Alibaba Cloud E-MapReduce

Cluster Planning and Configurations

Issue: 20200709

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 Cloud-authorized 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 due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall 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 the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. 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 bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary, incidental, 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 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 contact Alibaba Cloud directly if you discover any errors in this document.

Document conventions

Style	Description	Example
	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.
	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.
	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.
	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 list --instanceid` *Instance_ID*
[] or [a|b]	This format is used for an optional value, where only one item can be selected.	`ipconfig [-all|-t]`

Style	Description	Example
{} or {a|b}	This format is used for a required value, where only one item can be selected.	switch {active|stand}

Contents

1 Cluster planning

1.1 Authorize roles

When you use E-MapReduce (EMR) for the first time, you must assign the default system roles **AliyunEMRDefaultRole** and **AliyunEmrEcsDefaultRole** to EMR by using your Alibaba Cloud account.

Background information

For more information about the roles, see #unique_5.

- After the **AliyunEMRDefaultRole** role is assigned, you can use EMR to access other services such as ECS and OSS, create clusters, and store logs. For more information about the permissions of this role, see AliyunEMRDefaultRole.
- After the **AliyunEmrEcsDefaultRole** role is assigned, your EMR clusters can access Alibaba Cloud resources (such as OSS) without an AccessKey pair. For more information, see MetaService. After the role authorization is complete, you can use a default role or a custom role to create a cluster. For more information about the permissions of the AliyunEmrEcsDefaultRole role, see AliyunEMREcsDefaultRole.

> **Notice:**
> When you use EMR for the first time, you must assign default system roles to EMR by using your Alibaba Cloud account. Otherwise, your Alibaba Cloud account and RAM users under this account cannot use EMR.

Procedure

1. 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.

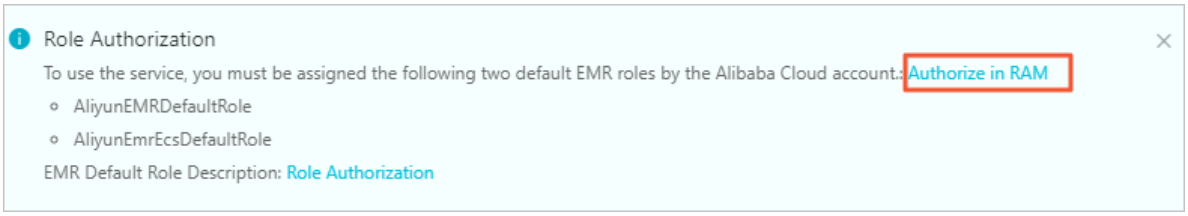

2. Click **Authorize in RAM**. On the page that appears, click **Confirm Authorization Policy** to assign the default roles **AliyunEMRDefaultRole** and **AliyunEmrEcsDefaultRole** to EMR.

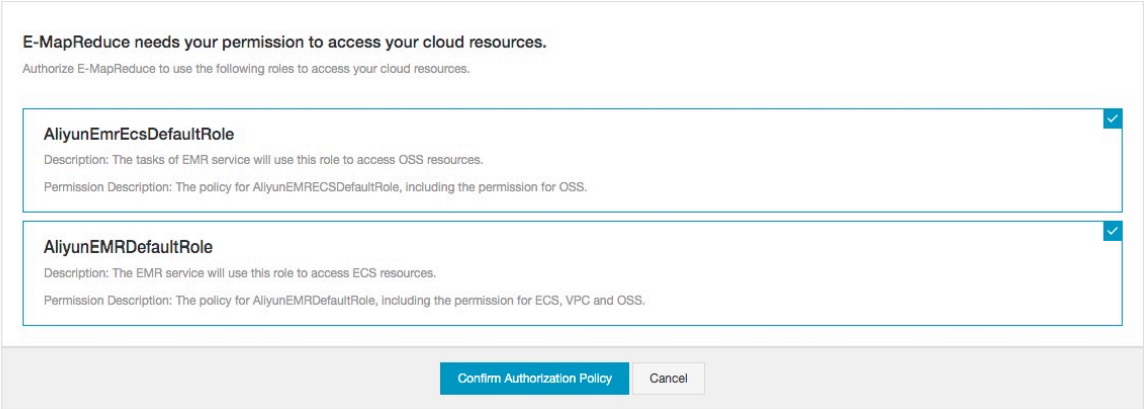

3. Refresh the EMR console to use the services in the console. To view policy details for the **AliyunEMRDefaultRole** and **AliyunEmrEcsDefaultRole** roles, log on to the RAM console.

Permissions of the default roles

- **AliyunEMRDefaultRole**

  AliyunEMRDefaultRole contains AliyunEMRRolePolicy. The following tables describe the permissions of this role.

  - ECS-related permissions

Permission (Action)	Description
ecs:CreateInstance	Creates an ECS instance.
ecs:RenewInstance	Renews an ECS instance.
ecs:DescribeRegions	Queries the region information of an ECS instance.
ecs:DescribeZones	Queries the zone information of an ECS instance.
ecs:DescribeImages	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.

