned to EMR, the following information appears.

0	Role Authorization	$\times$
	To use the service, you must be assigned the following two default EMR roles by the Alibaba Cloud account.: Authorize in RAM	
	AliyunEMRDefaultRole	
	AliyunEmrEcsDefaultRole	
	EMR Default Role Description: Role Authorization	

2. On the page that appears, click **Confirm Authorization Policy** to assign the default roles **AliyunEMRDefault Role** and **AliyunEmrEcsDefault Role** to EMR.

AliyunEmrEc	sDefaultRole				
Description: The t	tasks of EMR service will use this role	to access OSS resources.			
Permission Descri	iption: The policy for AliyunEMRECS	DefaultRole, including the permission	for OSS.		
an and and	- M. David, M.			 	
AliyunEMRD	efaultRole				
Description: The E	EMR service will use this role to acce	ss ECS resources.			

3. Refresh the EMR console to use the services.

To view policy details for the **AliyunEMRDefault Role** and **AliyunEmrEcsDefault Role** roles, log on to the **RAM console**.

#### Service roles

The following table describes the RAM roles that are associated with EMR.

Attribute	Default role	Description	System policy
EMR service role	AliyunEMRDefaultR ole	This role allows you to use EMR to access other Alibaba Cloud services when you configure resources and perform service- level operations on your EMR cluster. This role is required for all clusters and cannot be changed. For more information, see EMR service role.	AliyunEMRRolePolic Y
ECS application role (used in EMR V3.32.0 or an earlier minor version, or EMR V4.5.0 or an earlier minor version)	AliyunEmrEcsDefau ltRole	This role allows application processes that run on your cluster to access other Alibaba Cloud services. When you create a cluster, you can use this service role or use a custom role. For more information about this role, see ECS application role (used in EMR V3.32.0 and earlier V3.X.X versions as well as in EMR V4.5.0 and earlier V4.X.X versions).	AliyunEMRECSRoleP olicy

Attribute	Default role	Description	System policy
ECS application role (used in a minor version later than EMR V3.32.0 or EMR V4.5.0)	AliyunECSInstanceF orEMRRole	This role allows application processes that run on your cluster to access other Alibaba Cloud services. When you create a cluster, you can use this service role or use a custom role. For more information about this role, see ECS application role (used in V3.X.X versions later than EMR V3.32.0 as well as in V4.X.X versions later than EMR V4.5.0).	AliyunECSInstanceF orEMRRolePolicy
ECS application role (used in EMR Studio by default)	AliyunECSInstanceF orEMRStudioRole	This role allows you to use EMR Studio to access your resources in other Alibaba Cloud services. If this role is not assigned to your account, a window appears, which prompts you to assign this role when you create an EMR Studio cluster for the first time. To assign this role, use your Alibaba Cloud account.	AliyunECSInstanceF orEMRStudioRoleP olicy

### 3.3.2. EMR service role

The E-MapReduce (EMR) service role allows you to use EMR to access other Alibaba Cloud services when you configure resources or perform service-level operations on your EMR cluster. For example, the service role can be used to create an Elastic Compute Service (ECS) instance when you start an EMR cluster. This topic describes the EMR service role AliyunEMRDefaultRole and the policies of this role.

### Usage notes

To avoid impacts on the service stability of EMR, take note of the following points:

- The EMR service role cannot be changed.
- Do not delete or modify system policies of this role in the RAM console.

### Permissions

The default role **AliyunEMRDefaultRole** is configured with the policy **AliyunEMRRolePolicy**. The following tables describe the permissions of the **AliyunEMRRolePolicy** policy.

• ECS-related permissions

Permission (Action)	Description
ecs:CreateInstance	Creates an ECS instance.
ecs:RenewInstance	Renews an ECS instance.

Permission (Action)	Description
ecs:DescribeRegions	Queries the region information of an ECS instance.
ecs:DescribeZones	Queries the zone information of an ECS instance.
ecs:Describelmages	Queries the image information of an ECS instance.
ecs:CreateSecurityGroup	Creates a security group.
ecs:AllocatePublicIpAddress	Assigns a public IP address to an ECS instance.
ecs:DeleteInstance	Deletes an ECS instance.
ecs:StartInstance	Starts an ECS instance.
ecs:StopInstance	Stops an ECS instance.
ecs: DescribeInstances	Queries ECS instances.
ecs: DescribeDisks	Queries the disk information of an ECS instance.
ecs:AuthorizeSecurityGroup	Specifies inbound rules for a security group.
ecs:AuthorizeSecurityGroupEgres s	Specifies outbound rules for a security group.
ecs:DescribeSecurityGroupAttribu te	Queries details of a security group.
ecs:DescribeSecurityGroups	Queries security groups.

#### • OSS-related permissions

Permission (Action)	Description
oss:PutObject	Uploads a file or folder.
oss:GetObject	Obtains a file or folder.
oss:ListObjects	Queries files.

### 3.3.3. ECS application role (used in EMR V3.32.0 and earlier V3.X.X versions as well as in EMR V4.5.0 and earlier V4.X.X versions)

E-MapReduce (EMR) provides MetaService, which serves as a special ECS application role. In EMR V3.32.0 and earlier V3.X.X versions as well as in EMR V4.5.0 and earlier V4.X.X versions, when you create a cluster, this role is automatically bound to your cluster. Applications that run on your EMR cluster use this role to access other Alibaba Cloud resources without an AccessKey pair. This avoids the disclosure of the AccessKey pair in a configuration file.

### Prerequisites

This role is authorized. For more information, see Assign roles.

#### **Background information**

MetaService allows you to access only Object Storage Service (OSS), Log Service, and Message Service (MNS) without an AccessKey pair.

#### Permissions

The default role **AliyunEmrEcsDefault Role** is configured with the policy **AliyunEmrECSRolePolicy**. The following table describes OSS-related permissions.

Permission (Action)	Description
oss:Put Object	Uploads a file or folder.
oss:GetObject	Obtains a file or folder.
oss:ListObjects	Queries files.
oss:DeleteObject	Deletes a file.
oss:AbortMultipartUpload	Terminates a multipart upload event.

Notice Modify or delete the AliyunEmrEcsDefaultRole role with caution. Otherwise, your cluster fails to be created or jobs fail to be run.

### Data sources that support MetaService

MetaService allows you to access OSS, Log Service, and MNS. You can use an EMR SDK in your EMR cluster to read data from and write data to the preceding data sources without an AccessKey pair.

By default, only access to OSS is enabled. If you want to read data from and write data to Log Service and MNS, log on to the RAM console and configure the required permissions for the AliyunEmrEcsDefaultRole role. For more information, see .

For more information about how to authorize a RAM role, see Grant permissions to a RAM role.

#### Use MetaService

MetaService allows you to access OSS, Log Service, and MNS without an AccessKey pair. MetaService provides the following benefits:

- Reduces the risk of AccessKey information leak. To minimize the security risk, authorize roles in the RAM console based on the principle of least privilege.
- Improves user experience. MetaService shortens the OSS path that you need to enter during interactive access to OSS resources.
- Brings the following benefits for services in your EMR cluster:

The jobs that you run in the services can access Alibaba Cloud resources (OSS, Log Service, and MNS) without an AccessKey pair.

Comparison of operations before and after MetaService is used:

- Run the hadoop fs -ls command to view OSS data.
  - MetaService is not used:

```
hadoop fs -ls oss://ZaH*****As1s:Ba23N*********sdaBj2@bucket.oss-cn-hangzhou-int ernal.aliyuncs.com/a/b/c
```

MetaService is used:

```
hadoop fs -ls oss://bucket/a/b/c
```

#### • Create an external table in Hive.

MetaService is not used:

```
CREATE EXTERNAL TABLE test_table(id INT, name string)

ROW FORMAT DELIMITED

FIELDS TERMINATED BY '/t'

LOCATION 'oss://ZaH*****As1s:Ba23N*********sdaBj2@bucket.oss-cn-hangzhou

-internal.aliyuncs.com/a/b/c';
```

#### MetaService is used:

```
CREATE EXTERNAL TABLE test_table(id INT, name string)
ROW FORMAT DELIMITED
FIELDS TERMINATED BY '/t'
LOCATION 'oss://bucket/a/b/c';
```

#### • Use Spark to view OSS data.

MetaService is not used:

```
val data = sc.textFile("oss://ZaH*****As1s:Ba23N********sdaBj2@bucket.oss-cn-ha
ngzhou-internal.aliyuncs.com/a/b/c")
```

MetaService is used:

```
val data = sc.textFile("oss://bucket/a/b/c")
```

• Brings the following benefits for self-deployed services:

MetaService is an HTTP service. You can access the URL of this HTTP service to obtain a Security Token Service (STS) temporary credential. Then, you can use the STS temporary credential to access Alibaba Cloud resources without an AccessKey pair in self-managed systems.

Notice A new STS temporary credential is generated 30 minutes before the current one expires. Both STS credentials can be used within the 30 minutes.

### For example, you can run **curl http://localhost:10011/cluster-region** to obtain the region where your cluster resides.

You can use MetaService to obtain the following information:

- Region: /cluster-region
- Role name: /cluster-role-name
- AccessKey ID: /role-access-key-id
- AccessKey secret: /role-access-key-secret

- Security token: /role-security-token
- Network type: /cluster-network-type

# 3.3.4. ECS application role (used in V3.X.X versions later than EMR V3.32.0 as well as in V4.X.X versions later than EMR V4.5.0)

In V3.X.X versions later than E-MapReduce (EMR) V3.32.0 as well as in V4.X.X versions later than EMR V4.5.0, MetaService is replaced with the ECS application role AliyunECSInstanceForEMRRole. This role is automatically assigned to each ECS instance in your EMR cluster when you create or scale out the cluster. Applications that run on your EMR cluster use this role to access other Alibaba Cloud resources without an AccessKey pair. This avoids the disclosure of the AccessKey pair in a configuration file.

### Prerequisites

This role is authorized. For more information, see Assign roles.

### Permissions

The role **AliyunECSInstanceForEMRRole** is configured with the policy **AliyunECSInstanceForEMRRolePolicy**. The following table describes OSS-related permissions of this role.

Permission (Action)	Description
oss:Put Object	Uploads a file or folder.
oss:GetObject	Obtains a file or folder.
oss:ListObjects	Queries files.
oss:DeleteObject	Deletes a file.
oss:AbortMultipartUpload	Terminates a multipart upload event.

Notice Modify or delete the AliyunEmrEcsDefaultRole role with caution. Otherwise, your cluster fails to be created or jobs fail to be run.

### Use the ECS application role to obtain an STS temporary credential

You can use a Security Token Service (STS) temporary credential to access other Alibaba Cloud services within your account. For more information, see Use RAM roles to access other Alibaba Cloud services.

### 3.3.5. Use a custom ECS application role to access other cloud resources in your Alibaba Cloud account
E-MapReduce (EMR) allows you to configure a custom ECS application role in the Advanced Settings section of the Basic Settings step when you create a cluster in the EMR console. You can use this role to access other cloud resources, such as Object Storage Service (OSS) and Log Service, in your Alibaba Cloud account in password-free mode. This topic describes how to create and use a custom ECS application role.

#### Context

To manage access to external resources from an EMR cluster, you can attach policies to the custom ECS application role that is configured when you create the cluster. For example, you can impose the following limits:

- A cluster can access only the OSS data directories that you specify.
- A cluster can access the external resources that you specify.

#### Prerequisites

An OSS bucket is created in the OSS console. Make sure that the bucket resides in the region where you want to create an EMR cluster. For more information about how to create a bucket, see Create buckets.

#### Procedure

- 1. Step 1: Create a policy
- 2. Step 2: Create a RAM role
- 3. Step 3: Create a cluster and access external resources

#### Step 1: Create a policy

- 1. Go to the Create Custom Policy page.
  - i. Log on to the RAM console by using an Alibaba Cloud account.
  - ii. In the left-side navigation pane, choose **Permissions > Policies**.
  - iii. On the Policies page, click Create Policy.
- 2. On the **Create Custom Policy** page, configure the parameters.

Parameter	Description
Policy Name	In this example, the policy name is test-emr.
Configuration Mode	Select Script.

3. Click OK.

#### Step 2: Create a RAM role

- 1. In the left-side navigation pane of the RAM console, choose **Identities > Roles**.
- 2. On the Roles page, click Create Role.
- 3. Create a RAM role.
  - i. In the Create Role panel, select Alibaba Cloud Service.

- ii. Click Next.
- iii. In the Configure Role step, configure the parameters.

	Parameter	Description	
RAM Role Name		In this example, the RAM role name is test-emr.	
Select Trusted Service		Select Elastic Compute Service from the drop-down list.	

- iv. Click OK.
- 4. (Optional)Change the trusted service.

✓ Notice If you want to create a cluster in EMR V5.X or later, a minor version later than EMR V4.5, or a minor version later than EMR V3.32, skip this step.

- i. On the Roles page, click the name of the RAM role that you created.
- ii. Click the Trust Policy Management tab.
- iii. Click Edit Trust Policy.



- v. Click OK.
- 5. Add the required permissions.
  - i. On the **Roles** page, find the **RAM role** that you created and click **Add Permissions** in the Actions column.
  - ii. In the Add Permissions panel, click Custom Policy and select the policy that you created in Step 1.
  - iii. Click OK.
  - iv. Click Complete.

#### Step 3: Create a cluster and access external resources

```
1.
```

- 2. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- 3. Click **Cluster Wizard** to create a cluster. In the **Advanced Settings** section of the **Basic Settings** step, enter the name of the RAM role that you created in Step 2: Create a RAM role in the ECS Role field. For more information about how to create a cluster, see Create a cluster.

Advanced Settir	ngs			
Add Use	r 🕜			+Add
	Set a password to continue.	Account	Actions	
			No Data	
Permissio	on Settings 🕜			
	EMR Role: 🕜 🛛 Al	iyunEMRDefaultRole		
	ECS Role: AI	iyunECSInstanceForEMRRole		

4. After the cluster is created, log on to the master node of the cluster by using SSH. For more information, see Log on to a cluster.

Run the following command to check whether you can access the OSS bucket that is specified in the policy:

hdfs dfs -ls oss://<yourBucketName>/

⑦ Note Replace < yourBucket Name> in the command with the name of the OSS bucket.

• If the following information is returned, access to the OSS bucket is denied.

```
[root@emr-header-1 ~]# hdfs dfs -ls oss://emr-logs2/
ls: java.io.IOException: ErrorCode : 403 , ErrorMsg: HTTP/1.1 403 Forbidden: <?xml version="1.0" enco
ding="UTF-8"?>
```

• If the following information is returned, access to the OSS bucket is allowed.

[root@emr-he	eader-	1 ~]# hd1	fs dfs -ls o	ss://em	r-logs2/
Found 10 ite	ems				
-rw-rw-rw-	1	30	2020-03-10	16:08	oss://emr-logs2/
drwxrwxrwx		0	2019-12-20	14:49	oss://emr-logs2/
drwxrwxrwx		0	2020-03-12	23:56	oss://emr-logs2/
-rw-rw-rw-	1	161	2020-05-18	22:01	oss://emr-logs2/
drwxrwxrwx		0	2019-12-20	14:22	oss://emr-logs2/
drwxrwxrwx		0	2020-05-18	22:28	oss://emr-logs2/

#### FAQ

• Q: What do I do if the NoPermission message appears when I create a cluster?

A: Perform the following operations in sequence to resolve the issue:

- i. Check whether the RAM user that you used to create the cluster is authorized to create clusters and change ECS application roles. If the RAM user is granted the AliyunEMRDevelopAccess permission, you must change the permission to AliyunEMRFullAccess.
- ii. Check whether the name of the ECS application role that you specified when you create the cluster is valid.
- iii. Check whether the trust policy is changed to *emr.aliyuncs.com*.
- Q: What do I do if I cannot access an OSS bucket from Hadoop Distributed File System (HDFS)?

A: Perform the following operations in sequence to resolve the issue:

- i. Check whether the OSS bucket that you want to access resides in the same region as your cluster.
   If the OSS bucket and your cluster do not reside in the same region, you must add the endpoint of the OSS bucket to the access links.
- ii. Check whether the OSS bucket that you want to access is one of the resources specified in the created policy. If the OSS bucket is not one of the resources, modify the policy.
- iii. Check whether the related permissions on the OSS bucket are configured in the OSS console. If the permissions are configured, you must cancel the settings in the OSS console and specify Action in the policy to configure the permissions.

### 3.3.6. Grant permissions to RAM users

If you want to allow a RAM user to use the E-MapReduce (EMR) console, you must grant the required permissions to the RAM user by using your Alibaba Cloud account in the RAM console.

#### Context

RAM is a resource access control service provided by Alibaba Cloud. For more information, see What is RAM? The following examples describe how RAM is used to implement access control in EMR:

- RAM users: If you purchased multiple instances for an EMR cluster, you can create a policy that allows specific users who are responsible for O&M, development, or data analysis to use these instances. This eliminates the risk of AccessKey pair leaks and ensures account security.
- RAM user groups: You can create multiple user groups and grant different permissions to them. The authorization process is the same as that for RAM users. The user groups can be used to manage multiple RAM users at the same time.

#### Policies

Policies are categorized into system policies and custom policies.

• System policies

System policies are provided by Alibaba Cloud to meet various management purposes. The following table describes the system policies that are used in EMR.

System policy	Description	Permission
AliyunEMRFullAccess	Provides RAM users with full access to EMR.	This policy allows RAM users to perform all operations on all EMR resources.
AliyunEMRDevelopAc cess	Provides RAM users with the developer permissions of EMR.	This policy allows RAM users to perform operations on all EMR resources, except for the operations to create and release clusters.
AliyunEMRFlowAdmin	Provides RAM users with the administrator permissions on the Data Platform module in EMR.	This policy allows RAM users to create projects and develop and manage jobs. This policy does not allow RAM users to add members to projects or manage clusters.

#### • Custom policies

Custom policies are the policies that you design based on your business requirements. Custom policies are suitable for users who are familiar with Alibaba Cloud service APIs and require fine-grained access control. For more information about how to create a custom policy, see Policy structure and syntax.

#### Grant permissions to a RAM user

Perform the following steps to grant permissions on EMR resources to a RAM user in the RAM console:

- 1. Log on to the RAM console by using your Alibaba Cloud account.
- 2. In the left-side navigation pane, choose **Identities > Users**.
- 3. On the Users page, find the RAM user to which you want to grant permissions and click Add **Permissions** in the Actions column.
- 4. In the Add Permissions panel, configure the parameters that are described in the following table based on your business requirements.
- 5. Click OK.
- 6. Click Complete.

The granted permissions immediately take effect. You can log on to the RAM console by using the RAM user to which you granted permissions to check the permissions.

# 3.4. Manage user accounts

This topic describes how to manage E-MapReduce (EMR) user accounts on the Users page of the EMR console.

#### **Background information**

Information about EMR user accounts is stored in the built-in OpenLDAP service of an EMR cluster. You can use the information to authenticate EMR users in the EMR cluster.

If you click the link of an open source component on the Connect Strings page to access the web UI of the component, you must use an EMR user account for identity authentication. If you enable LDAP authentication, you must also use an EMR user account for identity authentication. If you configure LDAP as the user source for Ranger, you can manage the permissions of user accounts that are listed on the Users page. You can use EMR user accounts to run kinit commands on a high-security cluster.

The Users page lists all EMR user accounts. RAM users that correspond to the EMR user accounts are classified into the following types based on the permissions that are granted to the RAM users in the EMR console:

- Administrator: an Alibaba Cloud account, or a RAM user that is granted the emr:ManageUserPlatform and emr:CreateLdapUser permissions, such as a RAM user to which the AliyunEMRFullAccess policy is attached. An administrator can view the information about all user accounts that are configured in a cluster. The administrator can also add or remove a user account, reset the password of a user account, modify the remarks of a user account, and download the authentication credentials of a user account. Authentication credentials can be downloaded only in a high-security cluster.
- Common user: RAM users to which other policies, such as AliyunEMRDevelopAccess, are attached. A common user can only view the information about the EMR user account whose username is the same as the username of the common user, reset the password, modify the remarks, and download the authentication credentials of the user account. A common user cannot add or remove a user account.

#### Prerequisites

• An EMR cluster is created. For more information, see Create a cluster.

#### • RAM users are created. For more information, see Create a RAM user.

**Note** You must create a RAM user first. Only an EMR user account whose username is the same as the username of a RAM user can be added to the Users page of the EMR console.

#### Limits

Only Hadoop clusters, Data Science clusters, Flink clusters, and Druid clusters support the management of EMR user accounts.

#### Add a user account

**Notice** If you use a RAM user to log on to the EMR console, you must grant the ram:ListUsers permission to the RAM user before you add a user account. You can attach the AliyunRAMReadOnlyAccess policy to the RAM user in the RAM console by using your Alibaba Cloud account. You can also configure a custom policy to grant the ram:ListUsers permission to the RAM user.

- 1. Go to the Cluster Overview page.
  - i.
  - ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
  - iii. Click the Cluster Management tab.
  - iv. On the Cluster Management page, find your cluster and click Details in the Actions column.
- 2. In the left-side navigation pane, click Users.
- 3. On the Users page, click Add User in the upper-left corner.
- 4. In the Add User dialog box, select an existing RAM user as an EMR user account from the Username drop-down list and specify Password and Confirm password.
- 5. Click OK.

#### Remove a user account

1. Go to the Cluster Overview page.

i.

- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the Cluster Management tab.
- iv. On the **Cluster Management** page, find your cluster and click **Details** in the Actions column.
- 2. In the left-side navigation pane, click Users.
- 3. On the **Users** page, find the user account that you want to remove and click **Delete** in the Action column.
- 4. In the Delete message, click OK.

#### Reset the password of a user account

You can reset the password of a user account.

Notice This operation may cause tasks that are running to fail.

- 1. Go to the Cluster Overview page.
  - i.
  - ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
  - iii. Click the Cluster Management tab.
  - iv. On the Cluster Management page, find your cluster and click Details in the Actions column.
- 2. In the left-side navigation pane, click Users.
- 3. On the Users page, find the user account whose password you want to reset and click Reset Password in the Action column.
- 4. In the **Set Password** dialog box, enter the new password in the **Password** and **Confirm Password** fields.
- 5. Click OK.

#### Download the authentication credentials of a user account

**Notice** Authentication credentials can be downloaded only in high-security clusters. You can download the keytab file of a user account.

- 1. Go to the Cluster Overview page.
  - i.
  - ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
  - iii. Click the Cluster Management tab.
  - iv. On the Cluster Management page, find your cluster and click Details in the Actions column.
- 2. In the left-side navigation pane, click Users.
- 3. On the **Users** page, find the user account whose authentication credentials you want to download and click **Download Authentication Credentials** in the Action column.

#### Update user account configurations

You can update the user account configurations that fail to take effect at the earliest opportunity due to network latency. You can also click Update to synchronize the added OpenLDAP user accounts to the user account list on the Users page.

On the Users page, click Update to update the user account configurations.

#### Manage Linux user accounts

You can use this feature for high-security clusters with a self-managed LDAP server deployed. When you add a Linux user account, a Linux user account with a specific name is automatically created for each node in the cluster. This account is also available on the nodes that are added when you scale out the cluster in the future.

- 1. In the upper-right corner of the Users page, click Linux Users.
- 2. In the Linux User Management dialog box, enter usernames. Separate multiple usernames with semicolons (;).

3. Click Add.

#### FAQ

Q: Can different clusters share an EMR user account?

A: No, different clusters cannot share an EMR user account. EMR user accounts listed on the Users page are valid only for the current cluster. For example, EMR user account A created in cluster-1 cannot be shared with cluster-2. To use EMR user account A in cluster-2, you must create this account in cluster-2.

# 3.5. Manage and use tags

Tags are used to label clusters. You can use tags to identify and manage your cluster resources. You can add tags when you create a cluster. You can also add tags on the Cluster Overview page of a created cluster. You can add up to 20 tags to a single cluster. This topic describes how to add, remove, and view tags and how to use tags to search for clusters.

#### Context

An E-MapReduce (EMR) cluster contains multiple Elastic Compute Service (ECS) instances. When you create an EMR cluster, two system tags are automatically added to each ECS instance in the cluster. To identify the cluster to which an ECS instance belongs and to identify the role of the ECS instance in the cluster, perform the following steps: Log on to the Alibaba Cloud ECS console and move the pointer over the tag icon of an ECS instance on the Instances page.

	Instance ID/Name	Tag	Monito	oring Zone
	acs:emr:clusterld:C-6F6E5C5 i-op1383wskm EMR_C-6F6E5( RKI		nr:hostGroupTyp	e:CORE Edit Tags Hangzhou Zone I
		ER- 🗣	c) 🛆	Hangzhou Zone I

For example, an ECS instance has the following system tags:

- acs:emr:clusterId=C-A510C93EA117\*\*\*\*
- acs:emr:hostGroupType=CORE

In this example, the ECS instance is a core node of an EMR cluster whose ID is C-A510C93EA117\*\*\*\*.

#### Precautions

- If you update the tags of an EMR cluster, the tags of the ECS instances in the cluster are also updated in the ECS console.
- If you update the tags of ECS instances in the ECS console, the tags of the EMR cluster to which the ECS instances belong are not updated. To make sure that the tags of ECS instances are consistent with the tags of the EMR cluster, we recommend that you do not update the tags of ECS instances in the ECS console. If the number of tags of an ECS instance in a cluster reaches the upper limit, you can no longer create tags for the cluster.
- Tag information is not shared across regions.

For example, tags that are created in the China (Hangzhou) region are invisible in the China (Shanghai) region.

• You can add up to 20 tags to a cluster. If the number of tags that are added to a cluster exceeds the upper limit, you must remove some of the tags before you add new tags.

#### Add a new tag or an existing tag when you create a cluster

1.

- 2. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- 3. Click Cluster Wizard in the Clusters section. In the Advanced Settings section of the Basic Settings step, find Tag and click + Add.

For more information about how to create a cluster, see Create a cluster.

4. In the **Create/Bind Tag** dialog box, select an existing tag key or enter a new tag key for Tag Key.

If you select an existing tag key, you can add an existing tag. If you enter a new tag key, you can create a tag and then add the tag.

Create/Bind Tag	1	×
* Tag Key:	y: TagKey1	
Tag Value:	e: Select an existing key or enter a new key Confirm	
TagKey2 ×		*
		*
		Confirm Cancel

**?** Note Each tag consists of a key-value pair. Tag keys and tag values have the following limits:

- A tag key cannot be empty and can be up to 128 characters in length. A tag value can be empty and can be up to 128 characters in length.
- Tag keys and tag values cannot start with aliyun or acs: and cannot contain http:// or https://.
- A tag key must be unique in a cluster. For example, the **city/shanghai** tag is added to a cluster. If the **city/newyork** tag is then added to the cluster, the **city/shanghai** tag is automatically removed.
- 5. Click Confirm.

Confirm the tag to be added.

6. Click Confirm.

The tag is added.

#### Use tags to search for clusters

On the Cluster Management page, search for a cluster by using a tag key or tag value.

1. Go to the Cluster Management page.

i.

- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the Cluster Management tab.

2. On the Cluster Management page, click Tag and select a tag key.

If you do not select a tag value, all clusters to which the tag key is added are displayed.

3. Click Search next to Tag.

All the clusters that meet the search condition are displayed.

#### Manage the tags of an existing cluster

You can add, remove, or view the tags of an existing cluster.

- 1. Go to the Cluster Overview page of your cluster.
  - i.
  - ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
  - iii. Click the Cluster Management tab.
  - iv. Find the cluster whose tags you want to manage and click **Details** in the Actions column.
- 2. On the Cluster Overview page, click Edit Tag in the Cluster Info section.
- 3. In the Edit Tag dialog box, perform the following steps:
  - Add a new tag.
    - a. Click Add Tag in the lower-left corner.
    - b. In the Edit Tag dialog box, specify Tag Key and Tag Value.
    - c. Click Confirm.
    - d. Click Confirm.
  - Add an existing tag.
    - a. Click Existing Tags in the lower-left corner.
    - b. Select an existing tag key and the tag value that matches the tag key.
    - c. Click Confirm.
  - Remove a tag.

Click the x icon next to a tag to remove this tag.

(?) Note After you remove a tag from a cluster, the system automatically deletes the tag if the tag is not added to other clusters.

# 3.6. Use resource groups

You can group your resources based on usage, permissions, and regions. This way, you can manage resources based on users and projects in a hierarchical manner. Each resource belongs to only one resource group. Resource groups do not affect the association among resources. In E-MapReduce (EMR), you can specify resource groups only for clusters and projects. This topic describes how to specify resource groups for resources and provides related examples.

#### Context

When you use resource groups, take note of the following points:

• A resource group can contain resources in different regions. For example, Resource Group A contains

clusters in the China (Hangzhou) region and clusters in the China (Shanghai) region.

- A resource group can contain different types of resources. For example, Resource Group A contains resources such as clusters, ECS instances, and projects.
- Clusters and projects that belong to the same Alibaba Cloud account and reside in the same region can be associated across resource groups. For example, your Alibaba Cloud account can access Resource Group A and Resource Group B. Resource Group A contains a project that resides in the China (Beijing) region. Resource Group B contains a cluster that also resides in the China (Beijing) region. In this case, you can run a job of the project in the cluster.
- RAM users have access to resource groups as authorized. For example, if you authorize a RAM user to manage all Alibaba Cloud resources in your Alibaba Cloud account, all resource groups in your Alibaba Cloud account are accessible to the RAM user.

#### Limits

- You can create and manage a resource group and authorize RAM users to access the resource group only in the Resource Management console. For more information, see What is Resource Management?
- In EMR, you can specify resource groups only for clusters and projects. When you create or scale out a cluster or move a cluster across resource groups, all nodes in the cluster are moved to the same resource group as the cluster. The following node resources can be managed in resource groups: Elastic Compute Service (ECS) instances, disks, images, Elastic Network Interfaces (ENIs), security groups, and key pairs.
- You cannot move resources across resource groups that belong to different Alibaba Cloud accounts.

Notice If you move a node resource to a different resource group, the cluster to which the node resource belongs remains in the original resource group. To manage resources and grant permissions in a centralized manner, we recommend that you manage node resources in the same resource group as the cluster.

#### Specify a resource group

A resource must belong to a resource group. If you do not specify a resource group when you create a resource, the resource is added to the default resource group. This section describes how to specify resource groups when you create a cluster and a project.

**?** Note When you create a resource, you can specify only a resource group to which you have access.

• Specify a resource group when you create a cluster

i.

- ii. In the top navigation bar, select the region in which you want to create a cluster.
- iii. Click Clust er Wiz ard.
- iv. After you configure the software and hardware parameters, go to the Basic Settings step and configure the parameters in the Basic Information section. Then, expand the Advanced Settings section and select an existing resource group from the Select Resource Group drop-down list.

If you want to create a resource group, click the **link** below the Select Resource Group drop-down list. For more information, see Create a resource group.

Onte For more information about how to create a cluster, see Create a cluster.

- Specify a resource group when you create a project
  - i.
  - ii. In the top navigation bar, select the region in which you want to create a project.
  - iii. Click the Data Platform tab.
  - iv. On the Data Platform tab, click Create Project in the upper-right corner.
  - v. In the **Create Project** dialog box, configure the **Project Name** and **Project Description** parameters and select an existing resource group from the **Select Resource Group** drop-down list.

If you want to create a resource group, click the **link** below the Select Resource Group drop-down list. For more information, see Create a resource group.

vi. Click Create.

For more information about how to create a project, see Manage projects.

#### Scenarios

You can use resource groups in the following scenarios:

- Add resources that serve different purposes to different resource groups for separate management. For more information, see Scenario 1: Add resources that serve different purposes to different resource groups.
- Configure an independent administrator for each resource group to manage users and user permissions by resource group. For more information, see Scenario 2: Manage users and user permissions by resource group.

**(?)** Note RAM users can view resources as authorized. A RAM user may have access only to some of the resource groups in an Alibaba Cloud account. In this case, if the RAM user selects All Resources in the top navigation bar, a message that indicates the RAM user is not authorized to access all resources appears.

# Scenario 1: Add resources that serve different purposes to different resource groups

You can add clusters to be used in a test environment and clusters to be used in a production environment to different resource groups. During tests, select a cluster from the resource group for the test environment. This prevents misoperations that may occur on production clusters. When you launch a service, select a cluster from the resource group for the production environment to ensure that the service works as expected.

1. Create two resource groups named Test Environment and Production Environment.

For more information, see Create a resource group.

2. Configure the same administrator for the **Test Environment** and **Production Environment** resource groups.

For more information, see Add RAM authorization.

3. Create two clusters named TestEnv1 and TestEnv2.

Specify the **Test Environment** resource group for the clusters.

4. Create two clusters named ProdEnv1 and ProdEnv2.

Specify the **Production Environment** resource group for the clusters.

- 5. Use the administrator account of the **Test Environment** and **Production Environment** resource groups to
- 6. Select one of the resource groups from the top navigation bar.

The clusters in the selected resource group are displayed in the cluster list. For example, if you select **Test Environment**, the TestEnv1 and TestEnv2 clusters are displayed.

#### Scenario 2: Manage users and user permissions by resource group

You can add clusters and projects of a department to an independent resource group and configure an independent administrator for the resource group. This way, you can separately manage the users and user permissions of each resource group. In the following procedure, resource groups for a **development department** and a **test department** are created.

1. Create two resource groups named **Develop Dept** and **Test Dept**.

For more information, see Create a resource group.

2. Configure different administrators for the **Develop Dept** and **Test Dept** resource groups.

For more information, see Add RAM authorization.

3. Create a cluster named **IT Cluster** and a project named **IT FlowProject**. Select the **Develop Dept** resource group for the cluster and project.

For more information about how to create a cluster, see Create a cluster. For more information about how to create a project, see Manage projects.

- 4. Create two clusters named FinanceCluster1 and FinanceCluster2. Select the Test Dept resource group for the clusters.
- 5. Use the administrator account of the Test Dept resource group to
- 6. Select the **Test Dept** resource group in the top navigation bar.

The FinanceCluster1 and FinanceCluster2 clusters are displayed in the cluster list.

# 3.7. Manage security groups

Security groups are an important means for network security isolation. Security groups are used to configure network access control for Elastic Compute Service (ECS) instances in a cluster. This topic describes how to add an ECS instance to a security group and add security group rules.

#### Context

When you create a cluster, you must create a security group or select an existing security group. You can add security group rules to control outbound and inbound network access for all ECS instances in the security group.

We recommend that you add ECS instances to different security groups and configure access control policies for each security group based on the use scenarios of the ECS instances. In this topic, the security groups that exist before you use E-MapReduce (EMR) are referred to as user security groups, and the security groups that are created when you create EMR clusters are referred to as EMR security groups.

#### Usage notes

- When you add security group rules, you must allow access only from specific IP addresses. To prevent attacks, you are not allowed to set the IP address to 0.0.0.0/0.
- When you configure inbound and outbound rules for applications, follow the principle of least privilege. You can allow access only from the current public IP address when you configure a security

group rule.

To obtain the current public IP address, visit <a href="http://myip.ipip.net/">http://myip.ipip.net/</a>.

#### Create a security group

When you create an EMR cluster, you can create a security group or select an existing security group.

Notice Do not use an advanced security group that is created in the ECS console.

#### Add an instance to a security group

- ? Note
  - An ECS instance of the classic network type must be added to a security group of the classic network type in the same region.
  - An ECS instance of the virtual private cloud (VPC) type must be added to a security group in the same VPC.
- 1. Go to the Cluster Overview page.

i.

- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the Cluster Management tab.
- iv. On the **Cluster Management** page, find your cluster and click **Details** in the Actions column.
- 2. In the Instance Info section, click Master Instance Group (MASTER) or Core Instance Group (CORE). Then, click the ID in the ECS ID column for an ECS instance.
- 3. On the Instances page, click the **Security Groups** tab.
- 4. On the Security Groups tab, click Add to Security Group.
- 5. In the Add to Security Group dialog box, select a security group from the Security Group dropdown list.

If you want to add the ECS instance to multiple security groups at the same time, click **Join Multiple Security Groups** after you select a security group. The security group is added to the box that appears. Then, perform the same operations to add other security groups to the box.

6. Click OK.

Repeat Step 2 to Step 6 until all the ECS instances in the EMR cluster are added to security groups.

#### Add a security group rule

1. Obtain the public IP address of your on-premises machine.

For security purposes, we recommend that you allow access only from the current public IP address when you configure a security group rule. To obtain your current public IP address, visit <a href="http://myip.ipip.net/">http://myip.ipip.net/</a>.

2. Go to the Cluster Overview page of your EMR cluster.

i.

ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.

- iii. Click the Cluster Management tab.
- iv. On the Cluster Management page, find your cluster and click Details in the Actions column.
- v. In the Network Info section of the Cluster Overview page, click the link of Security Group ID.
- 3. On the Security Group Rules page, add security group rules.

Configure the **Port Range** and **Authorization Object** parameters. Retain the default values of the other parameters. For more information, see Add a security group rule.

Parameter	Description	
Port Range	Set this parameter to the port that is used to access the ECS instance.	
Authorization Object	Set this parameter to the public IP address that is obtained in Step 1.	
	<b>Notice</b> To prevent attacks from external users, you are not allowed to set <b>Authorization Object</b> to <b>0.0.0.0/0</b> .	

4. Click Save.

# 3.8. Manage Metadata

### 3.8.1. Configure DLF as a metadatabase

This topic describes how to specify Data Lake Formation (DLF) as the metadatabase of a Hadoop cluster in the E-MapReduce (EMR) console.

#### Context

DLF is a fully managed service that is developed by Alibaba Cloud to help you easily build and manage a cloud-native data lake. DLF allows you to manage metadata and user permissions in a centralized manner, ingest data into data lakes with ease, and explore data in data lakes with a few clicks. For more information, see Overview.

DLF also allows you to connect the cloud-native data lake to various compute engines. This way, you can break down data silos and gain insight into the business value of data.

#### Prerequisites

The DLF service is activated in the DLF console.

#### Limits

- DLF is compatible only with Hive 2.x, Hive 3.x, Presto, and Spark SQL in EMR.
- You can specify DLF as the Hive metadatabase only for a Hadoop cluster in EMR V3.33.0 or a later minor version, or in EMR V4.5.0 or a later minor version.
- DLF can be used to manage the metadata of Hadoop clusters in EMR only in the China (Beijing), China (Shanghai), China (Hangzhou), and China (Shenzhen) regions.

#### Change the storage type of the metadata

You can change the value of a Hive parameter to change the storage type of the metadata.

Onte To migrate metadata from a metadatabase, submit a ticket.

- 1. Go to the Hive service page.
  - i.
  - ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
  - iii. Click the Cluster Management tab.
  - iv. Find your cluster and click **Det ails** in the **Actions** column.
  - v. In the left-side navigation pane, choose **Cluster Service > Hive**.
- 2. Change the value of the hive.imetastoreclient.factory.class parameter.
  - i. Click the **Configure** tab.
  - ii. In the Configuration Filter section, enter the hive.imetastoreclient.factory.class parameter in

the search box and click the o icon. Then, specify a value for the parameter based on the

storage type that you want to use.

 Use a built-in MySQL database or an independent ApsaraDB RDS for MySQL database as the metadatabase

Set the value of the hive.imetastoreclient.factory.class parameter to org.apache.hadoop.hive.ql.metadata.SessionHiveMetaStoreClientFactory.

Use DLF as the metadatabase

Set the value of the hive.imetastoreclient.factory.class parameter to com.aliyun.datalake.metastore.hive2.DlfMetaStoreClientFactory.

- 3. Save the configuration.
  - i. Click **Save** in the upper-right corner of the Hive service page.
  - ii. In the **Confirm Changes** dialog box, configure the Description parameter and click **OK**.
- 4. Restart the Hive metastore.
  - i. On the Hive service page, choose Actions > Restart Hive Metastore in the upper-right corner.
  - ii. In the Cluster Activities dialog box that appears, configure the parameters and click OK.
  - iii. In the **Confirm** message, click **OK**.

To view the task progress, click **History** in the upper-right corner.

# 3.8.2. Configure an independent ApsaraDB RDS for MySQL database

This topic describes how to configure an independent ApsaraDB RDS for MySQL database as the metadatabase of a Hadoop cluster in the E-MapReduce (EMR) console.

#### Prerequisites

An ApsaraDB RDS for MySQL instance is purchased. For more information, see Create an ApsaraDB RDS for MySQL instance.

(?) Note In this topic, an ApsaraDB RDS for MySQL instance that runs MySQL 5.7 is used.

#### Limits

When you create an ApsaraDB RDS for MySQL instance, you must set **Database Engine** to **MySQL 5.7** and **Edition** to **High-availability**.

#### Procedure

1. Step 1: Prepare a metadatabase

Prepare a metadatabase.

2. Step 2: Create a cluster

Create a cluster in the EMR console and associate the cluster with the metadatabase.

3. (Optional)Step 3: Initialize the metastore service

Initialize the metastore service based on the Hive version.

Notice If you created a Hadoop cluster, perform this step.

#### Step 1: Prepare a metadatabase

1. Create a database named hivemeta.

For more information, see the "Create a database" section of the Create accounts and databases for an ApsaraDB RDS for MySQL instance topic.

Create Database		
* Database Name		
hivemeta		
The database name must b	e 2 to 64 cha	aracters in length and can
* Supported Character S	Set	
utf8	~	~
Authorized Account		
Unauthorized Account (d	lefault) 🗸 🗸	<ul> <li>Create Account</li> </ul>
Description		
The description must be	0 to 256 char	aracters in length.
The description must be 0 t	to 256 charac	ctors in longth
The description must be of		cters in length.
Create Cancel		

2. Create an account and grant read and write permissions to the account.

#### ♥ Notice

For more information about how to create an account, see the "Create an account" section of the Create accounts and databases for an ApsaraDB RDS for MySQL instance topic.

* Database Account			
hiveuser			
An account name must be 1 to underscores (_). It must start		rs in length and can contain lowercase letters, digits, and and end with a letter or digit.	
* Account Type 🕖			
O Privileged Account 💿	Standard A	Account	
Authorized Databases:			
Unauthorized Databases		Authorized Databases:	
Not Found	>	hivemeta Read/Write ORead-only ODL O	nly
0 Item		<ul> <li>↓</li> <li>1 Item</li> </ul>	•
* Password			
•••••			
		a length and must contain at least three of the following types: ind special characters. Special characters include ! @ # \$ % ^	
Description			
The description must be 0 to 2	256 character		56 5

Record the username and password of the account that you created. The username and password are required when you create a cluster. For information about how to create a cluster, see Step 2: Create a cluster.

3. Obtain the internal endpoint of the database.

- i. Configure an IP address whitelist. For more information, see Use a database client or the CLI to connect to an ApsaraDB RDS for MySQL instance.
- ii. In the left-side navigation pane of the instance details page, click **Database Connection**.
- iii. On the **Database Connection** page, click the Copy icon on the right of the internal endpoint to copy the internal endpoint.

<	rm-bp1tb7otw (Running) & Back to Instances			
Basic Information	Database Connection			
Accounts				
Databases	Instance Connection Read/Write Splitting			
Backup and Restorati	Database Connection			
Database Connect	Network Type: VPC (VPC:vpc-bp1im0n59k8ibnjvr3w9e) 📀			
Database Proxy	Internal Endpoint rm- mysql.rds.aliyuncs.com			

Record the internal endpoint. This endpoint is required when you create a cluster. For more information about how to create a cluster, see Step 2: Create a cluster.

#### Step 2: Create a cluster

In the **Basic Settings** step, configure the parameters described in the following table. For more information about the configurations of other parameters, see **Create a cluster**.

Parameter	Description
Cluster Name	The name of the cluster. The name must be 1 to 64 characters in length and can contain only letters, digits, hyphens (-), and underscores (_).
Туре	Select Independent ApsaraDB RDS for MySQL.
Connection URL	<ul> <li>Set Connection URL to a value in the format of <i>jdbc:mysql://rm-xxxxx.mysql.rds.aliyuncs.com/<database name="">?</database></i> <i>createDatabaselfNotExist=true&amp;characterEncoding=UTF-8.</i></li> <li><i>rm-xxxxx.mysql.rds.aliyuncs.com</i> indicates the internal endpoint obtained in Step 1: Prepare a metadatabase.</li> <li><i><database name=""></database></i> indicates the name of the database created in Step 1: Prepare a metadatabase.</li> </ul>
Database Username	Set this parameter to the username of the account created in Step 1: Prepare a metadatabase.
Database Password	Set this parameter to the password of the account created in Step 1: Prepare a metadatabase.

#### Step 3: Initialize the metastore service

**Notice** For a Hadoop cluster, perform the following steps to initialize the metastore service based on the Hive version.

- 1. Go to the Cluster Overview page.
  - i.
  - ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
  - iii. Click the Cluster Management tab.
  - iv. On the Cluster Management page, find your cluster and click Details in the Actions column.
- 2. In the **Software Info** section of the **Cluster Overview** page, check the Hive version and initialize the metastore service of Hive.
  - If the Hive version is 2.3.X, perform the following steps to initialize the service:
    - a. Enter the specified directory.

cd /usr/lib/hive-current/scripts/metastore/upgrade/mysql/

b. Log on to the ApsaraDB RDS for MySQL database.

mysql -h {Private or public IP address of the ApsaraDB RDS for MySQL database}  $-u{U$  sername to log on to the ApsaraDB RDS for MySQL database}  $-p{Password to log on to the ApsaraDB RDS for MySQL database}$ 

c. Run the following command on the command line of the ApsaraDB RDS for MySQL database:

use {Name of the ApsaraDB RDS for MySQL database}; source /usr/lib/hive-current/scripts/metastore/upgrade/mysql/hive-schema-2.3.0.mysq l.sql;

#### Parameter description:

- Username to log on to the ApsaraDB RDS for MySQL database} indicates the username of the account created in Step 1: Prepare a metadatabase.
- Password to log on to the ApsaraDB RDS for MySQL database} indicates the password of the account created in Step 1: Prepare a metadatabase.
- Name of the ApsaraDB RDS for MySQL database indicates the name of the database created in Step 1: Prepare a metadatabase.
- If the Hive version is not 2.3.X, perform the following steps to initialize the service:
  - a. Log on to the master node of the cluster by using SSH. For more information, see Log on to a cluster.
  - b. Run the following command to switch to the hadoop user:

su hadoop

c. Run the following command to log on to the ApsaraDB RDS for MySQL database:

schematool -initSchema -dbType mysql

After the metastore service is initialized, you can use the created ApsaraDB RDS for MySQL database as the Hive metadatabase.

**Note** Before the metastore service is initialized, exceptions may occur in the MetaStore and HiveServer2 processes for Hive and the ThriftServer process for Spark. After the metastore service is initialized, the processes are recovered.

#### FAQ

- What do I do if the "Failed to get schema version" error is reported when the MetaStore client is initialized?
- What do I do if Hive metadata contains Chinese characters, such as Chinese characters in column comments and partition names?

# 3.9. Manage bootstrap actions

E-MapReduce (EMR) allows you to use bootstrap actions to install third-party software and modify the runtime environment of your clusters. This topic describes how to add, edit, clone, and remove bootstrap actions.

#### Context

Bootstrap actions can help you perform many operations that are not supported by EMR clusters. For example, you can use bootstrap actions to perform the following operations:

- Use Yellowdog Updater, Modified (YUM) to install software whose installation package is available.
- Download public software from the Internet.
- Read your dat a from Object Storage Service (OSS).
- Install and run services, such as Flink or Impala.

#### Limits

- You can add a maximum of 10 bootstrap actions. Bootstrap actions are performed in the order that you specified.
- By default, the script that you specified is run by the root user. You can add the **su hadoop** command to the script. This way, the script is run by the hadoop user.

#### Add a bootstrap action

You can use one of the following methods to add a bootstrap action:

- Method 1: Add a bootstrap action when you create a cluster.
  - i. Go to the Cluster Management page.
    - a.
    - b. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
    - c. Click the **Cluster Management** tab.
  - ii. Click Cluster Wizard in the upper-right corner.
  - iii. In the Advanced Settings section of the Basic Settings step, find Bootstrap Actions and click
     + Add.
  - iv. In the Add Bootstrap Actions dialog box, configure the following parameters.

Parameter	Description
Name	The name of the bootstrap action that you want to add.
Script Path	The OSS path where the script file is located. You must specify this parameter in the <i>oss://**/*.sh</i> format.
Parameter	The parameter of the bootstrap action script. The parameter is used to specify the value of the variable that is referenced in the script.
Target Nodes	<ul> <li>Cluster: The bootstrap action is applicable to the entire cluster.</li> <li>Host Group: The bootstrap action is applicable only to a specific machine group.</li> <li>You can select Core Instance Group, Master Instance Group, or an existing machine group.</li> </ul>
Execution At	<ul> <li>Before Component Startup: The system runs the script before the deployed components are started.</li> <li>After Component Startup: The system runs the script after the deployed components are started.</li> </ul>
Execution Failure Policy	<ul> <li>Proceed: If the script fails to be run, the system continues to run the next script.</li> <li>Stop: If the script fails to run, the system stops running scripts.</li> </ul>

#### v. Click OK.

Examples of bootstrap actions are provided in Examples for your reference.

**?** Note The added bootstrap action may fail to be performed. However, the failure does not affect the creation of the cluster.

After the cluster is created, you can check the value of **Bootstrap Actions/EMR Version** in the Cluster Info section of the **Cluster Overview** page to determine whether an exception occurs. If an exception occurs, you can log on to each node of the cluster and view the operational log. Operational logs are stored in the */var/log/bootstrap-actions* directory.

- Method 2: Add a bootstrap action after you create a cluster.
  - i. Go to the Bootstrap Actions page.

a.

- b. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- c. Click the Cluster Management tab.
- d. On the **Cluster Management** page, find your cluster and click **Details** in the Actions column.

- e. In the left-side navigation pane, click **Bootstrap Actions**.
- ii. On the Bootstrap Actions page, click Add Bootstrap Actions.
- iii. In the Add Bootstrap Actions dialog box, configure the following parameters.