Permission (Action)	Description
ecs:AuthorizeSecurityGroup	Sets inbound rules for a security group.
ecs:AuthorizeSecurityGroupEgress	Sets outbound rules for a security group.
ecs:DescribeSecurityGroupAttribute	Views 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.

- **AliyunEmrEcsDefaultRole**

  AliyunEmrEcsDefaultRole contains AliyunEMRRolePolicy. The following table describes OSS-related permissions of this role.

Permission (Action)	Description
oss:PutObject	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.

Note:

You can attach policies to the role based on your business requirements.

1.2 Grant permissions to a RAM user

To make sure that your Resource Access Management (RAM) users can use the E-MapReduce console, you must access the RAM console by using your Alibaba Cloud account and grant the required permissions to the RAM users. This topic describes how to

use RAM to control access to E-MapReduce cluster resources at the account level. Detailed operations include creating RAM users or user groups and granting specific permissions.

Context

RAM is a resource access control service provided by Alibaba Cloud. For more information, see #unique_8 The following examples describe how RAM is used to implement access control.

- RAM users: You can purchase multiple Elastic Compute Service (ECS) instances to build E-MapReduce clusters. Your organization may have multiple users who need to access the instances, for example, users who are responsible for O&M, development, and data analysis. You can create multiple RAM users and use policies to control the permissions of these users. This eliminates the risk of disclosing your AccessKey pair of your Alibaba Cloud account and helps maintain 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 RAM users in batches.

Policy

Policies are categorized to system policies and custom policies.

- System policies: the default policies provided by Alibaba Cloud. The following system policies are frequently used in E-MapReduce:
  - AliyunEMRFullAccess: provides RAM users with full access to E-MapReduce, including all permissions on all E-MapReduce resources.
  - AliyunEMRDevelopAccess: provides RAM users with the developer permissions of E-MapReduce. This policy does not authorize RAM users to perform some operations such as creating and releasing clusters.
  - AliyunEMRFlowAdmin: provides RAM users with the administrator permissions of E-MapReduce. This policy allows RAM users to create projects, and develop and manage jobs. It does not allow RAM users to add members to projects or manage clusters.
- Custom policies: the user-defined policies. These policies are suitable for users who are familiar with various Alibaba Cloud service APIs and require fine-grained access control. For more information, see #unique_9.

Procedure

To grant permissions on E-MapReduce to RAM users in the RAM console, follow these steps:

1. Log on to the RAM console by using an Alibaba Cloud account.
2. In the left-side navigation pane, choose **Identities** > **Users**.
3. On the Users page, find the target RAM user to which you want to grant permissions and click **Add Permissions** in the Actions column.
4. In the Add Permissions pane, click the policy you want to attach to the RAM user and click **OK**.

   For more information about policies, see Policy.
5. Click **Complete**.

   After authorization is completed, the granted permissions take effect immediately. The authorized RAM users can log on to the RAM console to check their permissions.

1.3 Manage users

This topic describes how to create Resource Access Management (RAM) users and manage service accounts for RAM users. Currently, Knox and Kerberos accounts are supported.

Prerequisites

To allow an E-MapReduce cluster to support Kerberos accounts, you must enable the Kerberos cluster mode when you create the cluster. Alternatively, you can click **Enable Kerberos Mode** on the **Cluster Management** page to enable the high-security mode after the cluster is created.

After the high-security mode is enabled, all components in the cluster use Kerberos for authentication. For more information, see #unique_11.

Create a RAM user

You can use RAM to create and manage user accounts, which include employee accounts, system accounts, and application accounts. You can also manage the operation permissions that these user accounts have on the resources of your account. To create a RAM user, follow these steps:

1. Log on to the Alibaba Cloud E-MapReduce console.
2. In the top navigation bar, select the region where your cluster resides. Select the resource group as required. By default, all resources of the account appear.

3. Click the **Cluster Management** tab.
4. On the **Cluster Management** page that appears, find the target cluster and click **Details** in the Actions column.
5. In the left-side navigation pane of the Cluster Overview page that appears, click **Users**.
6. On the Users page that appears, click **Create RAM User** in the upper-right corner to access the Overview page of the RAM console.

   You can use RAM to create users and groups, and add permissions. For more information, see the RAM documentation.

Activate a Knox account

After you create a RAM user, you can follow these steps to activate a Knox account for the RAM user:

1. On the **Users** page, find the target the RAM user and click **Set Knox Password** in the **Action** column.
2. In the **Set Password** dialog box that appears, set **Password** and **Confirm password**.
3. Click **OK**.
4. Click Refresh in the upper-right corner to refresh the **Users** page and view the status of the Knox account.

   If **Synchronized** appears in the **Knox Account** column, the Knox account is activated. Then you can use the **username** and **password** to log on to Knox.

   For information, see #unique_12.