Parameter	Description
Name	The name of the bootstrap action that you want to add.
Script Path	The OSS path where the script file is located. You must specify this parameter in the <i>oss://**/*.sh</i> format.
Parameter	The parameter of the bootstrap action script. The parameter is used to specify the value of the variable that is referenced in the script.
Target Nodes	<ul> <li>Cluster: The bootstrap action is applicable to the entire cluster.</li> <li>Host Group: The bootstrap action is applicable only to a specific machine group.</li> <li>You can select Core Instance Group, Master Instance Group, or an existing machine group.</li> </ul>
Execution At	<ul> <li>Before Component Startup: The system runs the script before the deployed components are started.</li> <li>After Component Startup: The system runs the script after the deployed components are started.</li> </ul>
Execution Failure Policy	<ul> <li>Proceed: If the script fails to be run, the system continues to run the next script.</li> <li>Stop: If the script fails to run, the system stops running scripts.</li> </ul>

#### iv. Click OK.

Examples of bootstrap actions are provided in Examples for your reference.

#### Edit a bootstrap action

1. Go to the Bootstrap Actions page.

i.

- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the Cluster Management tab.
- iv. On the **Cluster Management** page, find your cluster and click **Details** in the Actions column.
- v. In the left-side navigation pane, click **Bootstrap Actions**.
- 2. On the **Bootstrap Actions** page, find the bootstrap action that you want to edit and click **Edit** in the Actions column.

- 3. In the Edit Bootstrap Action dialog box, modify the configuration items based on your business requirements. All configuration items can be modified.
- 4. Click OK.

#### Clone a bootstrap action

1. Go to the Bootstrap Actions page.

i.

- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the Cluster Management tab.
- iv. On the Cluster Management page, find your cluster and click Details in the Actions column.
- v. In the left-side navigation pane, click **Bootstrap Actions**.
- 2. On the **Bootstrap Actions** page, find the bootstrap action that you want to clone and click **Clone** in the Actions column.
- 3. In the **Clone Bootstrap Action** dialog box, modify the configuration items based on your business requirements. All configuration items can be modified.
- 4. Click OK.

#### Remove a bootstrap action

1. Go to the Bootstrap Actions page.

i.

- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the Cluster Management tab.
- iv. On the **Cluster Management** page, find your cluster and click **Details** in the Actions column.
- v. In the left-side navigation pane, click **Bootstrap Actions**.
- 2. On the **Bootstrap Actions** page, find the bootstrap action that you want to remove and click **Delete** in the Actions column.
- 3. In the Delete Bootstrap Action message, click OK.

#### Examples

When you add a bootstrap action, you must specify a bootstrap action name, the OSS path of a script file, and a parameter of the script based on your business requirements. When the bootstrap action is performed, each node downloads the script from the specified OSS path and runs the script directly or based on optional parameters. This section provides two examples:

• Example 1

You can specify the file that you want to download from OSS in the script. In this example, you can use the following script to download the <myfile>.tar.gz file from the *oss://<yourbucket>/* directory and decompress the file to the */<yourdir>* directory on your computer.

Notice The OSS endpoint in the script can be an internal, public, or VPC endpoint. If you use the classic network, you must specify an internal endpoint. For example, the internal endpoint that corresponds to the China (Hangzhou) region is *oss-cn-hangzhou-internal.aliyuncs.com*. If you use a virtual private cloud (VPC), you must specify an endpoint that you can access from the VPC. For example, the endpoint that corresponds to the China (Hangzhou) region is *vpc100-oss-cn-hangzhou ou.aliyuncs.com*.

```
#!/bin/bash
osscmd --id=<yourid> --key=<yourkey> --host=oss-cn-hangzhou-internal.aliyuncs.com get oss:/
/<yourbucket>/<myfile>.tar.gz ./<myfile>.tar.gz
mkdir -p /<yourdir>
tar -zxvf <myfile>.tar.gz -C /<yourdir>
```

#### • Example 2

You can use YUM to install additional system software. For example, you can use the following script to install *ld-linux.so.2*:

```
#!/bin/bash
yum install -y ld-linux.so.2
```

# 3.10. Cluster scripts

After you create a cluster, especially a subscription cluster, you can run scripts on multiple nodes at the same time to meet your business requirements, such as to install third-party software or modify the running environment of the cluster.

#### Context

You can take bootstrap actions to initialize on-demand clusters. A cluster script is similar to a bootstrap action. After a cluster is created, you can use the cluster script feature to install software and services previously unavailable to your cluster. Example:

- Use YUM to install software already provided.
- Download public software from the Internet.
- Install software to read your data from Object Storage Service (OSS).
- Install and run a service, such as Flink or Impala. However, the script you want to write is more complex.

#### Prerequisites

- A cluster is created. For more information, see Create a cluster.
- The cluster is in the Idle or Running state. Scripts cannot run on clusters in other states.
- Cluster scripts are developed or obtained and uploaded to OSS. For more information about the cluster scripts, see Examples.

#### Precautions

- Only one cluster script can run on a cluster at a specific point in time. You cannot submit another cluster script if one is already in progress. You can retain a maximum of 10 cluster script records for each cluster. If more than 10 records exist, you must delete the previous records before you create new cluster scripts.
- A cluster script may succeed on some nodes, but fail on others. For example, if you restart a node, the

script may fail to run. After you resolve the issue, you can run the cluster script on the failed nodes again. After you scale out a cluster, you can also run cluster scripts on the added nodes.

• You can use the cluster script feature to download scripts from OSS to a specific node and run these scripts. If the scripts fail to be run, you can log on to the node to check the operational log. The operational log for each node is stored in the */var/log/cluster-scripts/clusterScriptId* path. If the cluster is configured with an OSS log storage directory, the operational log is also stored in the *osslogp ath/cluster-scripts/cluster-scr* 

#### Create a cluster script

1. Go to the Cluster Overview page.

- i.
- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the Cluster Management tab.
- iv. On the Cluster Management page, find your cluster and click Details in the Actions column.
- 2. In the left-side navigation pane of the page that appears, click Cluster Scripts.
- 3. On the Cluster Scripts page, click Create and Run in the upper-right corner.
- 4. In the Create Script dialog box, specify Name, Script, and Target Nodes.

(?) Note We recommend that you test the cluster script feature on a single node before you use the feature on the entire cluster.

#### 5. Click OK.

After a cluster script is created, it is displayed in the cluster script list and is in the Running state.

- Click Refresh in the upper-right corner to update the status of the cluster script.
- Click Details in the Actions column to view the running status of the cluster script.
- Click Delete in the Actions column to delete the cluster script that you created.

#### Examples

Similar to a bootstrap action, you can specify the file that you want to download from OSS in the script. In the following example, the myfile.tar.gz file in the *oss://yourbucket* directory is downloaded to your computer and decompressed to the */yourdir* directory.

```
#!/bin/bash
osscmd --id=<yourAccessKeyId> --key=<yourAccessKeySecret> --host=oss-cn-hangzhou-internal.ali
yuncs.com get oss://<yourBucketName>/<yourFile>.tar.gz ./<yourFile>.tar.gz
mkdir -p /<yourDir>
tar -zxvf <yourFile>.tar.gz -C /<yourDir>
```

(?) Note The specified OSS address can be an internal, public, or VPC endpoint. If the node is of the classic network type, you must specify an internal endpoint. For example, the internal endpoint for a node that resides in the China (Hangzhou) region is *oss-cn-hangzhou-internal.aliyuncs.com*. If the node is of the VPC type, you must specify a domain name that you can access from the VPC. For example, the domain name for a node that resides in the China (Hangzhou) region is *vpc100-oss-cn-hangzhou.aliyuncs.com*.

You can also use YUM to install additional system software packages, such as *ld-linux.so.2*.

```
#!/bin/bash
yum install -y ld-linux.so.2
```

By default, the root account is used to run specified scripts on a cluster. You can run the su hadoop command in the script to switch to the hadoop user.

# 3.11. Manage cluster resources

# 3.11.1. Cluster resource overview

The cluster resource management feature of E-MapReduce is ideal for scenarios that involve large-scale clusters with many tenants. Currently, you can only manage resources of E-MapReduce Hadoop clusters.

#### Background

The cluster resource management feature provides the following benefits:

- Different departments or users can use different resource queues. This facilitates resource isolation.
- Resource queues are elastic, which improves the use efficiency of clusters.

#### **Resource management**

Currently, you can use Capacity Scheduler or Fair Scheduler to manage E-MapReduce cluster resources.

- Capacity Scheduler: For more information, see Use Capacity Scheduler.
- Fair Scheduler: For more information, see Use Fair Scheduler.

**Note** You can also visit the official websites of Capacity Scheduler and Fair Scheduler to learn more about the components.

## 3.11.2. Use Capacity Scheduler

#### Capacity Scheduler

Capacity Scheduler is a built-in scheduler in Apache YARN. The YARN service that is deployed in an E-MapReduce (EMR) cluster uses Capacity Scheduler as the default scheduler. Capacity Scheduler is a multi-tenant, hierarchical resource scheduler. Resources used in each child queue are allocated based on the configured capacity.

#### Prerequisites

An EMR Hadoop cluster is created. For more information, see Create a cluster.

#### Precautions

After you turn on Enable Resource Queue, you can no longer configure cluster resources on the **capacity-scheduler** tab in the Service Configuration section of the **Configure** tab on the YARN service page. Existing configurations are synchronized to the **Cluster Resources** page. If you want to configure cluster resources on the **Configure** tab of the YARN service page, turn off Enable Resource Queue on the **Cluster Resources** page.

#### **Configure Capacity Scheduler**

1. Go to the Cluster Resources page.

i.

- ii. In the top navigation bar, select the region where you want to create a cluster. The region of a cluster cannot be changed after the cluster is created.
- iii. Click the Cluster Management tab.
- iv. On the Cluster Management page, find your cluster and click Details in the Actions column.
- v. In the left-side navigation pane of the page that appears, click **Cluster Resources**.
- 2. Enable Capacity Scheduler.
  - i. On the Cluster Resources page, turn on Enable Resource Queue.
  - ii. Select Capacity Scheduler for Queue Type.

**Note** If you use the cluster resource management feature for the first time, the Queue Type parameter is set to **Capacity Scheduler** by default.

#### iii. Click Save.

- 3. In the upper-right corner of the Cluster Resources page, click Queue Settings.
- 4. Configure queue information on the **Queue Settings** tab.
  - Find your queue and click Edit in the Actions column to modify resource queue information.
  - Find your queue and choose **More > Create Child Queue** in the Actions column to create a child queue.

You cannot create a child queue for the default queue.

**root** is a level-1 queue. It is the parent queue of all other queues and manages all resources of YARN. By default, only the default queue is available within the root queue.

#### ○ Notice

- The sum of the Capacity values set for all child queues at the same level within the same parent queue must be 100, in percentage. For example, two child queues default and department are available within the root queue. The sum of the Capacity values set for the default and department queues must be 100. Child queues market and dev are available within the department queue. The sum of the Capacity values set for the market and dev queues must also be 100.
- If you do not specify a queue when applications are running, jobs are submitted to the default queue.
- After you create a level-2 queue within the root queue, you need only to click **Deploy** to make the configuration take effect.
- After you create or modify a level-3 queue within the root queue, you must restart ResourceManager to make the configuration take effect.
- You must restart ResourceManager after you modify the name of a queue.

#### Switch the scheduler type

After Enable Resource Queue is turned on, you can perform the following steps to switch the scheduler type:

**Notice** After the switchover is complete, you must restart ResourceManager to make the configuration take effect.

- 1. In the upper-right corner of the **Cluster Resources** page, click **Select Scheduler**.
- 2. Select the required scheduler for Queue Type.
- 3. Click Save.
- 4. Restart ResourceManager.
  - i. In the upper-right corner of the Cluster Resources page, choose Actions > Restart ResourceManager.
  - ii. In the Cluster Activities dialog box, configure the parameters and click OK.
  - iii. In the **Confirm** message, click **OK**.

When a success message appears, the scheduler type is switched.

#### Disable the cluster resource management feature

(?) Note After you disable the cluster resource management feature, you cannot perform operations on the Cluster Resources page. If you want to use the cluster resource management feature again, turn on Enable Resource Queue on the Cluster Resources page or configure the xml-direct-to-file-content parameter on the capacity-scheduler tab in the Service Configuration section of the Configure tab on the YARN service page.

- 1. On the Cluster Resources page, turn off Enable Resource Queue.
- 2. In the Disable Resource Queue dialog box, click OK.

#### Submit a job

- If you do not specify a queue when applications are running, jobs are submitted to the default queue.
- You must specify a child queue. Tasks cannot be submitted to the parent queue.
- You must use the **mapreduce.job.queuename** parameter to specify the queue to which you want to submit a job. Example:

`hadoop jar /usr/lib/hadoop-current/share/hadoop/mapreduce/hadoop-mapreduce-examples-2.8.5.
jar pi -Dmapreduce.job.queuename=test 2 2`

## 3.11.3. Use Fair Scheduler

#### Fair Scheduler

Fair Scheduler is a built-in scheduler in Apache YARN. It is used to fairly allocate resources to applications running on YARN. The resources used in each queue are allocated based on the configured weight.

#### Prerequisites

An EMR Hadoop cluster is created. For more information, see Create a cluster.

#### Precautions

After you turn on Enable Resource Queue, you can no longer configure cluster resources on the **fair-scheduler** tab in the Service Configuration section of the **Configure** tab on the YARN service page. Existing configurations are synchronized to the **Cluster Resources** page. If you want to configure cluster resources on the **Configure** tab of the YARN service page, turn off Enable Resource Queue on the **Cluster Resources** page.

#### **Configure Fair Scheduler**

1. Go to the Cluster Resources page.

i.

- ii. In the top navigation bar, select the region where you want to create a cluster. The region of a cluster cannot be changed after the cluster is created.
- iii. Click the Cluster Management tab.
- iv. On the Cluster Management page, find your cluster and click Details in the Actions column.
- v. In the left-side navigation pane of the page that appears, click **Cluster Resources**.
- 2. Enable Fair Scheduler.
  - i. On the Cluster Resources page, turn on Enable Resource Queue.
  - ii. Select Fair Scheduler for Queue Type.

**?** Note If you use the cluster resource management feature for the first time, the Queue Type parameter is set to **Capacity Scheduler** by default.

#### iii. Click Save.

- 3. In the upper-right corner of the Cluster Resources page, click Queue Settings.
- 4. Configure queue information on the Queue Settings tab.
  - Find your queue and click Edit in the Actions column to modify resource queue information.
  - Find your queue and choose **More > Create Child Queue** in the Actions column to create a child queue.

You cannot create a child queue for the default queue.

**root** is a level-1 queue. It is the parent queue of all other queues and manages all resources of YARN. By default, only the default queue is available within the root queue.

Notice Parameters of a parent queue have a higher priority than those of its child queues because the queue configurations are nested. If the resource usage configured for a child queue is higher than that of the parent queue, the scheduler allocates resources to the child queue based on the parameter settings of the parent queue.

- Queue Name is required. A queue name cannot contain periods (.).
- Weight is required. If the weight is reached, it is considered that resources are fairly allocated to queues. Weights take effect on level-2, level-3, and lower-level queues. For example, a parent queue has a child queue with a weight of 2 and another child queue with a weight of 1. If tasks are running in the child queues and the resource allocation ratio reaches 2:1, resources are fairly allocated.
- Maximum Resources specifies the maximum number of vCores and the maximum memory space that can be allocated to a queue. The values of this parameter must be greater than the values of Minimum Resources but less than the resource scale that the YARN service can provide. If the values of the Maximum Resources parameter are greater than the resource scale, the resource scale takes effect for the queue. For example, the number of vCores that the YARN service can provide is 16, but the value of vCores for Maximum Resources is 20. In this case, 16 vCores are allocated to the queue.
- If you do not specify a queue when applications are running, jobs are submitted to the default queue.
- If you do not restart ResourceManager after you modify the name of a child queue, tasks can still be submitted to the original queue. However, the queue configuration is no longer displayed in the EMR console. After you restart ResourceManager, the original queue becomes unavailable.
- After you delete a queue that is not a level-2 queue, click Deploy to make the modification take effect. After you delete a level-2 queue, choose Actions > Restart ResourceManager in the upper-right corner of the Cluster Resources page to make the modification take effect.

#### Switch the scheduler type

After Enable Resource Queue is turned on, you can perform the following steps to switch the scheduler type:

**Notice** After the switchover is complete, you must restart ResourceManager to make the configuration take effect.

- 1. In the upper-right corner of the **Cluster Resources** page, click **Select Scheduler**.
- 2. Select the required scheduler for Queue Type.
- 3. Click Save.
- 4. Restart ResourceManager.
  - i. In the upper-right corner of the Cluster Resources page, choose Actions > Restart ResourceManager.
  - ii. In the Cluster Activities dialog box, configure the parameters and click OK.
  - iii. In the **Confirm** message, click **OK**.

When a success message appears, the scheduler type is switched.

#### Disable the cluster resource management feature

**Note** After you disable the cluster resource management feature, you cannot perform operations on the Cluster Resources page. If you want to use the cluster resource management feature again, turn on Enable Resource Queue on the Cluster Resources page or configure the xml-direct-to-file-content parameter on the fair-scheduler tab in the Service Configuration section of the Configure tab on the YARN service page.

#### 1. On the Cluster Resources page, turn off Enable Resource Queue.

2. In the Disable Resource Queue dialog box, click OK.

#### Submit a job

You must use the **mapreduce.job.queuename** parameter to specify the queue to which you want to submit a job. Example:

```
`hadoop jar /usr/lib/hadoop-current/share/hadoop/mapreduce/hadoop-mapreduce-examples-2.8.5.ja
r pi -Dmapreduce.job.queuename=test 2 2`
```

# 3.12. Cluster management page

This topic describes the services and access address of the cluster.

Software	Service	Visit address		
	yarn resourcemanager	masternode1_private_ip:8088,ma sternode2_private_ip:8088		
Hadoop	jobhistory	masternode1_private_ip:19888		
Παυσομ	timeline server	masternode1_private_ip:8188		
	hdfs	<i>masternode1_private_ip:50070,m asternode2_private_ip:50070</i>		
Spark	spark ui	masternode1_private_ip:4040		
	history	masternode1_private_ip:18080		
Tez	tez-ui	<i>masternode1_private_ip:8090/tez -ui2</i>		
Hue	hue	masternode1_private_ip:8888		
Zeppelin	zeppelin	masternode1_private_ip:8080		
Hbase	hbase	masternode1_private_ip:16010		
Presto	presto	masternode1_private_ip:9090		
Oozie	oozie	masternode1_private_ip:11000		

Software	Service	Visit address
Ganglia	ganglia	<i>masternode1_private_ip:9292/ga nglia</i>

# 3.13. View the cluster list and cluster details

This topic describes how to view the basic information about clusters that belong to your Alibaba Cloud account and the details about a cluster.

#### Prerequisites

An E-MapReduce (EMR) cluster is created. For more information, see Create a cluster.

#### View the cluster list

1.

- 2. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- 3. Click the Cluster Management tab.

The Cluster Management tab displays the basic information about all your clusters and the operations that you can perform on each cluster.

	0 T	<b>C</b>		Billing Meth	a	
Cluster ID/Name	Cluster Type	Status	Created At/Time Elapsed	od	Cluster Tag	Actions
C-7F570B31E6	Hadoop	O Idle	2021-12-14 10:37:43 1 Days 4 Hours 29 Minutes 7 Se conds	Subscription	۲	■ Details   ③ Renewal   More▼
C-1A91BCF	🕼 Hadoop	9 Idle	2021-12-15 10:53:50 4 Hours 44 Minutes 51 Seconds	Pay-As-You-Go	•	■ Details   More ▼
Parameter	Descript	tion				
Cluster ID/Name	The ID a	and name of	the cluster.			
Cluster Type	The typ	The type of the cluster.				
Status	The sta	The status of the cluster. For more information, see Status list.				list.
Created At/Time Elapsed	• Time creat					
Billing Method	The bill	The billing method of the cluster.				
Cluster Tag	The tag	The tag that is added to the cluster.				

Parameter	Description
	The operations that you can perform on the cluster. The following operations are included:
	• <b>Details</b> : View the details about the cluster on the Cluster Overview page.
	• <b>Renewal</b> : Renew the cluster. You can renew only subscription clusters.
	• More:
Actions	<ul> <li>CloudMonitor: Enable the system to monitor CPU idleness, memory capacity, and disk capacity of the EMR cluster.</li> </ul>
	• Auto Renewal: Enable automatic renewal for the subscription cluster.
	<ul> <li>Scale Out: Scale out the cluster. For more information, see Scale out a cluster.</li> </ul>
	<ul> <li>Release: Release the cluster. You can release only pay-as-you-go clusters. For more information, see Release a cluster.</li> </ul>
	Restart : Restart the cluster.

#### View the details about the cluster

1.

- 2. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- 3. Click the Cluster Management tab.
- 4. On the **Cluster Management** page, find your cluster and click **Details** in the Actions column.

The Cluster Overview page displays details about the cluster in the following sections: Cluster Info, Software Info, Network Info, and Instance Info.

• Cluster Info

Cluster Info			
Cluster Name: EMR-3-38 🔗	Cluster ID: C-10C8A31A	Region: cn-hangzhou	Status: 🕒 Idle
I/O Optimization: Yes	High Availability: No	Security Mode: Standard	Total number of nodes in the cluster: 3
Start Time: 2021-12-09 16:14:23	Billing Method: Subscription	Time Elapsed: 5 Days 22 Hours 54 Minutes 16 Seconds	Auto Renewal: Disable 🕦
Notification: No nodes will expire in the next 7 days.	Hive Metadata: Built-in MySQL	Bootstrap Actions/EMR Version: Standard	ECS Role: AliyunECSInstanceForEMRRole
Tag: - 🔗 Edit Tag			
1			

Parameter	Description
Cluster Name	The name of the cluster.
Cluster ID	The ID of the cluster.
Region	The region where the cluster resides.
Status	The status of the cluster. For more information, see Status list.
I/O Optimization	Specifies whether I/O optimization is enabled.
High Availability	Specifies whether high availability is enabled.

Parameter	Description
Security Mode	The security mode. The software of the cluster starts in Kerberos security mode. For more information, see Introduction to Kerberos.
Total number of nodes in the cluster	The total number of nodes in the cluster.
Start Time	The time when the cluster was created.
Billing Method	The billing method of the cluster.
Time Elapsed	The time elapsed from the time when the cluster was created to the current time.
Hive Metadata	The metadata storage center.
Bootstrap Actions/EMR Version	The information about custom scripts and software configuration.
ECS Role	The ECS application role. You can assign an application role to a cluster. Then, EMR applies for a temporary AccessKey pair when applications that run on the compute nodes of the cluster access other Alibaba Cloud services, such as Object Storage Service (OSS). You do not need to manually enter an AccessKey pair. You can grant the access permissions of the application role on specific Alibaba Cloud services based on your business requirements.
Tag	The tag that is added to the cluster. For more information, see Manage and use tags.

#### • Software Info

Software Inf	0 🕜						
EMR Version: E	MR-3.38.2						
Cluster Type: H	ladoop						
Software Info:	HDFS 2.8.5	YARN 2.8	8.5	Hive 2.3.9	Ganglia	3.7.2	Spark 2.4.8
	Hue 4.9.0	Tez 0.9.2	Sqo	op 1.4.7	RANGER	1.2.0	Hudi 0.9.0
	Knox 1.1.0	OpenLDA	P 2.4.4	44 Bigb	oot 3.8.0	Delta	Lake 0.6.1
	SmartData 3	.8.0 Icek	erg 0.	12.0			

Parameter	Description
EMR Version	The EMR version of the cluster.
Cluster Type	The type of the cluster.
Software Info	All the applications that are installed in the cluster and the application versions, such as HDFS 2.8.5, Hive 2.3.9, and Spark 2.4.8.
#### • Network Info

Parameter	Description
Zone ID	The region where the cluster resides.
Network Type	The network type of the cluster. By default, the VPC network type is selected.
Security Group ID	The ID of the security group to which the cluster is added.
VPC/VSwitch	The ID of the VPC where the cluster resides and the ID of the vSwitch to which the cluster belongs.

#### • Instance Info

Instance Info 🕑					
Master Instance Group (MASTER) (* Subscription	ECS ID	Deployment Stat us	Public IP Addres s	Internal IP Addre ss	Created At
<ul> <li>ECS Instance Type: ecs.g5.xlarge</li> <li>Instances: 1</li> </ul>	i-bp10mre6d9p0tz9je4iv 🗗 📮	Normal	47.96.	192.168	2021-12-09 16:16:30
♦ CPU: 4vCores ♦ Memory: 16G8 ♦ Data Disk Type: 80G8 ESSD Disk*1	Show all nodes 🖉			Items per Page::8Items	; < 1 > Total Items:
Core Instance Group (CORE) Subscription + ECS Instance Type: ecs.g5.xlarge + Instances: 2 + CPU: 4vCores + Memory: 16G8 + Data Disk Type: 80G8 ESSD Disk*4					
Parameter	Description				
ECS Instance Type	The type of the ECS instances.				
Instances	The current number of nodes in the cluster and the number of nodes that you requested. Theoretically, the two values must be the same. However, the current number is less than the requested number during the cluster creation process.				
СРИ	The number of vCPUs for a single instance.				
Memory	The memory capacity of a single instance.				
Data Disk Type	The data disk type and data disk capacity of a single instance.				

Parameter	Description
	The information about each ECS instance in a master node group, including the following parameters:
	• ECS ID: the ID of the ECS instance.
ECS instances	<ul> <li>Deployment Status: the deployment status of the ECS instance. Valid values: Initializing, Normal, and Scaling out.</li> </ul>
	Public IP Address: the public IP address of the ECS instance.
	<ul> <li>Internal IP Address: the private IP address of the ECS instance, which can be accessed by all nodes in the cluster.</li> </ul>
	• <b>Created At</b> : the time when the ECS instance was created.

# 4.Manage services4.1. View the status of services

You can view the status of all the services, such as HDFS and YARN, that are deployed in a cluster on the Cluster Management page. This topic describes how to view the status of the services that are deployed in a cluster in the E-MapReduce (EMR) console.

# Prerequisites

An EMR cluster is created. For more information, see Create a cluster.

## Procedure

- 1. Go to the Cluster Management page.
  - i.
  - ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
  - iii. Click the Cluster Management tab.
- 2. On the Cluster Management page, find your cluster and click the cluster ID.

In the Services section, you can view the status of all services that are deployed in the cluster.

Services 🕝				Add Service
😤 HDFS i Good	 🕼 YARN 🕒 Good		🐔 HIVE 🕒 Good	
55 GANGLIA 🖲 Good	 🞝 SPARK 🕒 Good		🖶 HUE 🔍 Good	
🖳 TEZ 🕒 Good	 SQOOP Good		HUDI 🕒 Good	•••
KNOX 🖲 Good	 S OPENLDAP Good	•••	BIGBOOT Good	•••
DELTALAKE Good	 SMARTDATA Good		ICEBERG • Good	

Only the services that you selected when you created the cluster are displayed in the Services section.

3. In the Services section, click HDFS.

On the page that appears, you can view the service status, deployment topology, service configuration, and configuration history.

Onte The HDFS service is used as an example in this topic.

# 4.2. Add a service

You can add a service to an existing cluster in the E-MapReduce (EMR) console. This topic describes how to add a service to an existing cluster in the EMR console.

# Prerequisites

An EMR cluster is created. For more information, see Create a cluster.

# Procedure

- 1. Go to the Clusters and Services page.
  - i.
  - ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
  - iii. Click the Cluster Management tab.
  - iv. On the **Cluster Management** page, find the cluster to which you want to add a service and click the **cluster ID**.
- 2. In the upper-right corner of the Clusters and Services page, click Add Service.
- 3. In the Add Service dialog box, select the service that you want to add and click OK.

You can view the added service in the **Services** section.

# 4.3. Restart a service

After you modify the configuration items of a service, you must restart the service to make the configurations take effect. This topic describes how to restart a service in the E-MapReduce (EMR) console.

# Prerequisites

A cluster is created. For more information, see Create a cluster.

# Precautions

- To minimize impacts on business operations, we recommend that you restart a service in rolling mode. For instances deployed in primary/secondary mode, restart the service on the secondary instance and then the service on the primary instance.
- If you clear Rolling Execute, the service is restarted on all nodes at the same time. In this case, the service stops running during the restart. Proceed with caution.

# Procedure

1. Go to the Cluster Overview page.

i.

- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the Cluster Management tab.
- iv. On the **Cluster Management** page, find your cluster and click **Details** in the Actions column.
- 2. In the left-side navigation pane of the page that appears, choose **Cluster Service > HDFS**.

The HDFS service is used as an example in this topic.

- 3. In the upper-right corner of the HDFS service page, choose Actions > Restart All Components.
- 4. In the **Cluster Activities** dialog box, specify **Description** and click **OK**.
- 5. In the **Confirm** message, click **OK**.

# 4.4. Manage parameters for services

E-MapReduce (EMR) allows you to modify and add parameters for services such as HDFS, YARN, and Spark. This topic describes how to modify and add the parameters.

## Prerequisites

An EMR cluster is created. For more information, see Create a cluster.

# **Modify parameters**

- 1. Go to the Cluster Overview page of your cluster.
  - i.
  - ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
  - iii. Click the Cluster Management tab.
  - iv. On the Cluster Management page, find your cluster and click Details in the Actions column.
- 2. Go to the **Configure** tab.
  - i. Click the service for which you want to modify parameters.

For example, in the left-side navigation pane, choose **Cluster Service > HDFS**.

ii. On the HDFS service page, click the **Configure** tab.

Home Page > Custer Management > Cluster Anagement > Cluster > Service > HDPS < Back Normal @ HDFS ~ Current Cluster;					
Configuration Filter	Service Configuration	Deploy Client Configuration     Save			
Search: Please Input	ALL hdfs-site kms-site httpfs-site hadoop-env httpfs-env < >				
Scope:	dfs.datanode.cache.revocation.timeout.ms 900000	0			
Default Cluster Configuration 🐣	dfs.ha.automatic-failover.enabled false	0			
Туре	dfs:replication 2	0			
BASIC ADVANCED INFORMATION DATA_PATH	dfs.datanode.available-space-volume-choosing-policy.balanced-space-threshold 10737418240	0			
PERFORMANCE TIME CODEC OSS PORT	dfs.namenode.write-lock-reporting-threshold-ms 5000	0			
MEMORY DISK NETWORK PATH URI	dfs.namenode.avoid.read.stale.datanode false	0			
	nfs.kerberos.principal	0			
	dfs.namenode.lease-recheck-interval-ms 2000	0			
	dfs.namenode.audit.log.debug.cmdlist	0			
	dfs.client.block.write.locateFollowingBlock.initial.delay.ms 400	0			
	dfs.mover.max-no-move-interval 60000	0			

- 3. Modify parameters.
  - i. In the **Configuration Filter** section, enter the keyword of the parameter you want to modify in the search box and click the **Q** icon.
  - ii. Find the parameter you want to modify and change its value.
- 4. Save the configurations.
  - i. In the upper-right corner of the Service Configuration section, click Save.
  - ii. In the **Confirm Changes** dialog box, specify Description and turn on **Auto-update Configuration**.
  - iii. Click OK.
- 5. Make the configurations take effect.

Perform the following operations based on the type of configurations you added to make the configurations take effect:

• Client-side configurations

- a. After you add client-side configurations, click **Deploy Client Configuration** in the upperright corner of the **Configure** tab.
- b. In the Cluster Activities dialog box, specify Target Nodes, Actions on Failures, and Description.
- c. Click OK.

You can click History in the upper-right corner to view the execution status and progress.

- Server-side configurations
  - a. After you add server-side configurations, click **Actions** in the upper-right corner of the **HDFS** page and select an option to restart the related service.
  - b. In the upper-right corner of the Service Configuration section, click **Deploy Client Configuration**.
  - c. In the Cluster Activities dialog box, specify Target Nodes, Actions on Failures, and Description.
  - d. Click OK.

You can click History in the upper-right corner to view the execution status and progress.

# Add parameters

1. Go to the Cluster Overview page of your cluster.

i.

- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the Cluster Management tab.
- iv. On the **Cluster Management** page, find your cluster and click **Details** in the Actions column.
- 2. Go to the **Configure** tab.
  - i. Click the service to which you want to add parameters.

For example, in the left-side navigation pane, choose **Cluster Service > HDFS**.

ii. On the HDFS service page, click the **Configure** tab.

Status Component Deployment Configure	History	
Configuration Filter	Service Configuration	Deploy Client Configuration     Sav
iearch: Please Input	ALL höfs-site kms-site httpfs-site hadoop-env httpfs-env < >	
	dfs.datanode.cache.revocation.timeout.ms 900000	0
cope: Default Cluster Configuration \vee	dfs.ha.automatic-failover.enabled false	6
/pe	dfs.replication 2	0
BASIC ADVANCED INFORMATION DATA	dis.datanode.available-space-volume-choosing-policy.balanced-space-threshold 10737418240	0
LOG_PATH LOG JVM DATA DATABAS	dfs.namenode.write-lock-reporting-threshold-ms 5000	0
MEMORY DISK NETWORK PATH URI	dfs.namenode.avoid.read.stale.datanode false	0
	nfs.kerberos.principal	0
	dfs.namenode.lease-recheck-interval-ms 2000	0
	dfs.namenode.audit.log.debug.cmdlist	0
	dfs.client.block.write.locateFollowingBlock.initial.delay.ms 400	0
	dfs.mover.max-no-move-interval 60000	0

- 3. Add parameters.
  - i. In the Service Configuration section, click the tab where you want to add parameters.

ii. Click **Custom Configuration** in the upper-right corner. In the Add Configuration Item dialog box, configure the parameters you want to add.

Status Component Deployment Configure History Configuration Filter Search Please Input	Service Configuration ALL   hdfs-site   hadoo	p-env ⊨ httpfs-env ⊂core-site	Deploy Ciert Configuration     Sour     Custom Configuration
Add Configuration Item			×
* Key	* Value	Description	Actions
			Delete
Add			
			OK Cancel

Parameter	Description
Кеу	The name of the parameter that you want to add.
Value	The value of the parameter.
Description	The description of the parameter.
Actions	You can click Delete in this column to remove the parameter.

- iii. Click OK.
- 4. Save the configurations.
  - i. In the upper-right corner of the Service Configuration section, click Save.
  - ii. In the **Confirm Changes** dialog box, specify Description and turn on **Auto-update Configuration**.
  - iii. Click OK.
- 5. Make the configurations take effect.

Perform the following operations based on the type of configurations you added to make the configurations take effect:

- Client-side configurations
  - a. After you add client-side configurations, click **Deploy Client Configuration** in the upperright corner of the **Configure** tab.
  - b. In the Cluster Activities dialog box, specify Target Nodes, Actions on Failures, and Description.
  - c. Click OK.

You can click History in the upper-right corner to view the execution status and progress.

- Server-side configurations
  - a. After you add server-side configurations, click **Actions** in the upper-right corner of the **HDFS** page and select an option to restart the related service.
  - b. In the upper-right corner of the Service Configuration section, click **Deploy Client Configuration**.
  - c. In the Cluster Activities dialog box, specify Target Nodes, Actions on Failures, and Description.

d. Click OK.

You can click **History** in the upper-right corner to view the execution status and progress.

# 4.5. Customize software configurations

E-MapReduce (EMR) provides the Custom Software Settings feature for you to customize the configurations of software, such as Hadoop, Hive, and Pig, when you create a cluster. This topic describes how to customize software configurations.

## Limits

You can use the Custom Software Settings feature only when you create a cluster.

## Procedure

1. Go to the Cluster Management page.

i.

- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the Cluster Management tab.
- 2. Click Cluster Wizard in the upper-right corner.
- 3. In the Advanced Settings section of the Software Settings step, turn on Custom Software Settings.

Advanced Settings	
Kerberos Mode: 👩	All components in a cluster that is enabled with the Kerberos mode will use Kerberos authentication to validate requests. For more information, see Kerberos introduction of
Custom Software Settings: 🍘	Before creating a cluster, you can use a JSON file to customize parameters for the cluster components. For more information, see Software Settings 🗗

You can specify a configuration file in the JSON format and overwrite or add default cluster parameters. The following example shows the content of a configuration file in the JSON format :

L		
	{	
		"ServiceName":"YARN",
		"FileName":"yarn-site",
		"ConfigKey":"yarn.nodemanager.resource.cpu-vcores",
		"ConfigValue":"8"
	},	
	{	
		"ServiceName":"YARN",
		"FileName":"yarn-site",
		"ConfigKey":"aaa",
		"ConfigValue":"bbb"
	}	
]		

- ServiceName: the service name. You must specify the service name in all uppercase.
- FileName: the name of the file. The name of the file that is actually passed. You need to remove the suffix.
- ConfigKey : the name of a configuration item.
- ConfigValue : the value of the configuration item.

The following table lists the configuration files of each service.

Service	Configuration file
Hadoop	<ul> <li>core-site.xml</li> <li>log4j.properties</li> <li>hdfs-site.xml</li> <li>mapred-site.xml</li> <li>yarn-site.xml</li> <li>httpsfs-site.xml</li> <li>capacity-scheduler.xml</li> <li>hadoop-env.sh</li> <li>httpfs-env.sh</li> <li>mapred-env.sh</li> <li>yarn-env.sh</li> </ul>
Pig	<ul><li>pig.properties</li><li>log4j.properties</li></ul>
Hive	<ul> <li>hive-env.sh</li> <li>hive-site.xml</li> <li>hive-exec-log4j.properties</li> <li>hive-log4j.properties</li> </ul>

After you customize software configurations, you can continue to create the cluster. For more information, see Create a cluster.

# 4.6. View the deployment information of service components

When you create an E-MapReduce (EMR) cluster, the service components that are deployed on cluster nodes vary based on the cluster type. For example, the NameNode component of the HDFS service is deployed on the master node of a Hadoop cluster. This topic describes how to view the deployment information of service components on each node of an EMR cluster.

# Prerequisites

An EMR cluster is created. For more information, see Create a cluster.

# Procedure

- 1. Go to the Cluster Overview page of your cluster.
  - i.
  - ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
  - iii. Click the **Cluster Management** tab.
  - iv. Find your cluster and click **Det ails** in the **Actions** column.
- 2. In the left-side navigation pane, click **Cluster Service** and select the service whose component deployment information you want to view.
- 3. On the page that appears, click the **Component Deployment** tab.

The deployment information is displayed on the Component Deployment tab. The following sections provide the deployment information of each cluster type:

- Hadoop cluster
- Druid cluster
- Dataflow-Kafka cluster
- Flink cluster
- Data Science cluster
- ClickHouse clust er

# Hadoop cluster

The following tables provide the deployment information of service components on each node of a Hadoop cluster of EMR V3.29.0.

• Required services

Service	Component on the master node	Component on the core nodes
HDFS	<ul> <li>KMS</li> <li>SecondaryNameNode</li> <li>HttpFS</li> <li>HDFS Client</li> <li>NameNode</li> </ul>	<ul><li>DataNode</li><li>HDFS Client</li></ul>
YARN	<ul> <li>ResourceManager</li> <li>App Timeline Server</li> <li>JobHistory</li> <li>WebAppProxyServer</li> <li>Yarn Client</li> </ul>	<ul><li>Yarn Client</li><li>NodeManager</li></ul>
Hive	<ul> <li>Hive MetaStore</li> <li>HiveServer2</li> <li>Hive Client</li> </ul>	Hive Client

Service	Component on the master node	Component on the core nodes
Spark	<ul> <li>Spark Client</li> <li>SparkHistory</li> <li>ThriftServer</li> </ul>	Spark Client
Knox	Knox	N/A
Tez	<ul><li> Tomcat</li><li> Tez Client</li></ul>	Tez Client
Ganglia	<ul> <li>Gmond</li> <li>Httpd</li> <li>Gmetad</li> <li>Ganglia Client</li> </ul>	<ul><li>Gmond</li><li>Ganglia Client</li></ul>
Sqoop	Sqoop Client	Sqoop Client
Bigboot	<ul><li>Bigboot Client</li><li>Bigboot Monitor</li></ul>	<ul><li>Bigboot Client</li><li>Bigboot Monitor</li></ul>
OpenLDAP	OpenLDAP	N/A
Hue	Hue	N/A
SmartData	<ul> <li>Jindo Namespace Service</li> <li>Jindo Storage Service</li> <li>Jindo Client</li> </ul>	<ul><li> Jindo Storage Service</li><li> Jindo Client</li></ul>

# • Optional services

Service	Component on the master node	Component on the core nodes
LIVY	Livy	N/A
Superset	Superset	N/A
Flink	<ul><li> FlinkHistoryServer</li><li> Flink Client</li></ul>	Flink Client
RANGER	<ul> <li>RangerPlugin</li> <li>RangerAdmin</li> <li>RangerUserSync</li> <li>Solr</li> </ul>	RangerPlugin

Service	Component on the master node	Component on the core nodes
Storm	<ul> <li>Storm Client</li> <li>UI</li> <li>Nimbus</li> <li>Logviewer</li> </ul>	<ul><li>Storm Client</li><li>Supervisor</li><li>Logviewer</li></ul>
Phoenix	Phoenix Client	Phoenix Client
Kudu	<ul><li>Kudu Master</li><li>Kudu Client</li></ul>	<ul> <li>Kudu Tserver</li> <li>Kudu Master</li> <li>Kudu Client</li> </ul>
HBase	<ul> <li>HMaster</li> <li>HBase Client</li> <li>ThriftServer</li> </ul>	<ul><li> HBase Client</li><li> HRegionServer</li></ul>
ZooKeeper	<ul><li> ZooKeeper follower</li><li> ZooKeeper Client</li></ul>	<ul> <li>ZooKeeper follower</li> <li>ZooKeeper leader</li> <li>ZooKeeper Client</li> </ul>
Oozie	Oozie	N/A
Presto	<ul><li>Presto Client</li><li>PrestoMaster</li></ul>	<ul><li>Presto Client</li><li>PrestoWorker</li></ul>
Impala	<ul> <li>Impala Runtime and Shell</li> <li>Impala Catalog Server</li> <li>Impala StateStore Server</li> </ul>	<ul><li>Impala Runtime and Shell</li><li>Impala Daemon Server</li></ul>
Pig	Pig Client	Pig Client
Zeppelin	Zeppelin	N/A
FLUME	<ul><li> Flume Agent</li><li> Flume Client</li></ul>	<ul><li> Flume Agent</li><li> Flume Client</li></ul>

# Druid cluster

The following tables provide the deployment information of service components on each node of a Druid cluster of EMR V3.29.0.

• Required services

Service	Component on the master node	Component on the core nodes

Service	Component on the master node	Component on the core nodes
Druid	<ul> <li>Druid Client</li> <li>Coordinator</li> <li>Overlord</li> <li>Broker</li> <li>Router</li> </ul>	<ul> <li>MiddleManager</li> <li>Historical</li> <li>Druid Client</li> </ul>
HDFS	<ul> <li>KMS</li> <li>SecondaryNameNode</li> <li>HttpFS</li> <li>HDFS Client</li> <li>NameNode</li> </ul>	<ul><li>DataNode</li><li>HDFS Client</li></ul>
Ganglia	<ul> <li>Gmond</li> <li>Httpd</li> <li>Gmetad</li> <li>Ganglia Client</li> </ul>	<ul><li>Gmond</li><li>Ganglia Client</li></ul>
ZooKeeper	<ul> <li>ZooKeeper follower</li> <li>ZooKeeper Client</li> </ul>	<ul> <li>ZooKeeper leader</li> <li>ZooKeeper follower</li> <li>ZooKeeper Client</li> </ul>
OpenLDAP	OpenLDAP	N/A
Bigboot	<ul><li>Bigboot Client</li><li>Bigboot Monitor</li></ul>	<ul><li>Bigboot Client</li><li>Bigboot Monitor</li></ul>
SmartData	<ul> <li>Jindo Namespace Service</li> <li>Jindo Storage Service</li> <li>Jindo Client</li> </ul>	<ul><li> Jindo Storage Service</li><li> Jindo Client</li></ul>

#### • Optional services

Service	Component on the master node	Component on the core nodes
YARN	<ul> <li>ResourceManager</li> <li>App Timeline Server</li> <li>JobHistory</li> <li>WebAppProxyServer</li> <li>Yarn Client</li> </ul>	<ul><li>Yarn Client</li><li>NodeManager</li></ul>
Superset	Superset	N/A

# Dataflow-Kafka cluster

The following tables provide the deployment information of service components on each node of a Dataflow-Kafka cluster of EMR V3.29.0.

• Required services

Service	Component on the master node	Component on the core nodes
Kafka-Manager	Kafka Manager	N/A
Kafka	<ul> <li>Kafka Client</li> <li>KafkaMetadataMonitor</li> <li>Kafka Rest Proxy</li> <li>Kafka Broker broker</li> <li>Kafka Schema Registry</li> </ul>	<ul><li> Kafka Broker broker</li><li> Kafka Client</li></ul>
Ganglia	<ul> <li>Gmond</li> <li>Httpd</li> <li>Gmetad</li> <li>Ganglia Client</li> </ul>	<ul><li>Gmond</li><li>Ganglia Client</li></ul>
ZooKeeper	<ul><li> ZooKeeper follower</li><li> ZooKeeper Client</li></ul>	<ul> <li>ZooKeeper leader</li> <li>ZooKeeper follower</li> <li>ZooKeeper Client</li> </ul>
OpenLDAP	OpenLDAP	N/A

#### • Optional services

Service	Component on the master node	Component on the core nodes
RANGER	<ul> <li>RangerPlugin</li> <li>RangerUserSync</li> <li>RangerAdmin</li> <li>Solr</li> </ul>	RangerPlugin
Кпох	Кпох	N/A

# Flink cluster

The following tables provide the deployment information of service components on each node of a Flink cluster of EMR V3.30.0.

• Required services

Service	Component on the master node	Component on the core nodes

Service	Component on the master node	Component on the core nodes
HDFS	<ul> <li>KMS</li> <li>SecondaryNameNode</li> <li>HttpFS</li> <li>HDFS Client</li> <li>NameNode</li> </ul>	<ul><li>DataNode</li><li>HDFS Client</li></ul>
YARN	<ul> <li>ResourceManager</li> <li>App Timeline Server</li> <li>JobHistory</li> <li>WebAppProxyServer</li> <li>Yarn Client</li> </ul>	<ul><li>Yarn Client</li><li>NodeManager</li></ul>
Ganglia	<ul> <li>Gmond</li> <li>Httpd</li> <li>Gmetad</li> <li>Ganglia Client</li> </ul>	<ul><li>Gmond</li><li>Ganglia Client</li></ul>
ZooKeeper	<ul><li> ZooKeeper</li><li> ZooKeeper Client</li></ul>	<ul><li> ZooKeeper</li><li> ZooKeeper Client</li></ul>
Knox	Knox	N/A
Flink-Vvp	Flink-Vvp	N/A
OpenLDAP	OpenLDAP	N/A

#### • Optional services

Service	Component on the master node	Component on the core nodes
PAI-Alink	Alink	N/A

# Data Science cluster

The following tables provide the deployment information of service components on each node of a Data Science cluster of EMR V3.29.1.

• Required services

Service	Component on the master node	Component on the core nodes

Service	Component on the master node	Component on the core nodes
HDFS	<ul> <li>HDFS Client</li> <li>KMS</li> <li>HttpFS</li> <li>NameNode</li> <li>SecondaryNameNode</li> </ul>	<ul><li> HDFS Client</li><li> DataNode</li></ul>
YARN	<ul> <li>WebAppProxyServer</li> <li>JobHistory</li> <li>App Timeline Server</li> <li>ResourceManager</li> <li>Yarn Client</li> </ul>	<ul> <li>NodeManager</li> <li>Yarn Client</li> </ul>
ZooKeeper	<ul><li> ZooKeeper Client</li><li> ZooKeeper follower</li></ul>	<ul> <li>ZooKeeper Client</li> <li>ZooKeeper leader</li> <li>ZooKeeper follower</li> </ul>
Кпох	Кпох	N/A
Tensorflow on YARN	<ul> <li>TensorFlow-On-YARN- Gateway</li> <li>TensorFlow-On-YARN-History- Server</li> <li>TensorFlow-On-YARN</li> </ul>	<ul> <li>TensorFlow-On-YARN-Client</li> <li>TensorFlow-On-YARN- Gateway</li> </ul>
SmartData	<ul> <li>o Jindo Namespace Service</li> <li>o Jindo Storage Service</li> <li>o Jindo Client</li> </ul>	<ul><li> Jindo Storage Service</li><li> Jindo Client</li></ul>
Bigoot	<ul><li>Bigboot Monitor</li><li>Bigboot Client</li></ul>	<ul><li>Bigboot Monitor</li><li>Bigboot Client</li></ul>
PAI-EASYREC	Easyrec	Easyrec
PAI-EAS	PAIEAS	PAIEAS
PAI-Faiss	Faiss	Faiss
PAI-Redis	Redis	Redis
PAI-Alink	Alink	N/A
Flink-Vvp	Flink-Vvp	N/A
OpenLDAP	OpenLDAP	N/A

Service	Component on the master node	Component on the core nodes
Jindo SDK	Jindo SDK	Jindo SDK

#### • Optional services

Service	Component on the master node	Component on the core nodes
Zeppelin	Zeppelin	N/A
PAI-REC	Rec	N/A
AUTOML	AUTOML	AUTOML
TensorFlow	TensorFlow	TensorFlow

# ClickHouse cluster

The following table provides the deployment information of service components on each node of a ClickHouse cluster of EMR V3.35.0.