Delete a Knox account

1. On the **Users** page, find the target RAM user and click **Delete Knox Account** in the **Action** column.
2. Click **OK**.
3. Click Refresh in the upper-right corner to refresh the **Users** page and check whether the Knox account is deleted.

   If **Not Synchronized** appears in the Knox Account column, the Knox account is deleted.

Activate a Kerberos account

After you create a RAM user, you can follow these steps to activate a Kerberos account for the RAM user:

1. On the **Users** page, find the target RAM user and click **Set Kerberos Password** in the **Action** column.
2. In the **Set Password** dialog box that appears, set **Password** and **Confirm password**.
3. Click **OK**.
4. Click Refresh in the upper-right corner to refresh the **Users** page and view the status of the Kerberos account.

   If **Synchronized** appears in the **Kerberos Account** column, the Kerberos account is activated.

Delete a Kerberos account

1. On the **Users** page, find the target RAM user and click **Delete Kerberos Account** in the **Action** column.
2. Click **OK**.
3. Click Refresh in the upper-right corner to refresh the **Users** page and check whether the Kerberos account is deleted.

   If **Not Synchronized** appears in the Kerberos Account column, the Kerberos account is deleted.

Note

Different clusters cannot share a Knox account. This is because Knox accounts are created in clusters separately. For example, Knox account A created in cluster-1 cannot be shared with cluster-2. To use Knox account A in cluster-2, you must create this account in cluster-2.

1.4 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 only performs 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:**
> From EMR V3.2.0, clusters support task nodes.

Master nodes

Master nodes are the nodes deployed with components that are used to manage cluster service deployment, such as ResourceManager of Hadoop YARN. To view the running status of services in a cluster, you can use SSH to connect to the master node and then access the Web UIs of components. To quickly test or run a job, you can also connect to the master node and submit the job on the command line.

A cluster with high availability (HA) enabled has two master nodes. A cluster with HA disabled has only one master node. HA is disabled by default.

Core nodes

Core nodes are managed by master nodes. Core nodes run the DataNode service 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 nodes

Task nodes only run 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 effect on the running of the cluster. However, the reduction of task nodes may cause MapReduce and Spark jobs to fail. Whether the jobs are affected depends on the retry and fault tolerance capabilities of the computing service.

1.5 Deployment of service components

This topic uses EMR-3.26.0 as an example to describe the deployment details of service components on each node in different EMR clusters.

Background information

When you create an EMR cluster, different service components are deployed on different categories of nodes. For example, the NameNode component of Hadoop HDFS is deployed on the master node. For more information about node categories, see Node categories. You can use bootstrap actions to complete custom deployment and allocation for large-scale clusters. For example, a stand-alone ZooKeeper cluster uses Alibaba Cloud ApsaraDB for RDS as the Hive metastore.

View deployment details of service components on a cluster

1. Log on to the Alibaba Cloud E-MapReduce console.
2.
3. Click the **Cluster Management** tab.
4. Find the target cluster and click **Details** in the Actions column.
5. In the left-side navigation pane of the Cluster Overview page that appears, click **Cluster Service** and choose the service for which you want to view deployment details of components.
6. On the service page that appears, click the **Component Deployment** tab.

Hadoop clusters

Deployment details of service components that run on a Hadoop cluster

- Required services

Service name	Master node	Core node
HDFS	- KMS - SecondaryNameNode - HttpFS - HDFS Client - NameNode	- DataNode - HDFS Client

Service name	Master node	Core node
YARN	- ResourceManager - App Timeline Server - JobHistory - WebAppProxyServer - Yarn Client	- Yarn Client - NodeManager
Hive	- Hive MetaStore - HiveServer2 - Hive Client	Hive Client
Spark	- Spark Client - SparkHistory - ThriftServer	Spark Client
Knox	Knox	N/A
Zeppelin	Zeppelin	N/A
Tez	- Tomcat - Tez Client	Tez Client
Ganglia	- Gmond - Httpd - Gmetad - Ganglia Client	- Gmond - Ganglia Client
Pig	Pig Client	Pig Client
Sqoop	Sqoop Client	Sqoop Client
Bigboot	- Bigboot Client - Bigboot Monitor	- Bigboot Client - Bigboot Monitor
OpenLDAP	OpenLDAP	N/A
Hue	Hue	N/A
SmartData	- Jindo Namespace Service - Jindo Storage Service - Jindo Client	- Jindo Storage Service - Jindo Client

- Optional services

Service name	Master node	Core node
Flume	- Flume Agent - Flume Client	- Flume Agent - Flume Client
Livy	Livy	N/A
Superset	Superset	N/A
Flink	- FlinkHistoryServer - Flink Client	Flink Client
Ranger	- RangerPlugin - RangerAdmin - RangerUserSync	RangerPlugin
Storm	- Storm Client - UI - Nimbus - Logviewer	- Storm Client - Supervisor - Logviewer
Phoenix	Phoenix Client	Phoenix Client
Analytics Zoo	- Analytics-Zoo-Scala - Analytics-Zoo-Python	- Analytics-Zoo-Scala - Analytics-Zoo-Python
Kudu	- Kudu Master - Kudu Client	- Kudu Tserver - Kudu Master - Kudu Client
HBase	- HMaster - HBase Client - ThriftServer	- HBase Client - HRegionServer
ZooKeeper	- ZooKeeper - ZooKeeper Client	- ZooKeeper leader - ZooKeeper follower - ZooKeeper Client
Oozie	Oozie	N/A
Presto	- Presto Client - PrestoMaster	- Presto Client - PrestoWorker