Service	Component on the master node	Component on the core nodes
Ganglia	<ul><li>Gmond</li><li>Httpd</li><li>Gmetad</li><li>Ganglia Client</li></ul>	<ul><li>Gmond</li><li>Ganglia Client</li></ul>
ZooKeeper	<ul><li>ZooKeeper Client</li><li>ZooKeeper follower</li></ul>	<ul><li>ZooKeeper Client</li><li>ZooKeeper leader</li><li>ZooKeeper follower</li></ul>
ClickHouse	<ul><li>ClickHouse Server</li><li>ClickHouse Client</li></ul>	<ul><li>ClickHouse Server</li><li>ClickHouse Client</li></ul>

# 4.7. Access the Web UI4.7.1. Create an SSH tunnel to access web UIs of open source components

This topic describes how to create an SSH tunnel to access the web UIs of open source components.

# Background information

The ports over which you can access the web UIs of open source components are disabled for security purposes. The open source components include Hadoop, Spark, and Flink in an E-MapReduce (EMR) cluster. You can access these web UIs from the EMR console. You can also create an SSH tunnel on your on-premises server and enable port forwarding to access these web UIs. Dynamic port forwarding and local port forwarding are supported.

For more information about how to access the web UIs of open source components from the EMR console, see Access the web UIs of open source components.

# Prerequisites

- An EMR cluster is created. For more information, see Create a cluster.
- Your on-premises server is connected to the master node of the cluster. You can turn on Assign Public IP Address during cluster creation to associate an elastic IP address (EIP) with your cluster. You can also assign a fixed public IP address or an EIP address to the master node of your cluster in the ECS console after the cluster is created. For more information, see Bind an ENI.

# Obtain the public IP address of the master node

- 1. Go to the Cluster Overview page.
  - i.
  - ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
  - iii. Click the Cluster Management tab.
  - iv. On the **Cluster Management** page, find your cluster and click **Details** in the Actions column.
- 2. In the **Instance Info** section of the **Cluster Overview** page, view the public IP address of the master node.

I Instance Info 🕝					
Master Instance Group (MASTER) 🖉 Pay-As-You-Go	ECS ID	Deployment Status	Public Network IP	Internal Network IP	Created At
<ul> <li>Instances: 1</li> <li>CPU: 4vCores</li> </ul>	i-bp 🗗 🖬 🖬	<ul> <li>Normal</li> </ul>	116.	192.	May 25, 2020, 10:08:42
+ Memory: 16GB + Data Disk Type: 80GB ESSD Disk*1	Show all nodes Ø			Items per	Page::8Items (1) Total Items: 1

# Obtain the name of the master node

1. Go to the Cluster Overview page.

i.

- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the **Cluster Management** tab.
- iv. On the Cluster Management page, find your cluster and click Details in the Actions column.
- 2. In the left-side navigation pane of the Cluster Overview page, click Instances.
- 3. On the **Instances** page, view the **node name** that corresponds to the public IP address of the master node.

For information about how to obtain the public IP address of the master node, see Obtain the public IP address of the master node.

6	Cluster Service 🗸	Instance Name	ECS ID	IP Addresses	Role 🍸
ø	Cluster Resources		i-br	Internal Network ID:102	0005
=	Instances	emr-worker-1	i-pi	Internal Network IP:192	CORE
\$	Cluster Scripts	emr-header-1	i-t	Internal Network IP:192.	MASTER
¢	Connect Strings			Public Network:120.	
×	Auto Scaling	emr-worker-2	i-b	Internal Network IP:192.	CORE
20	Users				

# Enable dynamic port forwarding

Create an SSH tunnel to allow communication between a port of your on-premises server and the master node of an EMR cluster. Run the on-premises SOCKS proxy server that listens on the port. The port data is forwarded to the master node of the EMR cluster by using the SSH tunnel.

- 1. Create an SSH tunnel.
  - Use a key:

```
ssh -i <Storage path of the key file> -N -D 8157 root@<Public IP address of the master node>
```

• Use a username and a password:

ssh -N -D 8157 root@<Public IP address of the master node>

Parameter description:

- 8157 : Port 8157 is used in this example. You can replace this port with an unoccupied port on your on-premises server in actual configuration.
- -D : Dynamic port forwarding is enabled. Start the SOCKS proxy process to listen on the port.
- <Public IP address of the master node> : For more information about how to obtain the public IP address of the master node, see Obtain the public IP address of the master node.
- Storage path of the key file> : the path where the key file is stored.
- 2. Configure the Google Chrome browser.

Notice Keep your on-premises server running after the tunnel is created. No responses are returned.

You can use one of the following methods to configure the Google Chrome browser:

• Use the CLI

a. Open the CLI and go to the local installation directory of the Google Chrome browser client.

The default installation directory of Google Chrome depends on the operating system.
--

Operating system	Google Chrome installation directory
Mac OS X	/Applications/Google Chrome.app/Contents/macOS/Google Chr ome
Linux	/usr/bin/google-chrome
Windows	<i>C</i> :\ <i>Program Files (x86)\Google\Chrome\Application\</i>

#### b. Run the following command in the default installation directory of Google Chrome:

```
chrome --proxy-server="socks5://localhost:8157" --host-resolver-rules="MAP * 0.0.0.
0 , EXCLUDE localhost" --user-data-dir=/tmp/
```

#### Parameter description:

- Image: two provides the format of a file directory is /tmp/.
  the format of a file directory is /tmp/.
- 8157 : Port 8157 is used in this example. You can replace this port with an unoccupied port on your on-premises server in actual configuration.
- c. Enter *http://<Name of the master node>:<Port number>* in the address bar of the browser and press Enter to access a specific web UI.

For more information about the ports of components, see Common ports of services. For more information about how to obtain the name of the master node, see Obtain the name of the master node.

For example, enter *http://emr-header-1:8088* in the address bar of the browser and press Enter to access the web UI of YARN.

• Use a Google Chrome extension

Extensions allow you to easily manage and use proxies in your web browser. You can use an extension to browse web pages and access web UIs at the same time.

- a. Add the Google Chrome extension Proxy SwitchyOmega.
- b. Click this extension and select **Options** from the short cut menu.
- c. On the SwitchyOmega page, click **New profile** in the left-side navigation pane. In the New Profile dialog box, enter a profile name, such as SSH tunnel, in the **Profile name** field, select **PAC Profile**, and then click Create.

d. Enter the following content in the PAC Script editor:

```
function regExpMatch(url, pattern) {
  try { return new RegExp(pattern).test(url); } catch(ex) { return false; }
}
function FindProxyForURL(url, host) {
    // Important: replace 172.31 below with the proper prefix for your VPC subnet
    if (shExpMatch(url, "*localhost*")) return "SOCKS5 localhost:8157";
    if (shExpMatch(url, "*emr-header*")) return "SOCKS5 localhost:8157";
    if (shExpMatch(url, "*emr-worker*")) return "SOCKS5 localhost:8157";
    return 'DIRECT';
}
```

- e. In the left-side navigation pane, click Apply changes to complete the configurations.
- f. Open Google Chrome. Click the SwitchyOmega extension. Then, select the created SSH tunnel.
- g. Enter *http://<Name of the master node>:<Port number>* in the address bar of the browser and press Enter to access a specific web UI.

For more information about the ports of components, see Common ports of services. For more information about how to obtain the name of the master node, see Obtain the name of the master node.

For example, enter *http://emr-header-1:8088* in the address bar of the browser and press Enter to access the web UI of YARN.

# Enable local port forwarding

Notice If you use this method to access a web UI, you cannot go to the job details page.

You can use the local port forwarding method to forward data on a port of the master node to the local port and access the web application interface running on the master node. The SOCKS proxy is not required.

- 1. Run the following command on your on-premises server to create an SSH tunnel:
  - Use a key:

ssh -i <Storage path of the key file> -N -L 8157:<Name of the master node>:8088 root@<P ublic IP address of the master node>

Use a username and a password:

```
ssh -N -L 8157:<Name of the master node>:8088 root@<Public IP address of the master nod e> \ensuremath{\mathsf{e}}\xspace
```

Parameter description:

- Local port forwarding is enabled. You can specify a local port to forward data to the remote port that is hosted on the on-premises web server of the master node.
- 8088 : the port that is used to access ResourceManager on the master node. You can replace this port as required.

For more information about the ports of components, see Common ports of services. For more information about how to obtain the name of the master node, see Obtain the name of the master node.

- 8157 : Port 8157 is used in this example. You can replace this port with an unoccupied port on your on-premises server in actual configuration.
- <Public IP address of the master node> : For more information about how to obtain the public IP address of the master node, see Obtain the public IP address of the master node.
- Storage path of the key file> : the path where the key file is stored.
- 2. Keep your on-premises server running. Open a browser, enter *http://localhost:8157/* in the address bar of the browser, and then press Enter.

# Common ports of services

Service	Port	Description
	50070	The web UI port of HDFS. Parameter: dfs.namenode.http-address or dfs.http.address. <b>Note</b> The dfs.http.address parameter has expired but can still be used.
	50075	The web UI port of DataNode.
	50010	The service port of DataNode. This port is used to transfer data.
	50020	The port of the inter-process communication (IPC) service.
	8020	The remote procedure call (RPC) port of HDFS in a high- availability (HA) cluster.
	8025	The port of ResourceManager. Parameter: yarn.resourcemanager.resource-tracker.address.
Hadoop 2.X	9000	The RPC port of HDFS in a non-HA cluster. Parameter: fs.defaultFS or fs.default.name. <b>Note</b> The fs.default.name parameter has expired but can still be used.
	8088	The web UI port of YARN.
	8485	The RPC port of JournalNode.
	8019	The port of ZKFailoverController (ZKFC).
	19888	The web UI port of JobHistory Server. Parameter: mapreduce.jobhistory.webapp.address.

		Description
	10020	The web UI port of JobHistory Server. Parameter: mapreduce.jobhistory.address.
	8020	The port of NameNode.
	9870	Parameter: dfs.namenode.http-address or dfs.http.address.           Once         The dfs.http.address parameter has expired but can still be used.
	9871	The port of NameNode.
Hadoop 3.X	9866	The port of DataNode.
	9864	The port of DataNode.
	9865	The port of DataNode.
	8088	The port of ResourceManager. Parameter: yarn.resourcemanager.webapp.address.
MapReduce	8021	The port of JobTracker. Parameter: mapreduce.jobtracker.address.
	2181	The port that is used to connect a client to ZooKeeper.
ZooKeeper	2888	The internal communication port of a ZooKeeper cluster. The leader listens on this port.
	3888	The ZooKeeper port that is used to elect a leader.
	16010	The web UI port of the master node of HBase. Parameter: hbase.master.info.port.
	16000	The port of HMaster. Parameter: hbase.master.port.
HBase	16030	The web UI management port of RegionServer of HBase. Parameter: hbase.regionserver.info.port.
	16020	The port of HRegionServer. Parameter: hbase.regionserver.port.
	9099	The port of Thrift Server.

The Java The JDB The JDB The p with The p with The p Schedul The wel resource The wel The wel The wel The wel	b UI port of Driver. This port is used to schedule tasks. b UI port of Spark History Server. C port that is used for communication among the nodes ka cluster.
The JDB The JDB The p with t The p The well schedul The well resource The well The well resource The well The Well	C port of Spark Thrift Sever. Dort on which the master node of Spark communicates the worker nodes. Doort on which a standalone cluster submits applications. Doort on which a standalone cluster submits applications. D UI port of the master node. This port is used to e resources. D UI port of a worker node. This port is used to schedule es. D UI port of Driver. This port is used to schedule tasks. D UI port of Spark History Server. C port that is used for communication among the nodes Ka cluster.
<ul> <li>The pwith f</li> <li>The p</li> <li>The p</li> <li>The p</li> <li>The p</li> <li>The well</li> </ul>	bort on which the master node of Spark communicates the worker nodes. bort on which a standalone cluster submits applications. b UI port of the master node. This port is used to e resources. b UI port of a worker node. This port is used to schedule es. b UI port of Driver. This port is used to schedule tasks. b UI port of Spark History Server. c port that is used for communication among the nodes is a cluster.
with the with the The period of a Kafe	the worker nodes. bort on which a standalone cluster submits applications. b UI port of the master node. This port is used to e resources. b UI port of a worker node. This port is used to schedule es. b UI port of Driver. This port is used to schedule tasks. b UI port of Driver. This port is used to schedule tasks. b UI port of Spark History Server. C port that is used for communication among the nodes ka cluster.
schedul The wel resource The wel The wel The wel The RPC of a Kaf	e resources. b UI port of a worker node. This port is used to schedule es. b UI port of Driver. This port is used to schedule tasks. b UI port of Spark History Server. C port that is used for communication among the nodes ka cluster.
The well The well The RPC of a Kaf	es. b UI port of Driver. This port is used to schedule tasks. b UI port of Spark History Server. C port that is used for communication among the nodes ka cluster.
The wel The RPC of a Kaf	b UI port of Spark History Server. C port that is used for communication among the nodes Ka cluster.
T he RPC of a Kaf	port that is used for communication among the nodes ka cluster.
of a Kaf	ka cluster.
The por	t of the Redic convice
	t of the Redis service.
The we	b UI port of Hue.
The we	b UI port of Oozie.
The we	b UI port of Druid.
	t of Overlord. ter: <b>druid.plaintextPort</b> on the <b>overlord.runtime</b> tab.
Parame	t of MiddleManager. ter: <b>druid.plaintextPort</b> on the <b>Manager.runtime</b> tab.
	t of Coordinator. ter: <b>druid.plaintextPort</b> on the <b>coordinator.runtime</b>
	t of Historical. ter: <b>druid.plaintextPort</b> on the <b>historical.runtime</b>
	Parame middle The por Parame tab. The por Parame

Service	Port	Description
	18082	The port of Broker. Parameter: <b>druid.plaintextPort</b> on the <b>broker.runtime</b> tab.
Ganglia	9292	The web UI port of Ganglia.
Ranger	6080	The web UI port of Ranger.
Kafka Manager	8085	The port of Kafka Manager.
Superset	18088	The web UI port of Superset.
Impala	21050	The JDBC port that is used to connect to Impala.
Presto	9090	The web UI port of Presto.

# 4.7.2. Access the web UIs of open source

# components

This topic describes how to configure security group rules and access the web UIs of open source components in an E-MapReduce (EMR) cluster. After you create a cluster, EMR binds several domain names to the cluster for you to access the web UIs of open source components.

# Prerequisites

An EMR cluster is created. For more information, see Create a cluster.

Onte An elastic IP address (EIP) is associated with the EMR cluster.

# Configure security group rules

If you use a component for the first time, you must perform the following steps to configure security group rules:

1. Obtain the public IP address of your on-premises machine.

For security purposes, we recommend that you allow only access from the current public IP address when you configure a security group rule. To obtain your current public IP address, visit <a href="http://myip.ipip.net/">http://myip.ipip.net/</a>. You can view your public IP address.

2. Add security group rules.

i.

- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the Cluster Management tab.
- iv. On the **Cluster Management** page, find your cluster and click **Details** in the Actions column.
- v. In the **Network Info** section of the **Cluster Overview** page, record the value of **Network Type** and click the link of **Security Group ID**.

#### vi. Enable the required ports.

Notice To prevent attacks from external users, you are not allowed to set Authorization Object to 0.0.0/0.

The following table lists the ports you need to enable to access the web UIs of different components.

Component	Port
YARN UI	
HDFS UI	
Spark History Server UI	
Ganglia UI	
Oozie	
Tez	8443
ImpalaCatalogd	Note After Ranger is deployed in your cluster, you can access the web UI of Ranger.
ImpalaStatestored	
Storm	
Ranger UI	
Zeppelin	
Hue	8888

For example, you can perform the following operations to enable port 8443:

- a. On the **Security Group Rules** page, click **Add Security Group Rule** in the upper-right corner.
- b. In the Add Security Group Rule dialog box, set Port Range to 8443/8443.
- c. Set Authorization Object to the public IP address obtained in Step 1.
- d. Click OK.

#### ? Note

- If the network type of the cluster is VPC, set NIC Type to Internal Network and Rule Direction to Inbound. If the network type of the cluster is classic network, set NIC Type to Internet and Rule Direction to Inbound. In this topic, the VPC network type is used.
- When you configure inbound and outbound rules for applications, follow the principle of least privilege. Enable only the ports required by your applications.

#### vii. View the added rule on the **Inbound** tab.

Inbound	Outbound							▲ Import Rules	Ł Export Rules →
Action	Protocol Type	Port Range	Authorization Type(All) 👻	Authorization Object	Description	Priority	Created At		Actions
Allow	Customized TCP	8443/8443	IPv4 CIDR Block	11000	-	1	October 30, 2019, 16:40	Modify	Clone Delete
Allow	Customized TCP	8080/8080	IPv4 CIDR Block		-	1	October 30, 2019, 16:39	Modify	Clone   Delete
Allow	Customized TCP	8888/8888	IPv4 CIDR Block	10004-0		1	October 30, 2019, 16:38	Modify	Clone   Delete

Network access is securely enabled and network configuration is complete.

## Access the web UIs of open source components

#### 1.

- 2. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- 3. Click the Cluster Management tab.
- 4. On the **Cluster Management** page, find your cluster and click **Details** in the Actions column.
- 5. In the left-side navigation pane of the Cluster Overview page, click **Connect Strings**.
- 6. On the **Public Connect Strings** page, find the component whose web UI you want to access and click its link.
  - In V2.X.X versions later than V2.7.X or V3.X.X versions later than V3.5.X, you can use a Knox account to access the web UIs of open source components. For more information about how to create a Knox account, see Manage user accounts. For more information about how to use Knox, see Knox. To access the web UI of Hue, you must use the Hue username and password. For more information about how to use Hue, see Use Hue. You can directly access the web UI of Zeppelin without a username and password.
  - After Ranger is deployed in your cluster, you can use the default username and password to access the web UI of Ranger. For more information, see Overview.
  - You can access the web UI of Flink based on the version of your cluster:
    - Clusters of an EMR version earlier than V3.29.0:

Use an SSH tunnel. For more information, see Create an SSH tunnel to access web UIs of open source components.

(?) Note To access a Flink job on the web UI of YARN, go to the Connect Strings page in the EMR console, and click the link for the YARN UI in the Connect String column. In the Hadoop console, click the ID of the Flink job to view the details of the Flink job. For more information about how to view the details of a job on the web UI of YARN, see Quick start.

- Clusters of EMR V3.29.0 and later:
  - Flink-VVP: You can access the web UI of Flink-VVP from the EMR console. For more information, see Basic usage.
  - Flink (VVR): You can access the web UI of Flink (VVR) by using an SSH tunnel. For more information, see Create an SSH tunnel to access web UIs of open source components.

(?) Note To access a Flink job on the web UI of YARN, go to the Connect Strings page in the EMR console, and click the link for the YARN UI in the Connect String column. In the Hadoop console, click the ID of the Flink job to view the details of the Flink job. For more information about how to view the details of a job on the web UI of YARN, see Quick start.

# 5.Cluster Operations 5.1. Common file paths

This topic describes the paths of files that are frequently used in E-MapReduce (EMR). You can log on to the master node of your cluster to view the file paths.

# Big data software

Big data software is installed in the */usr/lib/xxx* directory. Examples:

- Hadoop: /usr/lib/hadoop-current
- Spark: /usr/lib/spark-current
- Hive: /usr/lib/hive-current
- Flink: /usr/lib/flink-current
- Flume: /usr/lib/flume-current

You can also log on to the master node of your cluster and run the **env |grep xxx** command to view a software installation directory.

For example, run the following command to view the installation directory of Hadoop:

env |grep hadoop

The following information is returned. /usr/lib/hadoop-current is the installation directory of Hadoop.

HADOOP LOG DIR=/var/log/hadoop-hdfs HADOOP HOME=/usr/lib/hadoop-current YARN PID DIR=/usr/lib/hadoop-current/pids HADOOP\_PID\_DIR=/usr/lib/hadoop-current/pids HADOOP MAPRED PID DIR=/usr/lib/hadoop-current/pids JAVA LIBRARY PATH=/usr/lib/hadoop-current/lib/native: PATH=/usr/lib/sqoop-current/bin:/usr/lib/spark-current/bin:/usr/lib/hive-current/hcatalog/bin :/usr/lib/hive-current/bin:/usr/lib/datafactory-current/bin:/usr/local/sbin:/usr/local/bin:/u sr/sbin:/usr/lib/b2monitor-current//bin:/usr/lib/b2monitor-current//bin:/usr/lib/b 2jindosdk-current//bin:/usr/lib/flow-agent-current/bin:/usr/lib/hadoop-current/bin:/usr/lib/h adoop-current/sbin:/usr/lib/hadoop-current/bin:/usr/lib/hadoop-current/sbin:/root/bin HADOOP CLASSPATH=/usr/lib/hadoop-current/lib/\*:/usr/lib/tez-current/\*:/usr/lib/tez-current/li b/\*:/etc/ecm/tez-conf:/opt/apps/extra-jars/\*:/usr/lib/spark-current/yarn/spark-2.4.5-yarn-shu ffle.jar HADOOP CONF DIR=/etc/ecm/hadoop-conf YARN LOG DIR=/var/log/hadoop-yarn HADOOP MAPRED LOG DIR=/var/log/hadoop-mapred

# Logs

Component logs are stored in the */mnt/disk1/log/xxx* directory. Examples:

- YARN ResourceManager logs: /mnt/disk1/log/hadoop-yarn in the master node
- YARN NodeManager logs: /mnt/disk1/log/hadoop-yarn in a core node or a task node
- HDFS NameNode logs: /mnt/disk1/log/hadoop-hdfs in the master node
- HDFS DataNode logs: /mnt/disk1/log/hadoop-hdfs in a core node or a task node

- Hive logs: /mnt/disk1/log/hive in the master node
- ESS logs: /mnt/disk1/log/ess/in the master node and core or task nodes.

# **Configuration files**

Configuration files are stored in the /etc/ecm/xxx directory. Examples:

- Hadoop: /etc/ecm/hadoop-conf/
- Spark: /etc/ecm/spark-conf/
- Hive: /etc/ecm/hive-conf/
- Flink: /etc/ecm/flink-conf/
- Flume: /etc/ecm/flume-conf/

If you log on to your cluster in SSH mode, you can only view the parameter settings in configuration files. To modify the parameters in configuration files, you must log on to the EMR console.

# Data directory

Cached data in JindoFS: /mnt/disk\*/bigboot/

# 5.2. Log on to a cluster

This topic describes how to log on to an E-MapReduce (EMR) cluster by using an SSH key pair or password. You can log on to an EMR cluster from an on-premises machine that runs a Linux or Windows operating system.

# Prerequisites

- An EMR cluster is created. For more information, see Create a cluster.
- Your on-premises machine is connected to the master node of the cluster. You can turn on **Assign Public IP Address** during cluster creation to associate an elastic IP address (EIP) with your cluster. You can also assign a fixed public IP address or an EIP address to the master node of your cluster in the ECS console after the cluster is created. For more information, see Bind an ENI.

# Context

After your on-premises machine is connected to the master node of your cluster in SSH mode, you can run Linux commands to monitor the cluster and interact with the cluster. You can also create an SSH tunnel to view the web UIs of open source components. For more information, see Create an SSH tunnel to access web UIs of open source components.

# Obtain the public IP address of the master node

1. Go to the Cluster Overview page of your cluster.

- ii. In the top navigation bar, select the region where you want to create a cluster. The region of a cluster cannot be changed after the cluster is created.
- iii. Click the Cluster Management tab.
- iv. On the Cluster Management page, find your cluster and click Details in the Actions column.
- 2. In the **Instance Info** section of the **Cluster Overview** page, view the public IP address of the master node.

i.

I Instance Info 🕝						
Master Instance Group (MASTER) 🖉	Pay-As-You-Go	ECS ID	Deployment Status	Public Network IP	Internal Network IP	Created At
<ul> <li>Instances: 1</li> <li>CPU: 4vCores</li> </ul>		i-bp 🗗 🖬	Normal	116.	192.	May 25, 2020, 10:08:42
Memory: 16GB     Data Disk Type: 80GB ESSD Disk*1		Show all nodes 📀			Items per	Page::8Items < 1 > Total Items: 1

# Log on to the cluster by using an SSH key pair

**Note** For more information about how to obtain the public IP address of the master node, see **Obtain the public IP address of the master node.** 

Use one of the following methods to log on to the cluster:

• Log on from your on-premises machine that runs a Linux operating system

In this example, the private key file *ecs.pem* is used.

i. Run the following command to modify the attribute of the private key file:

chmod 400 ~/.ssh/ecs.pem

~/.ssh/ecs.pem is the path where the *ecs.pem* file is stored on your on-premises machine.

ii. Run the following command to connect to the master node:

ssh -i ~/.ssh/ecs.pem root@<Public IP address of the master node>

• Use PuTTY to log on from your on-premises machine that runs a Windows operating system

Perform the following steps to log on to the cluster:

- i. Download PuTTY and PuTTYgen.
- ii. Convert the format of the private key file from .pemto .ppk.
  - a. Run PuTTYgen. In this example, PuTTYgen 0.73 is used.
  - b. In the Actions section, click Load to import the private key file that is saved when you create a cluster.

Make sure that the format of the file that you want to import is All files (\*.\*).

- c. Select the specific *.pem* file and click **Open**.
- d. Click Save private key.
- e. In the dialog box that appears, click **Yes**. Specify a name for the *.ppk* file and click **Save**.