Service name	Master node	Core node
Impala	- Impala Runtime and Shell - Impala Catalog Server - Impala StateStore Server	- Impala Runtime and Shell - Impala Daemon Server

Druid clusters

Deployment details of service components on a Druid cluster

- Required services

Service name	Master node	Core node
Druid	- Druid Client - Coordinator - Overlord - Broker - Router	- MiddleManager - Historical - Druid Client
HDFS	- KMS - SecondaryNameNode - HttpFS - HDFS Client - NameNode	- DataNode - HDFS Client
Ganglia	- Gmond - Httpd - Gmetad - Ganglia Client	- Gmond - Ganglia Client
ZooKeeper	- ZooKeeper follower - ZooKeeper Client	- ZooKeeper leader - ZooKeeper follower - ZooKeeper Client
OpenLDAP	OpenLDAP	N/A
Bigboot	- Bigboot Client - Bigboot Monitor	- Bigboot Client - Bigboot Monitor
SmartData	- Jindo Namespace Service - Jindo Storage Service - Jindo Client	- Jindo Storage Service - Jindo Client

- Optional services

Service name	Master node	Core node
YARN	- ResourceManager - App Timeline Server - JobHistory - WebAppProxyServer - Yarn Client	- Yarn Client - NodeManager
Superset	Superset	N/A

Kafka clusters

Deployment details of service components on a Kafka cluster

- Required services

Service name	Master node	Core node
Kafka-Manager	Kafka Manager	N/A
Kafka	- Kafka Client - KafkaMetadataMonitor - Kafka Rest Proxy - Kafka Broker broker - Kafka Schema Registry	- Kafka Broker broker - Kafka Client
Ganglia	- Gmond - Httpd - Gmetad - Ganglia Client	- Gmond - Ganglia Client
ZooKeeper	- ZooKeeper follower - ZooKeeper Client	- ZooKeeper leader - ZooKeeper follower - ZooKeeper Client
OpenLDAP	OpenLDAP	N/A

- Optional services

Service name	Master node	Core node
Ranger	- RangerPlugin - RangerUserSync - RangerAdmin	RangerPlugin
Knox	Knox	N/A

ZooKeeper clusters

Deployment details of service components on a ZooKeeper cluster

Service name	Master node	Core node
Ganglia	N/A	• Gmond • Httpd • Gmetad • Ganglia Client
ZooKeeper	N/A	• ZooKeeper leader • ZooKeeper Client • ZooKeeper follower

Flink clusters

Deployment details of service components on a Flink cluster

Required services

Service name	Master node	Core node
HDFS	• KMS • SecondaryNameNode • HttpFS • HDFS Client • NameNode	• DataNode • HDFS Client
YARN	• ResourceManager • App Timeline Server • JobHistory • WebAppProxyServer • Yarn Client	• Yarn Client • NodeManager
Ganglia	• Gmond • Httpd • Gmetad • Ganglia Client	• Gmond • Ganglia Client
ZooKeeper	• ZooKeeper follower • ZooKeeper Client	• ZooKeeper leader • ZooKeeper follower • ZooKeeper Client
Zeppelin	Zeppelin	N/A

Service name	Master node	Core node
Flink	• FlinkHistoryServer • Flink Client	Flink Client
HAS	HASServer	N/A
SmartData	• JindoFS Namespace Service • JindoFS Storage Service • SmartData Client	• JindoFS Storage Service • SmartData Client
Bigboot	• Bigboot Client • Bigboot Monitor	• Bigboot Client • Bigboot Monitor
OpenLDAP	OpenLDAP	N/A
Grafana	GrafanaServer	N/A
Prometheus	PrometheusServer	N/A
Alertmanager	AlertManager	N/A

1.6 Gateway clusters

Gateway clusters can serve as a separate job submission node for Hadoop clusters. When you create a gateway cluster, you must associate it with an existing cluster, so you can better run the associated cluster.

A gateway cluster is an independent cluster that consists of multiple nodes of 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 or 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 business requirements.

1.7 ECS instances

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

   **Note:**
  This instance family only applies to core nodes. Core nodes can be created only in Hadoop, Flink, and Druid clusters.
- Local SSD type

  This instance family uses local SSDs for storage, which provides high local IOPS and 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, not enterprise customers.
- GPU

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

Application 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 where data is stored on cloud disks provided by Alibaba Cloud. There are three replicas of data 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.
  - If local disks are used, the data reliability cannot be ensured because data is maintained on the EMR platform.
- Task nodes

  All instance families except the big data type apply to task nodes to supplement the computing capabilities of a cluster. The local SSD type is in development.

1.8 Storage

This topic describes data storage in E-MapReduce (EMR) clusters.

Disk roles

There are two types of disks for a node in an EMR cluster: the system disk that is used to install the operating system, and data disks that are used to store data. One system disk is deployed by default, while many data disks can be deployed. A maximum of 16 data disks can be attached to each node. The data disks can vary in configurations, types, and capacities.

By default, Enhanced SSD (ESSD) is used as the system disk, one cloud disk is used as the data disk for the master node, and four cloud disks are used as the data disks for each core node.

Cloud disks and local disks

The following two types of disks are available for data storage:

- Cloud disks

  Include standard SSDs, ultra disks, and ESSDs. Disks are not attached directly to local compute nodes. Instead, they 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, a replica is used automatically for recovery.

- Local disks

  Include local SATA disks and local SSDs of the big data type. Local disks are attached directly to compute nodes and have better performance than cloud disks. You cannot change the number of local disks. There is also no data backup at the backend. Upper-layer software is required to ensure data reliability.

When nodes in an EMR cluster are released, data on all 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 also uses all data disks as temporary storage for computing.

If your business volume is below terabytes, you can use cloud disks that provide lower IOPS and throughput. If you have terabytes of data, we recommend that you use local disks. The data reliability of local disks is ensured by EMR. If the throughput of cloud disks is insufficient, use local disks instead.

OSS

OSS can be used as HDFS in an EMR cluster. You can easily read and write data from and to OSS simply by modifying code that was originally used by HDFS. Example:

- Read data from Spark:

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

  Change the storage type from HDFS to OSS:

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

- You can directly run HDFS commands in OSS for MapReduce or Hive jobs.

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

  In this process, you do not need to enter the AccessKey pair or endpoint of OSS. EMR automatically completes this 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.

1.9 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. The cost is almost the same as physical hosts.

However, local disks cannot ensure data reliability. Alibaba Cloud provides a multi-replicas data storage policy for cloud disks to ensure data reliability, and you do not have 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.

The highlights of automated O&M solutions are as follows:

- 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.

Note:

If you want to create a Hadoop cluster by using d1 series instances, submit a ticket and contact EMR engineers for assistance. For information about how to submit a ticket, Click New Ticket.

1.10 Inter-access between classic networks and VPCs

This topic describes how to allow ECS instances in classic networks and instances of E-MapReduce (EMR) clusters in VPCs to access each other.

Alibaba Cloud provides two network types: classic network and VPC. EMR clusters run in VPCs, while business systems of many users still run in classic networks. To connect these two types of networks, Alibaba Cloud launched the ClassicLink solution. For more information, see ClassicLink overview.

The procedure to implement this solution is as follows:

1. Create a VSwitch based on the CIDR blocks described in the "ClassicLink overview" topic.
2. Use the VSwitch to deploy an EMR cluster.
3. In the ECS console, connect the required instances of a classic network to the VPC.
4. Set security group rules.

1.11 Disaster recovery in E-MapReduce clusters

This article will introduce disaster recovery of data and dervices in E-MapReduce clusters

Data

HDFS stores the data of each file in blocks, with each block holding multiple copies (three by default). HDFS also makes sure that these copies are stored in different frameworks. In most situations, HDFS stores the first copy in the local framework, the second in the same framework as the first but in different nodes, and the last copy in a different framework.

HDFS scans the data copies regularly. If it finds that a data copy has been lost, HDFS makes another to make sure the number of copies is stable. If a node that stores a copy has been lost, HDFS makes another node to recover the data in that node. In Alibaba Cloud, if you use cloud disks, each cloud disk has three data copies in the back-end. If any of them has an issue, the copies exchange and recover data to ensure reliability.

HDFS is a highly reliable file storage system that can store massive amounts of data. Based on the features of Alibaba Cloud, HDFS can also make backups of the data stored in OSS, providing even greater data reliability.

Services

The core components of HDFS guarantee high availability by making sure that there are at least two nodes to back each other up, such as YARN, HDFS, HiveServer, or Hive Meta. In this way, whenever a node experiences an issue, the nodes can exchange and recover data to ensure that services are not impacted.

2 Configure clusters

2.1 Create a cluster

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

Prerequisites

Authorization is completed in Resource Access Management (RAM). For more information, see Role authorization.

Procedure

1. Go to the cluster creation page.
   a) Log on to the Alibaba Cloud E-MapReduce console.
   b) 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.
   c) Click **Cluster Wizard** in the Clusters section.
2. Configure the cluster.

   To create a cluster, configure software, hardware, 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 correct when you create a cluster.

   a) Configure software parameters.

Parameter	Description
EMR Version	The major version of E-MapReduce. The latest version is selected by default.