Save the .ppk file to your on-premises machine. In this example, *kp-123.ppk* is used.

- iii. Run PuTTY.
- iv. In the left-side navigation pane, choose Connection > SSH > Auth. Click Browse below Private key file for authentication and select the .ppk file.
- v. Click Session. Enter the logon account and the public IP address of the master node in the Host Name (or IP address) field.

Session	Basic options for your Pu	TTY session		
Logging	Specify the destination you want to	connect to		
Terminal Keyboard	Host Name (or IP address)	Port		
Bell		22		
Features	Connection type:	SSH O Seria		
Appearance Behaviour Translation Selection	Load, save or delete a stored session Saved Sessions			
Colours	Default Settings	Load		
Data		Save		
Proxy Telnet Rlogin		Delete		
Serial	Close window on exit: Always      Never      O	nly on clean exit		

The format is *root@[Public IP address of the master node]*, such as *root@10.10.xx.xx*.

vi. Click Open.

If the following information appears, the logon is successful.



• Use Command Prompt to log on from your on-premises machine that runs a Windows operating system

Open Command Prompt and run the following command to log on to the cluster:

ssh -i <Storage path of the .pem file on your on-premises machine> root@<Public IP address of the master node>



# Log on to the cluster by using an SSH password

(?) Note The username and password used in the following operations are the root user and password you specified when you created a cluster. For more information about how to obtain the public IP address of the master node, see Obtain the public IP address of the master node.

Procedure:

• Log on from your on-premises machine that runs a Linux operating system

Run the following command in the command-line interface (CLI) of your on-premises machine to connect to the master node:

ssh root@[Public IP address of the master node]

- Log on from your on-premises machine that runs a Windows operating system
  - i. Download and install PuTTY.

Download link: PuTTY.

- ii. Start PuTTY.
- iii. Configure the parameters required to connect to a Linux instance.
  - Host Name (or IP address): Specify the fixed public IP address of the instance or the EIP address associated with the instance.
  - Port : Enter port number 22.
  - Connection type: Select SSH.
  - Saved Sessions: optional. Enter a name that helps you identify the session and click Save to save the session. This way, you do not need to enter session information such as the public IP address when you connect to the instance again.
- iv. Click Open.
- v. Specify the username and press Enter. The default username is root.

The characters of the password are hidden when you enter the password. After you enter the password, press Enter.

# FAQ

- Q: How do I log on to a cluster from my on-premises machine in password-free mode?
  - A: Perform the following steps:
    - i. Open Command Prompt and run the following command to generate a public key:

ssh-keygen

A public key file is generated on your on-premises machine.

id\_rsa 💼 id\_rsa.pub

- ii. Add the generated public key to the master node of the cluster to which you want to log on.
  - a. Go to the */.ssh* directory of the cluster.

cd ~/.ssh

b. Edit the key information of the master node of the cluster.

vim authorized\_keys

- c. Add the content of the generated public key file *id\_rsa.pub* to the *authorized\_keys* file.
- iii. Add the IP address of your on-premises machine to the security group to which the cluster

belongs.

a. Obtain the public IP address of your on-premises machine.

For security reasons, we recommend that you allow access only from the current public IP address when you configure a security group rule. To obtain your current public IP address, visit <a href="http://myip.ipip.net/">http://myip.ipip.net/</a>.

b. Add a security group rule in which port 22 is enabled.

For more information about how to add a security group rule, see Add a security group rule.

8 a 1 -	est-n / vpc-bp1h9b				Nev	v Version Lea	arn more. 🗯 Back	Add Security Group Rule Quick Rule Creation
Inbound	Outbound							± Import ± Export →
Action	Protocol Type	Port Range	Authorization Type(All) 👻	Authorization Object	Description	Priority	Creation Time	Action
Allow	Custom TCP	22/22	IPv4 CIDR Block	1.80.1		1	June 6, 2021, 10:10	Modify   Clone   Delete

iv. Run the following command in Command Prompt to log on to the cluster in password-free mode:

ssh root@<Public IP address of the master node>

- Q: How do I log on to a core node?
  - A: Perform the following steps:
    - i. On the master node, run the following command to switch to the hadoop user:

su hadoop

ii. Log on to the core node in password-free mode.

ssh emr-worker-1

iii. Run the following sudo command to obtain the root permissions:

sudo su - root

# 5.3. Scale out a cluster

You can scale out an E-MapReduce (EMR) cluster that has insufficient computing or storage resources. You can add core nodes and task nodes. By default, the configurations of an added node are the same as those of an existing node. This topic describes how to scale out a cluster.

# Prerequisites

An EMR cluster is created. For more information, see Create a cluster.

# Scale out a cluster

1. Go to the Cluster Management page.

i.

- ii. In the top navigation bar, select the region where you want to create a cluster. The region of a cluster cannot be changed after the cluster is created.
- iii. Click the Cluster Management tab.
- On the Cluster Management page, find the cluster that you want to scale out and choose More > Scale Out in the Actions column.
- 3. In the Scale Out dialog box, click the CORE (Core Instance Group) or TASK (Task Instance

Group) tab and configure related parameters.

Scale Out	x	
CORE (Core Instance Group)	TASK (Task Instance Group)	
Core Instance Group	Node Group Name: Core Instance Group	
+ Create Node Group	VSwitch: (vsw-bp1)	
	Configure: ecs.g6.xlarge 4vCores 16G ESSD Disk 80GB*4Disks	
	Billing Method: Pay-As-You-Go	
	Core Nodes: 2Nodes	
	Add: 5 Nodes	
	CE-MapReduce Service Terms)	
	Total	
	OK	
	Close	

The following table describes the parameters on the CORE (Core Instance Group) tab.

Parameter	Description			
Machine Group Name	The name of the machine group.			
VSwitch	The vSwitch of the cluster.			
Configure	The configurations of the existing node group.			
Billing Method	The billing method of the cluster. The billing method of a new node is the same as that of the cluster and cannot be changed. If the billing method is subscription, you can specify the <b>subscription period</b> for the new node.			
Core Nodes	By default, the total number of your core nodes is displayed.			
	The number of core nodes that you want to add. Click the upward or downward arrow or enter a number in the Add field.			
Add	<b>Note</b> When you adjust the number, the total scale-out fee of a cluster is displayed in the lower-right corner of the Scale Out dialog box.			
E-MapReduce Service Terms	Read and select the terms of service.			

#### 4. Click OK.

5. View the scale-out status.

In the **Instance Info** section of the **Cluster Overview** page, click the node group on which you performed the scale-out operation to view the status of each new node on the right side.

ECS ID	Deployment Status	Public IP Address	Internal IP Address	Created At
i-bp1dw	Scaling out	-	192.1	Sep 28, 2020, 15:35:44
i-bp11g	Normal	-	192.1	Sep 27, 2020, 17:47:21
i-bp19q	Normal	-	192.1	Sep 27, 2020, 17:34:42
i-bp10b	Normal	-	192.1	Sep 27, 2020, 17:34:41
Show all nodes 🕖			ltems per	Page::8Items 1 Total Items:

**Scaling out** indicates that the node is being added. **Normal** indicates that the node is added to the cluster and can provide services.

# Change the password

After you scale out the cluster, you can log on to the new node in SSH mode to change the password of the root user.

- 1. Log on to the master node of the cluster. For more information, see Log on to a cluster.
- 2. Run the following command to switch to the hadoop user:

```
su hadoop
```

3. Run the following command to log on to the new node:

ssh <ip.of.worker>

Obtain the private IP address of the new node. For more information, see View the cluster list and cluster details.

4. Run the following command to change the password of the root user:

sudo passwd root

# 5.4. Scale in a cluster

This topic describes how to reduce the number of task nodes to scale in an E-MapReduce (EMR) cluster.

#### Prerequisites

An EMR cluster is created. For more information, see Create a cluster.

#### Limits

You can reduce only the number of task nodes to scale in an EMR cluster. The cluster you want to scale in must meet the following requirements:

- The EMR version is V2.X later than V2.5.0 or V3.X later than V3.2.0.
- The cluster is in the Idle or Running state.
- The billing method of the cluster is pay-as-you-go.

## Procedure

1. Go to the Cluster Overview page.

i.
- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the Cluster Management tab.
- iv. On the Cluster Management page, find your cluster and click Details in the Actions column.
- 2. On the **Cluster Overview** page, choose **Change Configuration** > **Scale In** in the upper-right corner.
- 3. In the **Scale In** dialog box, click the downward arrow for Reserved Nodes or enter a number in this field to reduce the number of task nodes.
- 4. Click Scale In.
- 5. In the Confirm Scale-in Activity message, click OK.

# 5.5. Auto Scaling

# 5.5.1. Overview

This topic describes the auto scaling feature of E-MapReduce (EMR). You can configure auto scaling rules based on your business requirements and policies. Auto scaling applies only to Hadoop clusters of EMR. After this feature is enabled and configured, EMR automatically adds task nodes to ensure sufficient computing capabilities when your business loads increase and removes task nodes to reduce costs when your business loads decrease.

## Scenarios

You can enable auto scaling in the following scenarios to reduce costs and improve task execution efficiency:

- Add task nodes at the scheduled time to temporarily supplement the computing capabilities.
- Add task nodes based on the settings of certain cluster metrics to ensure that important jobs are completed on time.

### Features

Auto scaling allows you to purchase preemptible or pay-as-you-go instances as required.

# 5.5.2. Create an auto scaling machine group

If your business requirements fluctuate, we recommend that you enable auto scaling for your E-MapReduce (EMR) cluster and configure auto scaling rules to increase or decrease task nodes based on your business requirements.

### Prerequisites

An EMR Hadoop cluster is created. For more information, see Create a cluster.

#### Procedure

1. Go to the Cluster Overview page.

i.

- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the Cluster Management tab.

- iv. On the Cluster Management page, find your cluster and click Details in the Actions column.
- 2. In the left-side navigation pane, choose Auto Scaling > Auto Scaling Settings.
- 3. On the Auto Scaling Settings page, click Create Auto Scaling Machine Group.
- 4. In the Create Task Machine Group dialog box, specify Node Group Name.

Configure the other parameters based on the specifications of the auto scaling nodes that you want to use. You can modify the configurations after you create the auto scaling machine group.

- 5. Read the terms of service and select E-MapReduce Service Terms.
- 6. Click Create Node Group.

After auto scaling is configured and triggered due to business fluctuations, you can view the scaling history and the results of each auto scaling activity on the History page. For more information, see View auto scaling history.

# 5.5.3. Manage auto scaling

If your business requirements fluctuate, we recommend that you enable auto scaling for your E-MapReduce (EMR) cluster and configure auto scaling rules to increase or decrease task nodes based on your business requirements. Auto scaling not only ensures sufficient computing resources for your jobs but also reduces costs.

## Prerequisites

An auto scaling machine group is created. For more information, see Create an auto scaling machine group.

# Configure auto scaling

1. Go to the Cluster Overview page.

i.

- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the Cluster Management tab.
- iv. On the Cluster Management page, find your cluster and click Details in the Actions column.
- 2. In the left-side navigation pane, choose Auto Scaling > Auto Scaling Settings.
- 3. On the Auto Scaling Settings page, find your auto scaling machine group and click Configure Rules in the Actions column.
- 4. In the Auto Scaling Settings pane, configure the parameters.

Parameter	Description
Max Instances	The maximum number of task nodes in the auto scaling machine group. If an auto scaling rule is met but this upper limit has been reached, the system still does not trigger auto scaling. Maximum value: 500.
Minimum Instances	The minimum number of task nodes in the auto scaling machine group. If the number of task nodes set in a scale-out or scale-in rule is less than the value of this parameter, the cluster is scaled based on the value of this parameter when the rule is triggered for the first time.
	For example, if this parameter is set to 3 and a scale-out rule is that one node is added at 00:00 every day, the system adds three nodes at 00:00 on the first day. This way, the requirement for the minimum number of nodes is met.
	You can enable graceful deprecation and set a timeout period to deprecate the task node on which a job on YARN runs. If the period of time for which a job on YARN has run exceeds the timeout period or no job on YARN has run, the system deprecates the task node. The maximum value of Timeout is 3600, in seconds.
Graceful Deprecation	Note To enable graceful deprecation, you must first change the value of the yarn.resourcemanager.nodes.exclude-path parameter on the YARN service page to /etc/ecm/hadoop-conf/yarn- exclude.xml.

#### i. In the **Basic Information** section, configure the parameters described in the following table.

ii. In the **Cost Optimization Policy** section, select Single Billing Method or Cost Optimization Mode.

#### Single Billing Method

The system automatically searches for the instance types that match the vCPU and memory specifications you specified and displays the instance types in the Instance Type section. You must select one or more instance types in the Instance Type section so that the cluster can be scaled based on the selected instance types. Single Billing Method supports the following billing methods:

#### Pay-As-You-Go

The order in which you select instance types determines the priorities of the instances that are used. The hourly price of each instance is displayed below the disk specifications in the Instance Type section. The price is the sum of the EMR service price and ECS instance price.

Instance Selection M	Node Single Billing Method Cost Optimization Mode		
Billing Method 💿	Pay-As-You-Go O Spot Instance		
Instance Type			
Insta	nce Types: A maximum of three instance types can be selected. (The selection order represents priority.)		Filter 4C × 8GiB ×
	Instance Type	vCore	vMem
	ecs.c5.xlarge	4	8
	ecs.hfc6.xlarge	4	8
	ecs.c6.xlarge	4	8
	ecs.c6e.xlarge	4	8
	tem Disk. ESSD Disk SSD Ultra Disk. Details 6* Type: Jick Size: 40 GB * 1 Disks (Capacity Range: 60 ~ 500 GB) IOPS 3800		
D	Data Disk ESSD Disk SSD Ultra Disk Details G Type:		
Selec	Jisks Size:         40         GB *         4         Disks (Capacity Range: 40 ~ 32768 GB) IOPS 3800           cted Instance1         (Priority1) : ecs.c5.xlarge 4vCPU 8GB         Total ¥ 0.02         Save ¥ 1.826 @           cted Instance2         (Priority2) : ecs.ht/6.xlarge 4vCPU 8GB         Total ¥ 0.02         Save ¥ 1.43 @		

#### Spot Instance

**Notice** If you have high service level agreement (SLA) requirements for your jobs, do not select this option to use preemptible instances because the instances may be released due to a failed bid or other reasons.

The order in which you select instance types determines the priorities of the instances that are used. The hourly price of each instance based on the pay-as-you-go billing method is displayed below the disk specifications in the Instance Type section. You can also set an upper limit for the hourly price of each instance. The instance is displayed if its price does not exceed the upper limit. For more information about preemptible instances, see Overview.

Instance Selection Mode	Single Billing Method	Cost Optimization Mode	Ø			
			-			
Billing Method 🛛 🔾 Pay-As	-You-Go 🧿 Spot Instanc	e				
Instance Type						
Instance Tv	nes: A maximum of three i	instance types can be selecte	ed. (The selection order represents p	iority.)		× RGiR
					Filter 4C	✓ 8GiB
	Instance Type			vCore	vMem	
	ecs.c5.xlarge			4	8	
	ecs.hfc6.xlarge			4	8	
	ecs.c6.xlarge			4	8	
	ecs.c6e.xlarge			4	8	
System Dis	sk 🛛 ESSD Disk 본 S	SD Ultra Disk De	tails 🗗			
Тур						
Disk Size	e: 40 GB * 1 I	Disks (Capacity Range: 60 ~	500 GB) IOPS 3800			
Data Dis	sk 🛛 ESSD Disk 🍳 S	SD Ultra Disk De	tails 🗗			
Тур	e:					
Disk Siz	e: 40 GB * 4	1 Disks (Capacit	ry Range: 40 ~ 32768 GB) IOPS 3800			
Selected Ins	stance1 (Priority1) : ecs.o	5.xlarge 4vCPU 8GiB	Total ¥ 0.02 Save ¥ 1.826 Ø	Maximum Cost Per Instance ¥ 0.125~1.24/Hours	Type:: ¥ 1.24 Market Price Range for Current Node	е Туре:
Selected Ins	stance2 (Priority2) : ecs.ł	nfc6.xlarge 4vCPU 8GiB	Total ¥0.02 Save ¥1.43 @	¥ 0.125~1.24/Hours Maximum Cost Per Instance ¥ 0.117~0.896/Hours	Type:: ¥ 0.896 Market Price Range for Current Node	е Туре:

#### Cost Optimization Mode

In this mode, you can develop a detailed cost optimization policy to achieve a balance between cost and stability.

Cost Optimization Policy			
Instance Selection Mode Single Billing Method Cost Optimization Mode	0		
* Minimum Pay-As-You-Go Nodes 0 Percentage of Pay-As-You-G	o Nodes 0 % Lowest-Cost Instance Types 0 Types		
Supplement Preemptible Instances			
Parameter	Description		
Minimum Pay-As-You-Go Nodes	The minimum number of pay-as-you-go instances required by the auto scaling machine group. If the number of pay-as-you-go instances in the auto scaling machine group is less than this value, pay-as-you-go instances are preferentially created.		
Percentage of Pay-As-You-Go Nodes	The proportion of pay-as-you-go instances in the auto scaling machine group after the number of created pay-as-you-go instances reaches the value of Minimum Pay-As-You-Go Nodes.		
Lowest-Cost Instance Types	The number of instance types that have the lowest prices. If preemptible instances are required, the system evenly creates the preemptible instances of the instance types that have the lowest prices. The maximum value is 3.		
Supplement Preemptible Instances	Specifies whether to enable preemptible instance supplementation. If this feature is enabled, the system automatically replaces an existing preemptible instance with a new preemptible instance about five minutes before the existing instance is reclaimed.		

If you do not specify the Minimum Pay-As-You-Go Nodes, Percentage of Pay-As-You-Go Nodes, or Lowest-Cost Instance Types parameter, the machine group is a general cost optimization scaling group. If you specify the parameters, the machine group is a mixed-instance cost optimization scaling group. The two types of cost optimization scaling groups are fully compatible with each other in terms of interfaces and features.

You can use a mixed-instance cost optimization scaling group to achieve the same effect as a specific general cost optimization scaling group by configuring appropriate mixed-instance policies. Examples:

• In a general cost optimization scaling group, only pay-as-you-go instances are created.

In your mixed-instance cost optimization scaling group, set **Minimum Pay-As-You-Go Nodes** to 0, **Percentage of Pay-As-You-Go Nodes** to 100, and **Lowest-Cost Instance Types** to 1.  In a general cost optimization scaling group, preemptible instances are preferentially created.

In your mixed-instance cost optimization scaling group, set Minimum Pay-As-You-Go Nodes to 0, Percentage of Pay-As-You-Go Nodes to 0, and Lowest-Cost Instance Types to 1.

iii. In the Trigger Rules section, specify Trigger Mode.

Scale By Time: For information about this mode, see Configure auto scaling rules by time.

iv. Click OK.

## Enable auto scaling

After you configure auto scaling, find your auto scaling machine group on the **Auto Scaling Settings** page and turn on the switch in the **Auto Scaling Status** column to enable auto scaling.

Auto Scaling Settings								
Machine Groups Auto	o scaling applies only	to pay-as-you-go task	machine groups. You can	enable auto scaling only	for one machine group.			Create Auto Scaling Machine Group
Node Group Name	Instance Type	Current Nodes	Maximum Nodes	Minimum Nodes	Auto Scaling Status	Graceful Deprecation	Cost Reduced By (%)	Actions
12	4 vCPU, 16Gib	2				Disabled		Configure Rules

If you modify the basic information or trigger rules after you enable auto scaling, you must click **Use** Latest Configuration in the Actions column on the Auto Scaling Settings page to make the modifications take effect.

Home Page > Cluster Ma	anagement > Cluster	(C-7BE	> Auto Scaling Se	ttings				
Auto Scaling Settings								
Machine Groups Aut	to scaling applies only	to pay-as-you-go task	machine groups. You car	enable auto scaling only	y for one machine group.			Create Auto Scaling Machine Group
Node Group Name	Instance Type	Current Nodes	Maximum Nodes	Minimum Nodes	Auto Scaling Status	Graceful Deprecation	Cost Reduced By (%)	Actions
ww	8 vCPU, 32Gib	0	10	0		Disabled		Configure Rules
0616	4 vCPU, 16Gib	0	10	2		Disabled	-	Use Latest Configura tion Configure Rules Details
								configure rules becan

# Disable auto scaling

**Notice** You can disable auto scaling for an auto scaling machine group only if the number of instances in the group is 0. To disable auto scaling for an auto scaling machine group that contains instances, you must first configure a scale-in rule for the machine group or set the maximum number of instances to 0. After all instances in the machine group are removed, you can disable auto scaling for it.

Find your auto scaling machine group on the **Auto Scaling Settings** page and turn off the switch in the **Auto Scaling Status** column to disable auto scaling.

If you want to modify instance configurations or your business traffic becomes stable, you can disable auto scaling.

# 5.5.4. Configure auto scaling rules by time

If the computing workloads of a Hadoop cluster fluctuate on a regular basis, you can add and remove a specific number of task nodes at fixed points in time every day, every week, or every month to supplement the computing power. This ensures job completion at low costs.

### Prerequisites

An auto scaling machine group is created. For more information, see Create an auto scaling machine group.

# Configure auto scaling rules by time

For more information about how to configure basic information and cost optimization policies, see Manage auto scaling.

If you want to configure auto scaling rules by time in E-MapReduce (EMR), you can configure the relevant parameters based on the descriptions in the following table.

Auto scaling rules are divided into scale-out rules and scale-in rules. This topic uses a scale-out rule as an example. If you disable auto scaling for a cluster, all auto scaling rules are cleared. If you enable auto scaling for the cluster again, you must reconfigure auto scaling rules.

Add Scaling Rule - Scale	Out By Time			×
* Rule Name:	DThe rule name must be unique.			
	Run Periodically      Run Once     Daily     Y  Run Every     1     Days Scaling Activity Interval: 30 minutes			
* Execute At:	2021-06-17 09:43			
* Valid To:	2021-06-17 09:43			
* Retry Interval(Seconds):	0 Valid values: 0 to 21600 s	econds		
* Scale Out(Nodes):	1 Valid values: 1 to10 Node	5		
* Cooldown(Seconds):	0 Valid values: 0 to 86400 s	econds		
				Cancel
Parameter		Description		
Rule Name		The name of the auto scaling rule. The name of the auto scaling rule. The name	ame m	nust
Rule execution cycle	e	• <b>Run Periodically</b> : Auto scaling is perf specific point in time every day, every every month.		

• **Run Once**: Auto scaling is performed only once at a specific point in time.

Parameter	Description
	The retry interval. Auto scaling may not be performed at the specified point in time due to various reasons. If the retry interval is set, the system tries to perform auto scaling every 30 seconds during the period specified by this parameter until auto scaling succeeds. Valid values: 0 to 21600.
Retry Interval(Seconds)	For example, auto scaling operation A needs to be performed within a specified period of time. If auto scaling operation B is still in progress or is in the cooldown state during this period, operation A cannot be performed. In this case, the system tries to perform operation A every 30 seconds within the retry interval you specified. If required conditions are met, the cluster immediately runs auto scaling.
Scale Out(Nodes)	The number of task nodes you want to add to the cluster each time an auto scaling rule is triggered.
Cooldown(Seconds)	The interval between two auto scaling activities. Auto scaling is forbidden during the cooldown.

# Configure the specifications of nodes

You can specify the hardware specifications of the nodes that are used to scale in or scale out a cluster. You can configure the specifications only before you enable auto scaling. You are not allowed to modify the specifications after you enable auto scaling. If modifications are required due to special circumstances, disable auto scaling, modify the specifications, and then enable auto scaling again.

- The system automatically searches for the instance types that match the vCPU and memory specifications you specified and displays the instance types in the Instance Type section. You must select one or more instance types in the Instance Type section so that the cluster can be scaled based on the selected instance types.
- You can select a maximum of three instance types. This prevents auto scaling failures due to insufficient Elastic Compute Service (ECS) resources.
- Regardless of whether you select an ultra disk or a standard SSD, the minimum size of a data disk is 40 GB.

# 5.5.5. View auto scaling history

This topic describes how to view auto scaling history after auto scaling is triggered.

# Prerequisites

Auto scaling rules are configured. For more information, see Create a cluster.

# Procedure

1. Go to the Cluster Overview page.

i.

ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.

- iii. Click the Cluster Management tab.
- iv. On the Cluster Management page, find your cluster and click Details in the Actions column.
- 2. In the left-side navigation pane, choose **Auto Scaling > History**.

You can view information about each auto scaling activity, including the number and status of nodes after each auto scaling activity. An auto scaling activity can be in one of the following states:

- InProgress: The auto scaling activity is in progress.
- **Successful:** All nodes specified in the auto scaling rule are added or removed during the auto scaling activity.
- **part Of Success**: During the auto scaling activity, some of the nodes specified in the auto scaling rule are added or removed. The other nodes fail to be added or removed because of an insufficient disk quota or Elastic Compute Service (ECS) resources.
- Failed: No nodes specified in the auto scaling rule are added or removed during the auto scaling activity.
- **Rejected**: The auto scaling activity is rejected because the number of nodes in the auto scaling machine group is larger than the upper limit or less than the lower limit if the auto scaling activity is performed. Another possible cause is that the previous auto scaling activity has not been terminated when the auto scaling rule is triggered.