Parameter	Description
Cluster Type	The type of the cluster you want to create. E-MapReduce supports the following types of clusters: • **Hadoop**: Hadoop clusters provide multiple ecosystem components, such as Hadoop, Hive, Spark, Spark Streaming, Flink, Storm, Presto, Impala, Oozie, and Pig. Hadoop, Hive, and Spark are semi-hosted services and are used to store and compute large-scale distributed data offline. Spark Streaming, Flink, and Storm provide stream computing. Presto and Impala are used for interactive queries. For information about these components, see the Services section of the Status tab on the Clusters and Services page. • **Kafka**: Kafka clusters serve as a semi-hosted, distributed message system with high throughput and scalability. Kafka clusters provide a comprehensive service monitoring system that maintains cluster stability. Kafka clusters are professional, reliable, and secure. You do not need to deploy or maintain these clusters. These clusters are used in scenarios such as log collection and monitoring data aggregation. They can also be used for offline data processing, stream computing, and real-time data analysis. • **ZooKeeper**: ZooKeeper clusters provide a distributed and consistent lock service that facilitates coordination among large-scale Hadoop, HBase, and Kafka clusters. • **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 E-MapReduce Hadoop, E-MapReduce Spark, Object Storage Service (OSS), and ApsaraDB for RDS to build a flexible and stable system for real-time queries. • **Flink**: Flink clusters support all features of the open-source Flink ecosystem and can be used with Alibaba Cloud services such as OSS.
Required Services	The default components required for a specific cluster type. After a cluster is created, you can add, start, or stop services on the cluster management page.

Parameter	Description
Optional Services	The other components that you can specify as required. The relevant service processes for the components you specify are started by default. **Note:** 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.

Parameter	Description
Advanced Settings	• **Kerberos Mode**: specifies whether to enable Kerberos authentication for clusters. This feature is disabled by default. This is because it is not required by clusters created for common users. • **Custom Software Settings**: customizes software settings. You can use a JSON file to customize the parameters of basic components required for a cluster, such as Hadoop, Spark, and Hive. For more information, see Software configuration. This feature is disabled by default.

b) Configure hardware parameters.

Section	Parameter	Description
Billing Method	**Billing Method**	The billing method. Valid values: • **Pay-As-You-Go**: The system charges you for a cluster based on the hours the cluster is actually used. The fee is deducted on an hourly basis. The unit price of a pay-as-you-go cluster is higher than that of a subscription cluster of the same specification. We recommend that you use pay-as-you-go clusters for short-term test jobs or dynamically scheduled jobs. • **Subscription**: The system charges you for a cluster only once per subscription period. The unit price of a subscription cluster is lower than that of a pay-as-you-go cluster of the same specification. Subscription clusters are offered discounts, and longer subscription periods offer larger discounts. **Subscription Period**: the validity period of your subscription. You can get a 15% discount if you select a 12-month subscription period. **Note:** We recommend that you create a **pay-as-you-go** cluster for a test run. If the cluster passes the test, you can create a **subscription** cluster for production.

Section	Parameter	Description
Network Settings	**Zone**	The zone where a cluster is created. Zones are different geographical areas located in the same region. They are interconnected by an internal network. In most cases, you can use the zone selected by default.
	Network Type	The network type of the cluster. The Virtual Private Cloud (VPC) network type is selected by default. If no VPC is created, go to the VPC console to create one.
	VPC	The VPC selected in the region where your cluster resides. If no VPC is available in the region, click **Create VPC/VSwitch** to create one.
	VSwitch	The VSwitch selected in the zone where your cluster and VPC reside. If no VSwitch is available in the zone, create one.
	Security Group Name	The security group to which you want to add your cluster. If this is your first time using E-MapReduce, no security group exists. Enter a name to create a security group. If you have created security groups in E-MapReduce, select one as required. **Note:** The name must be 2 to 64 characters in length and can contain letters, digits, underscores (_), and hyphens (-). It must start with a letter.

Section	Parameter	Description
High Availability	**High Availability**	If the high availability mode is enabled, two master nodes are created in a Hadoop cluster to ensure the availability of the ResourceManager and NameNode processes. HBase clusters support high availability. In the original high availability mode, an HBase cluster uses a core node as the backup of the master node. If you enable the high availability mode described here, an HBase cluster uses two master nodes to achieve high availability, which is more secure and reliable.

Section	Parameter	Description
Instance	**Learn More**	• **Master Instance**: runs control processes, such as ResourceManager and NameNode. You can select an instance type as required. For more information, see Instance families. - **System Disk Type**: You can select an SSD, ESSD, or ultra disk based on your needs. - **Disk Size**: You can resize a disk based on your needs. The recommended minimum disk size is 120 GB. - **Data Disk Type**: You can select an SSD, ESSD, or ultra disk based on your needs. - **Disk Size**: You can resize a disk based on your needs. The recommended minimum disk size is 80 GB. - **Master Nodes**: One master node is configured by default. • **Core Instance**: stores all the data of a cluster. You can add core nodes as needed after a cluster is created. - **System Disk Type**: You can select an SSD, ESSD, or ultra disk based on your needs. - **Disk Size**: You can resize a disk based on your needs. The recommended minimum disk size is 120 GB. - **Data Disk Type**: You can select an SSD, ESSD, ultra disk, or local disk based on your needs. - **Disk Size**: You can resize a disk based on your needs. The recommended minimum disk size is 80 GB. - **Core Nodes**: Two core nodes are configured by default. You can change the number of core nodes as required. • **Task Instance**: stores no data but adjusts the computing capabilities of clusters. No task node is configured by default. You can add task nodes as required.

c) Configure basic parameters.

Section	Parameter	Description
Basic Information	**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 (_).
	Type	**Use User-created RDS** is recommended. • **Built-in MySQL**: Metadata is stored in the local MySQL database of a cluster. • **Unified Metabases**: Metadata is stored in an external metadatabase and is retained even after a cluster is released. For more information, see #unique_29. • **Use User-created RDS**: Metadata is stored in a user-created ApsaraDB for RDS database. .
	Assign Public Network IP	Specifies whether an Elastic IP Address (EIP) is associated with a cluster. **Note:** When **Type** is set to **Unified Metabases**, **Assign Public Network IP** is disabled by default. This feature can only be used for Kafka clusters. After a cluster is created, you can access the cluster only over the internal network. To access the cluster over the Internet, apply for a public IP address in ECS. For more information about how to apply for an EIP, see #unique_30.
	Remote Logon	Specifies whether to enable port 22 of a security group. Port 22 is disabled by default.
	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 must be 8 to 30 characters in length and contain uppercase letters, lowercase letters, digits, and special characters. The following characters are supported: ! @ # $ % ^ & *
Advanced Settings	**Add Knox User**	The user added to access the Web UIs of open-source big data software.

Section	Parameter	Description
	Permission Settings	The RAM roles that allow applications running in a cluster to access other Alibaba Cloud services. You can use the default RAM roles. • **EMR Role**: The value is fixed to AliyunEMRDefaultRole and cannot be changed. This RAM role authorizes a cluster to access other Alibaba Cloud services, such as ECS and OSS. • **ECS Role**: You can also assign an application role to a cluster. Then, E-MapReduce 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 on specific Alibaba Cloud services to the application role as required.
	Bootstrap Actions	Optional. You can configure bootstrap actions to run custom scripts before a cluster starts Hadoop. For more information, see Bootstrap actions.
	Tag	Optional. You can bind a tag when you create a cluster or bind a tag on the cluster details page after a cluster is created. For more information, see Cluster tags.
	Resource Group	Optional. Resource groups are designed for enterprise users. Enterprise users can use this feature to organize and manage their resources owned by multiple accounts and credited to multiple projects. If you do not specify a resource group, the cluster is automatically added to the default resource group. For more information, see Use resource groups.

Note:

The cluster configuration appears on the right side when you configure parameters. After you complete the preceding configuration and click Next: Confirm, you are directed to the Confirm step, in which you can confirm the configuration and the fee for the creation of your cluster. The fee varies with the billing method. For a pay-as-

you-go cluster, the price per hour appears. For a subscription cluster, the total fee appears.

d) Verify that the configuration is correct and click **Create**.

Notice:

- Pay-as-you-go clusters: Creation immediately starts after you click **Create**. On the **Cluster Management** tab that appears, a cluster whose status is **Initializing** appears in the cluster list. It takes several minutes to create a cluster. 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.

If a cluster fails to be created, **Creation Failed** appears in the Status column of the cluster. Move the pointer over the red exclamation point (!) to view the cause, as shown in the following figure.

You do not need to handle creation failures. If a cluster fails to be created, no computing resources are created. The cluster disappears from in the cluster list after three days.

3. Optional: Log on to a target core node.

a) On the master node, switch to the Hadoop user account by running the following command:

```
su hadoop
```

b) Log on to the target core node by using password-free SSH authentication.

```
ssh emr-worker-1
```

c) Run the following **sudo** command to obtain the root permissions:

```
sudo vi /etc/hosts
```

2.2 Cluster list and details

This topic describes the clusters in a cluster list, the cluster information contained in the cluster details, and supported cluster operations.

Prerequisites

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

Cluster Management tab

1. Log on to the Alibaba Cloud E-MapReduce console.
2.
3. Click the **Cluster Management** tab.

   The Cluster Management tab displays basic information about all of your clusters and the operations supported by each cluster.