# 5.5.6. Configure auto scaling monitoring alerts

In the Alibaba Cloud CloudMonitor console, you can configure event alerts for auto scaling activities in E-MapReduce (EMR). After you configure an event alert, if auto scaling fails or is rejected, CloudMonitor notifies the contacts in the contact group to handle the issue in a timely manner.

### Procedure

- 1. Log on to the CloudMonitor console.
- 2. In the left-side navigation pane, click Event Monitoring.
- 3. On the Event Monitoring page, click the Alert Rules tab.
- 4. On the Alert Rules tab, click Create Event Alert.
- 5. In the Create / Modify Event Alert panel, configure the alert rule parameters for system events.

Section	Parameter	Description
Basic Information	Alert Rule Name	The name of the event alert rule.
	Product Type	Select E-MapReduce.
	Event Type	Select SCALING.
	Event Level	Select CRITICAL.
Event alert	Event Name	Select All Events or a specific event.
	Resource Range	The range of the resources to which the event alert rule is applied.

For the configuration of other parameters, see Create a system event-triggered alert rule.

6. Click OK.

You can view the created rule on the Alarm Rules tab.

# 5.6. Expand disk capacity

界面上找不到该topic中的部分操作描述,跟需求人沟通,该topic先skip掉不处理,后续修改中文后重新提翻 处理

# 5.7. Add node groups

This topic describes how to add a task node group to an E-MapReduce (EMR) cluster.

### Prerequisites

A Hadoop cluster is created. For more information, see Create a cluster.

### Context

You can add node groups to meet your requirements for different instances. For example, you can use memory optimized instances with 1 vCPU and 8 GiB of memory to support big data offline processing and use compute optimized instances with 1 vCPU and 2 GiB of memory to support model training.

Notice If you want to add a core node group to a Hadoop cluster, submit a ticket.

### Limits

You can add a maximum of 10 task node groups to a Hadoop cluster.

#### Procedure

- 1. Go to the Cluster Overview page of your cluster.
  - i.
  - ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
  - iii. Click the Cluster Management tab.
  - iv. On the Cluster Management page, find your cluster and click Details in the Actions column.
- 2. On the Cluster Overview page, choose Change Configuration > Scale Out.
- 3. In the Scale Out dialog box, click the TASK (Task Instance Group) tab.
- 4. On the TASK (Task Instance Group) tab, perform the following operations:
  - i. (Optional)Click the Create Node Group tab.
  - ii. On the Create Node Group tab, specify Node Group Name.

(?) Note The node group name must be unique in the cluster.

- iii. Configure other parameters based on your requirements.
- iv. Read the terms of service and select E-MapReduce Service Terms.
- v. Click Create Node Group.

After the node group is added, you can view the node group information in the **Instance Info** section of the **Cluster Overview** page.

# 5.8. Remove abnormal ECS instances

When an exception occurs on Elastic Compute Service (ECS) instances in an E-MapReduce (EMR) cluster and you do not need to continuously use these instances, you can remove them. For example, you can remove the ECS instances that are in the Stopped state. This topic describes how to remove an abnormal ECS instance.

# Prerequisites

An EMR cluster is created. For more information, see Create a cluster.

## Procedure

1. Go to the Cluster Overview page.

i.

- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the Cluster Management tab.
- iv. On the **Cluster Management** page, find your cluster and click **Details** in the Actions column.
- 2. In the left-side navigation pane of the Cluster Overview page, click Instances.
- 3. On the Instances page, find the ID of the ECS instance that you want to remove and click **Delete** in the Action column.

Instances							Update Instance Inf
Instance Name: Please Input	ECS Instance ID:	Internal Network IP:	Pu	olic Network IP:	Search	Reset	
Instance Name	ECS ID	IP Addresses	Role 17	Node Group	Billing Method	Туре	Action
emr-header-2	Q	Sector and the sector	MASTER	MASTER	Subscription	ECS Instance Type: ecs.g6.xlarge Data Disk Type: SSD 80GB X 1Disks System Disk Type: SSD 120GB X 1Disks	
emr-header-1	Q		MASTER	MASTER	Subscription	CPU: 4 vCores   Memory: 16G ECS Instance Type: ecs.g6.xlarge Data Disk Type: SSD 80GB X 1Disks System Disk Type: SSD 120GB X 1Disks	
emr-worker-2	Q		CORE	CORE	Subscription	CPU: 4 vCores   Memory: 16G ECS Instance Type: ecs.g6.xlarge Data Disk Type: SSD 80GB X 4Disks System Disk Type: SSD 120GB X 1Disks	Delete

Onte You can remove only abnormal ECS instances.

4. Click OK.

The ECS instance is removed.

# 5.9. Upgrade node configurations

If the vCPUs or memory of the master node group or a core node group in an E-MapReduce (EMR) cluster is insufficient, you can upgrade the node configurations of the node group. This topic describes how to upgrade node configurations.

### Prerequisites

An EMR cluster is created. For more information, see Create a cluster.

### Limits

<sup>&</sup>gt; Document Version: 20220517

• You can upgrade only the node configurations of subscription clusters.

**?** Note If you use a high-availability cluster, only one of the master nodes is restarted. During the restart process, the other master node still runs to ensure normal services.

- Elastic Compute Service (ECS) instances with local disks, such as instances of the d1 and i2 instance families, do not support configuration upgrade. You can only add nodes when vCPUs or memory is insufficient.
- ECS instances that do not use local disks support configuration upgrade. However, they do not support configuration downgrade.

#### Procedure

1. Go to the Upgrade dialog box.

i.

- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the Cluster Management tab.
- iv. On the **Cluster Management** page, find the cluster for which you want to upgrade node configurations and click **Details** in the Actions column.
- v. On the Cluster Overview page, choose Change Configuration > Upgrade.
- 2. Upgrade node configurations.
  - i. In the Upgrade dialog box, modify the configurations based on your business requirements.
  - ii. Read the terms of service and select E-MapReduce Service Terms.
  - iii. Click OK.

An order is generated after a short period of time.

iv. Complete the payment.

After you complete the payment, a message that indicates the upgrade is successful and you must restart the node group is displayed in the **Instance Info** section of the **Cluster Overview** page.

I Instance Info 🕑 The configuration is upgraded. Res	start the node group for the new configu	ration to take effect.
Master Instance Group (MASTER) Subscription	ECS ID	Deployment Status
<ul> <li>Instances: 1</li> <li>CPU: 8vCores</li> <li>Memory: 32GB</li> </ul>	i-bp1gqhî	Normal
♦ Data Disk Type: 80GB ESSD Disk*1		

3. In the Instance Info section of the Cluster Overview page, click the message The configuration is upgraded. Restart the node group for the new configuration to take effect.

Notice When you restart a cluster, the ECS instances of the cluster are restarted. During the restart, the big data services deployed on the ECS instances are unavailable. Therefore, before you restart the cluster, make sure that the service unavailability does not affect your business.

Current Cluster: C		
	be restarted. All services will be suspended during the restart process. Make sure that your business will not be affected	
pefore you restart the	node.	
Rolling Restart @	Restart Scaled Nodes Only @	

Parameter	Description
Rolling Restart	<ul> <li>Rolling Restart is selected by default. If you select Rolling Restart, the system restarts an ECS instance only after the previous ECS instance is restarted and all big data services deployed on it are restored. It takes about 5 minutes to restart an ECS instance.</li> <li>If you clear Rolling Restart, the system restarts all ECS instances at the same time.</li> </ul>
Restart Scaled Nodes Only	<ul> <li>Restart Scaled Nodes Only is selected by default. If you select Restart Scaled Nodes Only, the system restarts only the nodes whose disk space is expanded or configurations are upgraded.</li> <li>For example, if you have upgraded only the configurations of the master node group, the system restarts the nodes in the master node group and do not restart those in core node groups.</li> <li>If you clear Restart Scaled Nodes Only, the system restarts all the nodes of the cluster.</li> </ul>

#### 4. Click OK.

In the **Instance Info** section of the **Cluster Overview** page, the message **Restarting** is displayed for the node group whose node configurations are upgraded.

Instance Info 😮 Restarting server group
Master Instance Group (MASTER) (Restarting Subscription
ECS Instance Type: ecs.g5.2xlarge
+Instances: 1
♦ CPU: 8vCores
♦ Memory: 32GB
Data Disk Type: 80GB ESSD Disk*1

After the message disappears, the restart is complete and the configurations have taken effect. You can log on to the cluster to verify the configurations.

### ? Note

- If you have upgraded only vCPU configurations, no more operations are required.
- If you have upgraded only memory configurations or both vCPU and memory configurations, we recommend that you perform steps in Modify configurations for YARN. This way, the YARN service can use the added resources.

# Modify configurations for YARN

- 1. In the left-side navigation pane, choose **Cluster Service > YARN**.
- 2. Modify vCPU configurations.
  - i. On the YARN service page, click the **Configure** tab.
  - ii. In the **Configuration Filter** section, search for the **yarn.nodemanager.resource.cpu-vcores** parameter. Then, change its value based on your business requirements.

In compute-intensive scenarios, we recommend that you set this parameter to a value that is equal to the number of vCPUs of each ECS instance. In scenarios that involve both compute-intensive and I/O intensive tasks, we recommend that you set this parameter to a value that ranges from the number of vCPUs of each ECS instance to twice the number. For example, if you use 32-vCPU ECS instances in compute-intensive scenarios, set this parameter to 32. If you use 32-vCPU ECS instances in scenarios that involve both compute-intensive and I/O intensive tasks, set this parameter to a value that ranges from 32 to 64.

- 3. Modify memory configurations.
  - i. On the YARN service page, click the **Configure** tab.
  - ii. In the **Configuration Filter** section, search for the **yarn.nodemanager.resource.memory-mb** parameter. Then, change its value to the memory size of each ECS instance multiplied by 0.8. Unit : MiB.

For example, if the memory size of each ECS instance is 32 GiB, set this parameter to 26214.

- 4. Save the configurations.
  - i. Click Save in the upper-right corner of the Service Configuration section.
  - ii. In the **Confirm Changes** dialog box, specify **Description** and click **OK**.
- 5. Issue the configurations.
  - i. On the YARN service page, choose Actions > Configure All Components.
  - ii. In the Cluster Activities dialog box, specify Description and click OK.
  - iii. In the **Confirm** message, click **OK**.

You can click **History** in the upper-right corner to view the status of the Configure YARN task. After the task enters the **Successful** state, perform the following step to restart the YARN service.

- 6. Restart the YARN service.
  - i. On the YARN service page, choose Actions > Restart All Components.
  - ii. In the Cluster Activities dialog box, specify Description and click OK.
  - iii. In the **Confirm** message, click **OK**.

You can click **History** in the upper-right corner to view the status of the Restart YARN task. After the task enters the **Successful** state, the YARN service is restarted.

# 5.10. Status list

# Cluster status list

⑦ Note You can view the cluster status in the cluster list or on the cluster details page.

Status	Status code	Description		
Creating	CREAT ING	The cluster is being created. The creation task includes two stages: creating physical ECS machines and activating Spark clusters. It takes a moment for the clusters to start running.		
Failed	CREAT E_FAILED	An exception occurred during creation. The ECS instance that you have created automatically rolls back. You can click the question mark (?) to the right of the status on the cluster list page to view exception details.		
Running	RUNNING	The computing cluster is running.		
Idle	IDLE	The cluster is not running any execution plan.		
Releasing	RELEASING	Click Release in the status list to set the cluster to this status. This status indicates that the cluster is in the releasing process. It may take a moment to complete this process.		
Release Failed	RELEASE_FAILED	An exception occurred when releasing the cluster. You can click the question mark (?) to the right of the status on the cluster list page to view the exception details. When the cluster is in this status, click Release to release the cluster again.		
Released	RELEASED	The computing cluster and the ECS instance that hosts the computing cluster have been released.		
Abnormal	ABNORMAL	Unrecoverable errors occurred on one or more nodes in the computing cluster. Click Release to release the cluster.		

# Job status list

? Note	View job status in the job status list	
--------	--	--

Status	Description
Status	Description
Ready	The creation information is complete, correct, and successfully saved. The job is ready to be added to the submission queue. It may take a moment for the job to change its status to Submitting.
Submitting	The job is in the submission queue of the computing cluster. It has not been submitted to the computing cluster.
Failed	An exception occurred when submitting the job to the computing cluster. You need to clone and submit the job if you want to submit this job again.
Running	The job is running in the cluster. Wait a moment and click the corresponding log button in the job list to view output log entries in real time.
Succeeded	The job has been successfully executed in the cluster. Click the corresponding log button to view the log entry.
Failed	An exception occurred when executing the job. Click the corresponding log button to view the log entry.

# 5.11. Cluster O&M guide

This topic describes how to view environment variables and how to start and stop the service processes of E-MapReduce (EMR) clusters. You can maintain your clusters based on the instructions in this topic.

# Prerequisites

An EMR cluster is created. For more information, see Create a cluster.

# View environment variables

- 1. Log on to your cluster in SSH mode. For more information, see Log on to a cluster.
- 2. Run the env command.

The configurations of environment variables are displayed. Example:

PRESTO HOME=/usr/lib/presto-current TEZ\_CONF\_DIR=/etc/ecm/tez-conf HUDI HOME=/usr/lib/hudi-current XDG SESSION ID=35918 SPARK HOME=/usr/lib/spark-current HOSTNAME=emr-header-1.cluster-23\*\*\*\* HADOOP LOG DIR=/var/log/hadoop-hdfs SMARTDATA CONF DIR=/usr/lib/b2smartdata-current//conf ECM\_AGENT\_STACK\_CACHE\_DIR=/usr/lib/emr/ecm-agent/cache/ecm TERM=xterm SHELL=/bin/bash HUE CONF DIR=/etc/ecm/hue-conf HADOOP HOME=/usr/lib/hadoop-current FLOW AGENT CONF DIR=/etc/ecm/flow-agent-conf HISTSIZE=1000 YARN PID DIR=/usr/lib/hadoop-current/pids ECM AGENT CACHE DIR=/usr/lib/emr/ecm-agent/cache SSH CLIENT=1.80.115.185 26289 22 HADOOP PID DIR=/usr/lib/hadoop-current/pids EMR HOME DIR=/usr/lib/emr HADOOP MAPRED PID DIR=/usr/lib/hadoop-current/pids SQOOP CONF DIR=/etc/ecm/sqoop-conf SQOOP HOME=/usr/lib/sqoop-current BIGBOOT MONITOR HOME=/usr/lib/b2monitor-current/ HCAT HOME=/usr/lib/hive-current/hcatalog DATA FACTORY CONF PATH=/etc/ecm/datafactory-conf HIVE HOME=/usr/lib/hive-current PWD=/root JAVA HOME=/usr/lib/jvm/java-1.8.0 EMR DATA DIR=/usr/lib/emr/data B2MONITOR CONF DIR=/usr/lib/b2monitor-current//conf HISTCONTROL=ignoredups SPARK\_PID\_DIR=/usr/lib/spark-current/pids SHLVL=1 HOME=/root HADOOP MAPRED LOG DIR=/var/log/hadoop-mapred ALLUXIO CONF DIR=/etc/ecm/alluxio-conf ECM AGENT LOG DIR=/usr/lib/emr/ecm-agent/log TEZ HOME=/usr/lib/tez-current DATA FACTORY HOME=/usr/lib/datafactory-current LOGNAME=root EMR LOG DIR=/usr/lib/emr/log EMR TMP DIR=/usr/lib/emr/tmp XDG RUNTIME DIR=/run/user/0 ECM AGENT HOME DIR=/usr/lib/emr/ecm-agent B2SDK CONF DIR=/usr/lib/b2smartdata-current/conf HIVE CONF DIR=/etc/ecm/hive-conf \_=/usr/bin/env

# Log on to the built-in MySQL

- 1. Log on to your cluster in SSH mode. For more information, see Log on to a cluster.
- 2. Run the following command to log on to the built-in MySQL:

```
mysql -uroot -pEMRroot1234
```

**Note** The username that is used to log on to the built-in MySQL is root and the password is EMRroot 1234.

## Start and stop a single service process in the EMR console

You can start, stop, or restart a service process in the EMR console. Operations on all service processes are similar. This section describes how to start, stop, and restart the HDFS service process DataNode on the emr-worker-1 node of your cluster.

- 1. Go to the Cluster Overview page.
  - i.
  - ii.
  - ...
  - iii.
  - iv. On the **Cluster Management** page, find your cluster and click **Details** in the Actions column.
- 2. In the left-side navigation pane, choose Cluster Service > HDFS.
- 3. Click the **Component Deployment** tab.

The HDFS service processes of the cluster are displayed.

- 4. Manage the DataNode process on the *emr-worker-1* node.
  - i. Start the process.
    - a. Find the process and click **Start** in the **Actions** column.
    - b. In the Cluster Activities dialog box, specify Description and click OK.
    - c. In the **Confirm** message, click **OK**.
  - ii. Restart the process.
    - a. Find the process and click Restart in the Actions column.
    - b. In the Cluster Activities dialog box, specify Description and click OK.
    - c. In the **Confirm** message, click **OK**.
  - iii. Stop the process.
    - a. Find the process and click **Stop** in the **Actions** column.
    - b. In the Cluster Activities dialog box, specify Description and click OK.
    - c. In the **Confirm** message, click **OK**.

# Manage multiple service processes at the same time in the EMR console

This section describes how to restart the DataNode processes of HDFS on all the nodes of your cluster at the same time.

- 1. Go to the Cluster Overview page.
  - i.
  - ii.
  - iii.
  - iv. On the **Cluster Management** page, find your cluster and click **Details** in the Actions column.

- 2. In the left-side navigation pane, choose **Cluster Service > HDFS**.
- 3. Click the Component Deployment tab.

The HDFS service processes of the cluster are displayed.

- 4. Choose Actions > Restart DataNode in the upper-right corner.
  - i. In the **Cluster Activities** dialog box, specify **Description** and click **OK**.

ii. In the **Confirm** message, click **OK**.

Notice After you perform a rolling restart, you cannot perform a common restart on the same processes. Otherwise, an error is reported.

## Start and stop a single service process by using the CLI

• YARN

Account: hadoop

- ResourceManager (master node)
  - Start ResourceManager

/usr/lib/hadoop-current/sbin/yarn-daemon.sh start resourcemanager

Stop ResourceManager

/usr/lib/hadoop-current/sbin/yarn-daemon.sh stop resourcemanager

#### • NodeManager (core node)

#### Start NodeManager

/usr/lib/hadoop-current/sbin/yarn-daemon.sh start nodemanager

Stop NodeManager

/usr/lib/hadoop-current/sbin/yarn-daemon.sh stop nodemanager

#### • JobHistory Server (master node)

#### Start JobHistory Server

/usr/lib/hadoop-current/sbin/mr-jobhistory-daemon.sh start historyserver

#### Stop JobHistory Server

/usr/lib/hadoop-current/sbin/mr-jobhistory-daemon.sh stop historyserver

#### • WebProxyServer (master node)

Start WebProxyServer

/usr/lib/hadoop-current/sbin/yarn-daemon.sh start proxyserver

#### Stop WebProxyServer

/usr/lib/hadoop-current/sbin/yarn-daemon.sh stop proxyserver

HDFS

#### Account: hdfs

#### • NameNode (master node)

Start NameNode

/usr/lib/hadoop-current/sbin/hadoop-daemon.sh start namenode

Stop NameNode

/usr/lib/hadoop-current/sbin/hadoop-daemon.sh stop namenode

#### • DataNode (core node)

Start DataNode

/usr/lib/hadoop-current/sbin/hadoop-daemon.sh start datanode

#### Stop DataNode

/usr/lib/hadoop-current/sbin/hadoop-daemon.sh stop datanode

#### • Hive

#### Account: hadoop

#### MetaStore (master node)

// Start MetaStore. You can set the memory size to a larger value based on your business
requirements.
UNDOOD UNDOCUMENTEEL2 (way/lib/bing guarant/bin/bing gamming metastang ) (way/log/bing/metastang)

HADOOP\_HEAPSIZE=512 /usr/lib/hive-current/bin/hive --service metastore >/var/log/hive/met astore.log 2>&1 &

#### • HiveServer2 (master node)

```
// Start HiveServer2.
HADOOP_HEAPSIZE=512 /usr/lib/hive-current/bin/hive --service hiveserver2 >/var/log/hive/h
iveserver2.log 2>&1 &
```

#### HBase

#### Account: hdfs

Notice You can perform the following operations only if the HBase service is deployed in your cluster. Otherwise, an error is reported.

#### • HMaster (master node)

Start HMaster

/usr/lib/hbase-current/bin/hbase-daemon.sh start master

Restart HMaster

/usr/lib/hbase-current/bin/hbase-daemon.sh restart master

#### Stop HMaster

/usr/lib/hbase-current/bin/hbase-daemon.sh stop master

#### • HRegionServer (core node)

#### Start HRegionServer

/usr/lib/hbase-current/bin/hbase-daemon.sh start regionserver

#### Restart HRegionServer

/usr/lib/hbase-current/bin/hbase-daemon.sh restart regionserver

#### Stop HRegionServer

/usr/lib/hbase-current/bin/hbase-daemon.sh stop regionserver

#### • Thrift Server (master node)

#### Start Thrift Server

/usr/lib/hbase-current/bin/hbase-daemon.sh start thrift -p 9099 >/var/log/hive/thrifts erver.log 2>&1 &

#### Stop Thrift Server

/usr/lib/hbase-current/bin/hbase-daemon.sh stop thrift

#### • Hue

#### Account: hadoop

#### • Start Hue

```
su -l root -c "${HUE_HOME}/build/env/bin/supervisor >/dev/null 2>&1 &"
```

#### • Stop Hue

ps aux   grep hue	//	Find	the	Hue	process.
kill -9 huepid	11	Kill	the	Hue	process.

#### • Zeppelin

#### Account: hadoop

#### • Start Zeppelin

```
// You can set the memory size to a larger value based on your business requirements.
su -l root -c "ZEPPELIN_MEM=\"-Xmx512m \" ${ZEPPELIN_HOME}/bin/zeppelin-daemon.sh
start"
```

#### • Stop Zeppelin

su -l root -c "\${ZEPPELIN\_HOME}/bin/zeppelin-daemon.sh stop"

• Presto

Account: hdfs

#### • PrestoServer (master node)

#### Start PrestoServer

/usr/lib/presto-current/bin/launcher --config=/usr/lib/presto-current/etc/coordinatorconfig.properties start

#### Stop PrestoServer

/usr/lib/presto-current/bin/launcher --config=/usr/lib/presto-current/etc/coordinatorconfig.properties stop

#### • PrestoServer (core node)

#### Start PrestoServer

/usr/lib/presto-current/bin/launcher --config=/usr/lib/presto-current/etc/worker-confi
g.properties start

Stop PrestoServer

/usr/lib/presto-current/bin/launcher --config=/usr/lib/presto-current/etc/worker-confi
g.properties stop

## Manage multiple service processes at the same time by using the CLI

You can manage service processes on all the core nodes of a cluster at the same time by running script commands. In an EMR cluster, all the worker nodes that are run by the hadoop account or the hdfs account are interconnected with the master node.

For example, you can run the following command to stop the NodeManager processes of all the core nodes. In this example, the number of core nodes is 10.

```
for i in `seq 1 10`;do ssh emr-worker-$i /usr/lib/hadoop-current/sbin/yarn-daemon.sh stop nod
emanager;done
```

# 5.12. Release a cluster

This topic describes how to release an E-MapReduce (EMR) cluster if it is no longer needed. You can release only pay-as-you-go clusters.

#### Prerequisites

- An EMR cluster is created. For more information, see Create a cluster.
- The cluster you want to release is in the Initializing, Running, or Idle state.

#### Limits

You can release only pay-as-you-go clusters. If you want to release a subscription cluster, submit a ticket.

# Precautions

After you confirm the release of a cluster, the system performs the following operations to handle the cluster:

- 1. Forcibly terminates all jobs in the cluster.
- 2. Terminates and releases all ECS instances that are created for the cluster.

The time needed to terminate and release the ECS instances depends on the size of the cluster but does not exceed 5 minutes. In most cases, the ECS instances can be released in seconds.

### Procedure

1. Go to the Cluster Management page.

i.

- ii. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- iii. Click the Cluster Management tab.
- 2. Find the cluster you want to release and choose **More > Release** in the Actions column.

Alternatively, find the cluster you want to release and click **Details** in the Actions column. On the **Cluster Overview** page, choose **Instance Status > Release** in the upper-right corner.

3. In the Cluster Management-Release message, click Release.

# 6.FAQ about cluster management

This topic provides answers to some frequently asked questions about cluster management.

- Error messages
  - How do I fix the error "The specified zone is not available for purchase"?
  - How do I fix the error "The request processing has failed due to some unknown error, exception or failure"?
  - How do I fix the error "The Node Controller is temporarily unavailable"?
  - How do I fix the error "No quota or zone is available"?
  - How do I fix the error "The specified InstanceType is not authorized for use"?
  - How do I fix the error "The specified DataDisk Size beyond the permitted range, or the capacity of snapshot exceeds the size limit of the specified disk category"?
  - How do I fix the error "Your account does not have enough balance"?
  - How do I fix the error "The maximum number of Pay-As-You-Go instances is exceeded: create ecs vcpu quot a per region limited by user quot a [xxx]"?
  - How do I fix the error "FAILED: SemanticException org.apache.hadoop.hive.ql.metadata.HiveException: java.lang.RuntimeException: Unable to instantiate org.apache.hadoop.hive.ql.metadata.SessionHiveMetaStoreClient"?
  - How do I fix the error "The specified instance Type exceeds the maximum limit for the PostPaid instances"?
- Billing
  - Why do I still receive renewal notifications after I renew my cluster?
  - Do EMR clusters support auto-renewal?
  - How do I apply for a refund for an EMR cluster?
- Product features
  - Do I need to handle a cluster creation failure?
  - Does EMR support preemptible instances?
  - How do I apply for an instance type with advanced configurations?
  - What do I do if the disk capacity of an EMR cluster is insufficient?
  - What do I do if the disk capacity of an EMR cluster is excessively large?
  - What do I do if the computing capability of an EMR cluster is excessively low?
  - What do I do if the computing capability of an EMR cluster is excessively high?
  - What do I do if the version of a component in an EMR cluster does not meet my business requirements?
  - How do I convert a non-high-availability (non-HA) cluster to an HA cluster?
  - How do I deploy third-party software or third-party services on EMR?
  - I upgraded the node configurations of a node group to increase the memory and vCPUs. Why are the added resources not applied to the YARN service?
  - How do I log on to a core node?
  - What is the division of work in an EMR cluster?
  - How do I handle disk exceptions in a Kafka cluster?
  - Can I install additional software on the master node of an EMR cluster?

- Do services on each node automatically start when I power on the server? Are services automatically resumed after they are interrupted unexpectedly?
- What is the port number of the HBase Thrift server?
- Why can I not set EMR Version to EMR-3.4.3 when I create a cluster?
- What are the differences between EMR and MaxCompute?
- Does EMR support automatic storage balancing? How do I manually rebalance storage?
- What do I do if resources for ApplicationMasters of a job that runs on YARN exceed the upper limit of a YARN queue?
- What are the differences between the two job submission modes Worker Node and Header/Gateway Node?
- How do Laccess the web UIs of open source services?
- Metadata management
  - What do I do if the "Failed to get schema version" error is reported when the MetaStore client is initialized?
  - What do I do if Hive metadata contains Chinese characters, such as Chinese characters in column comments and partition names?
- Permissions
  - How do I allow ECS instances in the classic network and ECS instances of EMR clusters in a virtual private cloud (VPC) to access each other?
  - How do I connect EMR clusters that belong to different VPCs of the same account?
  - How do I isolate the Object Storage Service (OSS) data of different Resource Access Management (RAM) users?

# How do I fix the error "The specified zone is not available for purchase"?

Pay-as-you-go Elastic Compute Service (ECS) instances are unavailable in the zone you selected. We recommend that you select a different zone.

# How do I fix the error "The request processing has failed due to some unknown error, exception or failure"?

Try again later. If the error persists, submit a ticket. You can also submit a ticket right after the error is reported.

# How do I fix the error "The Node Controller is temporarily unavailable"?

Wait a moment and create a cluster again.

### How do I fix the error "No quota or zone is available"?

ECS instances in the zone you selected are insufficient. Create a cluster again, but select a different zone or use the default zone when you create the cluster.

# How do I fix the error "The specified InstanceType is not authorized for use"?

> Document Version: 20220517

Submit applications to use pay-as-you-go instances with advanced configurations (each instance with more than eight vCPUs). Click submit a ticket to submit the applications. The supported specifications include eight vCPUs and 16 GiB of memory, eight vCPUs and 32 GiB of memory, 16 vCPUs and 32 GiB of memory, and 16 vCPUs and 64 GiB of memory.

# How do I fix the error "The specified instance Type exceeds the maximum limit for the PostPaid instances"?

Possible causes:

- Your quot a for pay-as-you-go instances has been reached.
- You do not have the permissions to create the current type of instance.

Solutions:

- Submit an application to increase your quota.
- Go to the ECS console and grant your account the permissions to create the current type of instance.

### How do I log on to a core node?

1. On the master node, run the following command to switch to the hadoop user:

su hadoop

2. Log on to the core node in password-free mode.

ssh emr-worker-1

3. Run the following sudo command to obtain the root permissions:

sudo su - root

# Why do I still receive renewal notifications after I renew my cluster?

Cause: You are charged for the E-MapReduce (EMR) service and the ECS instances that are deployed in your cluster. If you renew only the ECS instances, you can still receive renewal notifications for the EMR service.

Solution: Go to the **Cluster Overview** page for your EMR cluster. Select **Renewal** from the **Renewal** drop-down list in the upper-right corner. On the page that appears, check the expiration time of the ECS instances and EMR resources.

### Do EMR clusters support auto-renewal?

Yes, both EMR resources and ECS instances can be automatically renewed.

# How do I apply for a refund for an EMR cluster?

Submit a ticket to apply for a refund.

### Do I need to handle a cluster creation failure?

No, you do not need to handle cluster creation failures. No computing resources are created if you fail to create a cluster. The cluster record is automatically removed from the cluster list in the EMR console after three days.

# What is the division of work in an EMR cluster?

A standard EMR cluster consists of a single master node and multiple core nodes. Only core nodes store data and implement data computing. For example, a cluster consists of three instances. Each instance has four vCPUs and 8 GiB of memory. One instance serves as the master node and the other two instances serve as core nodes. The available computing resources of this cluster are two instances, each with four vCPUs and 8 GiB of memory.

# How do I handle disk exceptions in a Kafka cluster?

Cause: The disk is full or damaged.

#### Solution:

- If your disk is full, perform the following operations:
  - i. Log on to the server.
  - ii. Find the fully occupied disk and delete unnecessary data to free up some of the disk space. Take note of the following rules:
    - Do not delete Kafka data directories. Otherwise, you may lose all of your data.
    - Do not delete Kafka topics, such as consumer\_offsets and schema.
    - Find the topics that occupy a large space or that you no longer need. Delete historical log
      segments and the index and timeindex files of the segments from some partitions of the topics.

iii. Restart the Kafka broker.

- If your disk is damaged, perform the following operations:
  - If no more than 25% of disks are damaged on a machine, you do not need to take an action.
  - If more than 25% of disks are damaged on a machine, submit a ticket.

# How do I fix the error "The specified DataDisk Size beyond the permitted range, or the capacity of snapshot exceeds the size limit of the specified disk category"?

The disk size you specified is too small. We recommend that you set the disk size to a value larger than 40 GB.

# How do I fix the error "Your account does not have enough balance"?

The account balance is insufficient. Top up your account and try again.

# How do I fix the error "The maximum number of Pay-As-You-Go instances is exceeded: create ecs vcpu quota per region limited by user quota [xxx]"?

Your quot a for pay-as-you-go ECS instances has been reached. You can release some existing pay-asyou-go ECS instances or submit an application to increase your quot a.

# Can I install additional software on the master node of an EMR cluster?

We recommend that you do not install additional software. The installation may affect the stability and reliability of the cluster.

# Do services on each node automatically start when I power on the server? Are services automatically resumed after they are interrupted unexpectedly?

Yes, services automatically start when you power on the server, and services are automatically resumed after they are interrupted unexpectedly.

## What is the port number of the HBase Thrift server?

The port number of the HBase Thrift server is 9099.

## Does EMR support preemptible instances?

If you enable the auto scaling feature for a cluster, you can use preemptible instances. For more information, see Manage auto scaling.

#### How do I fix the error " FAILED: SemanticException

org.apache.hadoop.hive.ql.metadata.HiveException: java.lang.RuntimeException: Unable to

instantiate org.apache.hadoop.hive.ql.metadata.SessionHiveMetaStoreClient "?

During cluster creation, **Unified Metabases** is selected for Type. An error is reported during the execution of a Hive job.

The cluster fails to be created because the cluster does not have an elastic IP address. submit a ticket.

## Why can I not set EMR Version to EMR-3.4.3 when I create a cluster?

EMR is updated periodically. Some earlier versions are deprecated. Check whether the versions of services such as Hive and Spark in the available EMR versions meet your business requirements. If you want to use a deprecated EMR version, submit a ticket.

### What are the differences between EMR and MaxCompute?

Both of them are big data processing solutions. EMR is a big data platform that is built completely based on open source technologies. It is fully compatible with open source software. MaxCompute is a platform developed by Alibaba Cloud and is not open source. It offers easy-to-use features based on the encapsulation and low costs of operations and maintenance.

# Does EMR support automatic storage balancing? How do I manually rebalance storage?

Automatic storage balancing is not supported. You can perform the following steps to manually rebalance the storage of a cluster:

1.

- 2. In the top navigation bar, select the region where your cluster resides and select a resource group based on your business requirements.
- 3. Click the Cluster Management tab.
- 4. Find the cluster whose storage you want to rebalance, and click **Details** in the **Actions** column.
- 5. In the left-side navigation pane, choose **Cluster Service > HDFS**.
- 6. In the upper-right corner of the page that appears, choose **Actions > Rebalance**.
- 7. In the **Cluster Activities** dialog box, specify the related parameters and click **OK**.
- 8. In the **Confirm** message, click **OK**.

# How do I apply for an instance type with advanced configurations?

#### Submit a ticket.

# What do I do if the disk capacity of an EMR cluster is insufficient?

Expand the capacity of a disk or add core nodes to the EMR cluster. You are not allowed to add disks to a node of an EMR cluster.

# What do I do if the disk capacity of an EMR cluster is excessively large?

Purchase a new cluster and release the original one. You are not allowed to scale down the disk capacity of an EMR cluster. For more information, see Create a cluster.

# What do I do if the computing capability of an EMR cluster is excessively low?

Add task nodes to the cluster in the EMR console. For more information, see Scale out a cluster.

# What do I do if the computing capability of an EMR cluster is excessively high?

Resolve this issue based on the billing method of your cluster.

- For a pay-as-you-go cluster, remove one or more task nodes from the cluster in the EMR console.
- For a subscription cluster, stop the NodeManager of YARN on a specific task node, change the billing method of the ECS instance that serves as the task node to pay-as-you-go in the ECS console, and then release the instance.

# What do I do if the version of a component in an EMR cluster does not meet my business requirements?

Purchase a cluster of a later version. You are not allowed to update a specific component of an existing cluster. For more information, see Create a cluster.

# How do I convert a non-high-availability (non-HA) cluster to an HA cluster?

Non-HA clusters cannot be converted into HA clusters. We recommend that you purchase an HA cluster.

# How do I deploy third-party software or third-party services on EMR?

We recommend that you use bootstrap actions to install third-party software or third-party services when you create a cluster. If you manually install third-party software or third-party services after you create a cluster, you must reinstall the software or services when you add nodes.

# What do I do if an ECS instance of an EMR cluster reports an alert that an error occurred and the instance must be re-deployed?

Submit a ticket.

# I upgraded the node configurations of a node group to increase the memory and vCPUs. Why are the added resources not applied to the YARN service?

The memory and vCPU resources of the YARN service on an EMR cluster are the sum of the resources of the NodeManager on each node. However, the memory and vCPU configurations of the NodeManager on each node do not automatically change after you upgrade the node configurations. You can manually modify the configurations and restart the NodeManager on each node. This way, the total memory and vCPU resources of the YARN service increase. For more information about how to modify the parameters of a service, see Manage parameters for services.

You must modify the following parameters:

- Memory: yarn.nodemanager.resource.memory-mb
- vCPU: yarn.nodemanager.resource.cpu-vcores

	emr-	< Back 💯 YARN 🗸 🌒 Normal		() History
-	Cluster Overview	Status Component Deployment Configure History		
a	Cluster Management			
6	Cluster Service 🔷	Configuration Filter	Service Configuration	
	🖉 HDFS	Search yam.nodemanager.resource. O Q	ALL   yam-site	
	🌮 YARN	Scope	yam.nodemanager.resource.cpu-vcores	0
	🕷 Hive	Default Cluster Co 🕆	yam.nodemanager.resource.memory-mb 26144	0

# What do I do if resources for ApplicationMasters of a job that runs on YARN exceed the upper limit of a YARN queue?

- Problem description:
  - The job that runs on YARN remains in the ACCEPTED state.
  - The error Application is added to the scheduler and is not yet activated. Queue's AM resour ce limit exceeded. is reported.
- Solution: When you run a job on YARN, an ApplicationMaster starts and applies for the required resources. However, the yarn.scheduler.capacity.maximum-am-resource-percent parameter is used to limit the percentage of YARN resources that an ApplicationMaster can use. By default, the value of this parameter is 0.25, which indicates 25%. If your jobs do not consume a large number of resources, we recommend that you set the yarn.scheduler.capacity.maximum-am-resource-percent parameter to a value that ranges from 0.5 to 0.8. After you modify the parameter, you do not need to restart the YARN service. You can directly run the yarn rmadmin -refreshQueues command on a node. For more information about how to modify the parameters of a service, see Manage parameters for services.

	=	Home Page > Cluster Management > Cluster (C-E9	> Service > YARN		
	emr	< Back 😥 YARN 🗸 🌑 Normal		•	O History
≡	Cluster Overview	Status Component Deployment Configure History			
馬	Cluster Management				
6	Cluster Service	Configuration Filter	Service Configuration		
	HDFS	Search	ALL capacity-scheduler		
	🖉 HDFS	yarn.scheduler.capacity.maximum-am-resou 💿 🔍			
	SP YARN	Scope	xml-direct-to-file-content	<configuration> <property></property></configuration>	
	🐔 Hive	Default Cluster Co 💙		<pre><name>yarn.scheduler.capacity.maximum- applications</name> <value>10000</value></pre>	
	5 <sup>5</sup> Ganglia	Туре		<description>Maximum number of applications that can be pending and running. </description>	
	🖌 ZooKeeper	BASIC ADVANCED INFORMATION		 <property> <name>vam.scheduler.capacity.maximum-</name></property>	
	🞝 Spark	DATA_PATH LOG_PATH LOG JVM		am-resource-percent am-resource-percen	s
	🍋 HBase	CODEC OSS PORT MEMORY DISK		application masters i.e. controls number of concurrent running applications. 	
	<b>H</b> Hue	NETWORK PATH URI		<property> <name>yarn.scheduler.capacity.resource-</name></property>	•

# What are the differences between the two job submission modes Worker Node and Header/Gateway Node?

On the Edit Job page, click Job Settings in the upper-right corner. In the Job Settings panel, click the Advanced Settings tab. In the Mode section, select Worker Node or Header/Gateway Node for Job Submission.

**Worker Node**: In this mode, you cannot specify a node to submit the job. Before your actual task starts and is monitored, a task named Launcher starts. Therefore, resources for two ApplicationMasters are required. In most cases, the application IDs of the two tasks are continuous.

**Header/Gateway Node:** In this mode, you can specify a node to submit the job. Only the actual tasks start. However, if the number of tasks that you submit is excessively large, the node may be overloaded. If the memory of the node is insufficient, some tasks cannot start.

# How do I allow ECS instances in the classic network and ECS instances of EMR clusters in a virtual private cloud (VPC) to access each other?

Alibaba Cloud provides two network types: classic network and VPC. EMR clusters run in VPCs, but the business systems of many users run in the classic network. To connect these two types of networks, Alibaba Cloud launched the ClassicLink solution. For more information, see ClassicLink overview.

To implement this solution, perform the following steps:

- 1. Create a vSwitch for which a specific CIDR is configured. For more information, see ClassicLink overview.
- 2. Create an EMR cluster. Select the created vSwitch when you create the cluster.
- 3. In the ECS console, connect the required instances of the classic network to the VPC where the EMR cluster resides.
- 4. Configure security group rules.

### How do I access the web UIs of open source services?

You can access the web UIs of open source services on the **Public Connect Strings** page in the EMR console. For more information, see Access the web UIs of open source components.

# How do I isolate the Object Storage Service (OSS) data of different Resource Access Management (RAM) users?

To isolate the OSS data of different RAM users, perform the following steps in the RAM console:

- 1. Log on to the RAM console by using your Alibaba Cloud account.
- 2. Create a RAM user.
  - i. In the left-side navigation pane, choose Identities > Users.
  - ii. On the Users page, click **Create User**.

**?** Note You can create multiple RAM users at the same time.

- iii. On the Create User page, specify Logon Name and Display Name.
- iv. In the Access Mode section, select Console Access or OpenAPI Access.
  - Console Access: If you select this access mode, you must configure the basic settings for logon security. These settings specify whether to use a system-generated logon password or a custom logon password, whether to reset the password on the next logon, and whether to enable multi-factor authentication (MFA).
  - **OpenAPI Access:** If you select this access mode, an AccessKey pair is automatically created for the RAM user. The RAM user can call API operations or use other development tools to access Alibaba Cloud resources.

(?) Note To ensure the security of your Alibaba Cloud account, we recommend that you select only one access mode for the RAM user. This way, the RAM user cannot use an AccessKey pair to access Alibaba Cloud resources after the RAM user leaves the organization.

- v. Click OK.
- 3. Create permission policies.
  - i. In the left-side navigation pane, choose **Permissions > Policies**.
  - ii. Click Create Policy.
  - iii. On the Create Policy page, specify Policy Name.
  - iv. Select **Script** for Configuration Mode. For more information about how to configure a permission policy in **Script** mode, see Policy structure and syntax. In this example, two policies are created based on different environments. You can select one of the scripts to create a policy based on your business environment.

Test environment (test-bucket)	Production environment (prod-bucket)
<pre>{   "Version": "1",   "Statement": [   {   "Effect": "Allow",   "Action": [     "oss:ListBuckets"  ],   "Resource": [     "acs:oss:*:*:"  ]  },  {   "Effect": "Allow",   "Action": [     "oss:Listobjects",     "oss:PutObject",     "oss:PutObject",     "oss:DeleteObject"  ],  "Resource": [     "acs:oss:*:*:test-bucket/*"  ]  } ] }</pre>	<pre>{ "Version": "1", "Statement": [ { "Effect": "Allow", "Action": [ "acs:oss:*:*:" ] }, { "Effect": "Allow", "Action": [ "oss:Listobjects", "oss:GetObject", "oss:PutObject" ], "Resource": [ "acs:oss:*:*:prod-bucket/*" ] } ] }</pre>

After the preceding permission policies are attached to a RAM user, the RAM user is subject to the following limits in the EMR console:

- When the RAM user creates a cluster, job, or workflow, all buckets are displayed on the OSS file page. However, the RAM user can access only authorized buckets.
- Only the data in authorized buckets can be viewed.
- A job can read and write data only from and to an authorized bucket.

v. Click OK.

4. (Optional)Attach policies to the RAM user.

If the created RAM user is not attached policies, perform the following steps to attach policies to the RAM user:

- i. In the left-side navigation pane, choose Identities > Users.
- ii. On the Users page, find the RAM user to which you want to attach policies and click Add **Permissions** in the Actions column.
- iii. In the Add Permissions panel, click the policies that you want to attach to the RAM user and click OK.
- iv. Click Complete.
- 5. (Optional)Grant console logon permissions to the RAM user.

If you have not granted console logon permissions to the RAM user that you created, perform the following steps to grant the permissions to the RAM user:

- i. In the left-side navigation pane, choose **Identities > Users**.
- ii. On the Users page, find the RAM user to which you want to grant console logon permissions and click the logon name of the RAM user.
- iii. In the **Console Logon Management** section of the page that appears, click **Modify Logon** Settings.
- iv. In the Modify Logon Settings panel, select Enabled for Console Password Logon.
- v. Click OK.
- 6. Log on to the EMR console by using the RAM user.
  - i. Log on to the Alibaba Cloud Management Console by using the RAM user.
  - ii. Click the More icon in the upper-left corner and choose E-MapReduce.

# How do I connect EMR clusters that belong to different VPCs of the same account?

Each VPC is an isolated network environment. You can customize CIDR blocks, create subnets, and configure route tables and gateways for VPCs. You can create EMR clusters in different VPCs and use Express Connect to allow the VPCs to communicate with each other.

You must configure the following VPC-related parameters when you create an EMR cluster in a VPC:

- VPC: the VPC to which the cluster belongs. If no VPCs are available, create a VPC in the VPC console. You can create a maximum of two VPCs by using one account. To create more than two VPCs, .
- VSwitch: the vSwitch to which the cluster belongs. The vSwitch is used to support the communication of the ECS instances of the cluster. If no vSwitches are available, log on to the VPC console and click **vSwitch** in the left-side navigation pane to create a vSwitch. vSwitches are deployed in zones. When you create a vSwitch, you must select the zone in which you want to create a cluster.
- Security Group Name: the security group to which the cluster belongs. Security groups in the classic network cannot be used. Only security groups in the specified VPC can be used. For security purposes, only the security groups created in EMR are available in the drop-down list. To create a security group, you need only to enter a name in the Security Group Name field.

The following example describes how to create a Hive cluster and an HBase cluster that belong to different VPCs and use Cloud Enterprise Network to allow the Hive cluster to access the HBase cluster.

1. For more information about how to create a cluster, see Create a cluster.

In the EMR console, create Hive cluster C1 and HBase cluster C2 in the China (Hangzhou) region. C1 belongs to VPC 1 and C2 belongs to VPC 2.

2. Connect the two VPCs.

For more information, see Create a CEN instance.

3. Log on to the HBase cluster in SSH mode and run the following command in the HBase shell to create a table:

create 'testfromHbase','cf'

- 4. Log on to the Hive cluster in SSH mode and perform the following operations to enable communication between the clusters:
  - i. Add the following information to the hosts file:

 $zk_ip \ emr-cluster \ // \ zk_ip \ specifies the IP address of the ZooKeeper node in the HB ase cluster.$ 

ii. Run the following commands in the Hive shell to access the HBase cluster:

```
set hbase.zookeeper.quorum=172.*.*.111,172.*.*.112,172.*.*.113;
CREATE EXTERNAL TABLE IF NOT EXISTS testfromHive (rowkey STRING, pageviews Int, bytes
STRING) STORED BY 'org.apache.hadoop.hive.hbase.HBaseStorageHandler' WITH SERDEPROPER
TIES ('hbase.columns.mapping' = ':key,cf:cl,cf:c2') TBLPROPERTIES ('hbase.table.name'
= 'testfromHbase');
```

If the java.net.SocketTimeoutException error appears, add rules to the security group of the HBase cluster to enable all ports required for the Hive cluster to access the HBase cluster. The following figure shows an example.

Protocol type	Port range	Authorization type	Authorization object
ТСР	2181/2181	Address field access	192.168
ТСР	22/22	Address field access	and all the second seco
ТСР	16000/16000	Address field access	192.168.
ТСР	16020/16020	Address field access	192.168.
	тср тср тср	TCP         2181/2181           TCP         22/22           TCP         16000/16000	TCP     2181/2181     Address field access       TCP     22/22     Address field access       TCP     16000/16000     Address field access

# What do I do if the "Failed to get schema version" error is reported when the MetaStore client is initialized?

The following figure shows the error details.

```
[root@emr-header-1 ~J# su hadoop
[hadoop@emr-header-1 root]$ schematool -initSchema -dbTvpe mvsal
Metastore connection URL: jdbc:mysql://
Metastore Connection Driver : com.mysql.jdbc.Driver
Metastore connection User:
org.apache.hadoop.hive.metastore.HiveMetaException: Failed to get schema version.
Underlying cause: com.mysql.jdbc.exceptions.jdbc4.CommunicationsException : Communications link failure
The last packet sent successfully to the server was 0 milliseconds ago. The driver has not received any packets from the
server.
SQL Error code: 0
Use --verbose for detailed stacktrace.
*** schemaTool failed ***
```

Check the security group settings of your ApsaraDB RDS for MySQL instance. Make sure that the IP addresses of the ECS instances that are used for your EMR cluster are added to the whitelist of the ApsaraDB RDS for MySQL instance. For more information, see Configure an IP address whitelist for an ApsaraDB RDS for MySQL instance.

# What do I do if Hive metadata contains Chinese characters, such as Chinese characters in column comments and partition names?

Perform the following operations in the ApsaraDB RDS for MySQL database to encode the related fields in the UTF-8 format:

1. Change the data type of the COMMENT column:

alter table COLUMNS\_V2 modify column COMMENT varchar(256) character set utf8;

2. Change the data type of the PARAM\_VALUE column in the TABLE\_PARAMS table:

alter table TABLE\_PARAMS modify column PARAM\_VALUE varchar(4000) character set utf8;

3. Change the data type of the PARAM\_VALUE column in the PARTITION\_PARAMS table:

alter table PARTITION\_PARAMS modify column PARAM\_VALUE varchar(4000) character set utf8;

4. Change the data type of the PKEY\_COMMENT column:

alter table PARTITION\_KEYS modify column PKEY\_COMMENT varchar(4000) character set utf8;

5. Change the data type of the PARAM\_VALUE column in the INDEX\_PARAMS table:

alter table INDEX\_PARAMS modify column PARAM\_VALUE varchar(4000) character set utf8;