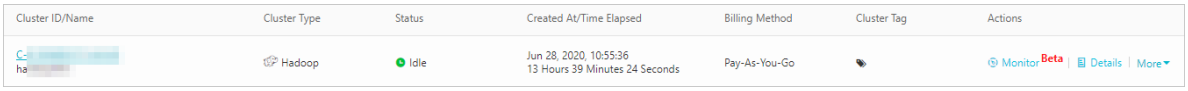

Parameter	Description
Cluster ID/Name	The ID and name of the cluster.
Cluster Type	The type of the cluster.
Status	The status of the cluster. For more information, see #unique_36. If a cluster experiences an exception, such as a creation failure, prompt information appears on the right. You can move the pointer over the prompt information to view detailed error information. You can also sort clusters by **Status**.
Created At/Time Elapsed	• Created At: the time the cluster was created. • Time Elapsed: the running time of the cluster from its creation to the current time. The time terminates after the created cluster is released.
Billing Method	The billing method of the cluster.

Parameter	Description
Actions	The operations supported on the cluster include: • **Monitor**: views the overview information about the cluster on the Cluster Overview page.  **Note:** EMR-APM is available only in the following regions: China (Hangzhou), China (Shanghai), China (Beijing), China (Zhangjiakou-Beijing Winter Olympics), China (Shenzhen), and US (Virginia). If you require EMR-APM in other regions, submit a ticket and contact the EMR product team. • **Details**: views detailed information about the cluster on the Cluster Overview page. • **More** - **CloudMonitor**: monitors the CPU idleness, memory capacity, and disk capacity of the EMR cluster to help users monitor the running status of the cluster. - **Scale Out**: scales out the cluster. For more information, see Scale out a cluster. - **Release**: releases the cluster. For more information, see Release a cluster. - **Restart**: restarts the cluster.

View cluster details

1. Log on to the Alibaba Cloud E-MapReduce console.
2. Click the **Cluster Management** tab.
3. Find the target cluster and click **Details** in the Actions column.

   The Cluster Overview page displays detailed information about your cluster and includes four sections: Cluster Info, Software Info, Network Info, and Instance Info.

   • Cluster Info section

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 #unique_36.

Parameter	Description
Tag	The cluster tag. For more information, see Cluster tags.
I/O Optimization	Specifies whether I/O optimization is enabled.
High Availability	Specifies whether high availability is enabled.
Security Mode	The security mode. The software of the cluster starts in Kerberos security mode. For more information, see Introduction to Kerberos.
Start Time	The time the cluster was created.
Billing Method	The billing method of the cluster.
Time Elapsed	The running time of the cluster.
Unified Metabases	Specifies whether the unified metadatabase is enabled.
Bootstrap Actions/ EMR Version	Information about custom scripts and software configuration.
ECS Role	The ECS application role. You can also assign an application role to a cluster, so 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 application role access permissions on specific Alibaba Cloud services as required.

- Software Info section

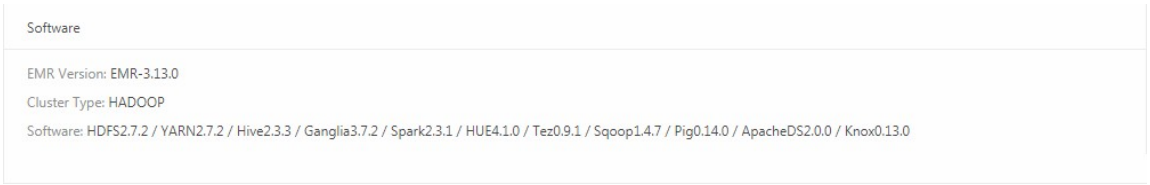

Parameter	Description
EMR Version	The major version of EMR.
Cluster Type	The type of the cluster.

Parameter	Description
Software Info	All user-installed applications and their versions, such as HDFS 2.7.2, Hive 2.3.3, and Spark 2.3.1.

- Network Info section

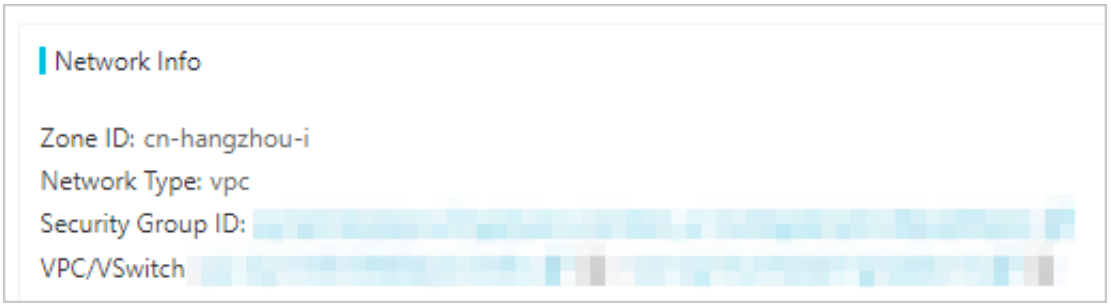

Parameter	Description
Zone ID	The zone where the cluster resides, such as cn-hangzhou-b, which is also the zone where the ECS instance resides.
Network Type	The network type of the cluster.
Security Group ID	The ID of the security group to which you want to add your cluster.
VPC/VSwitch	The ID of the VPC to which your cluster belongs and the ID of the subnet VSwitch.

- Instance Info section

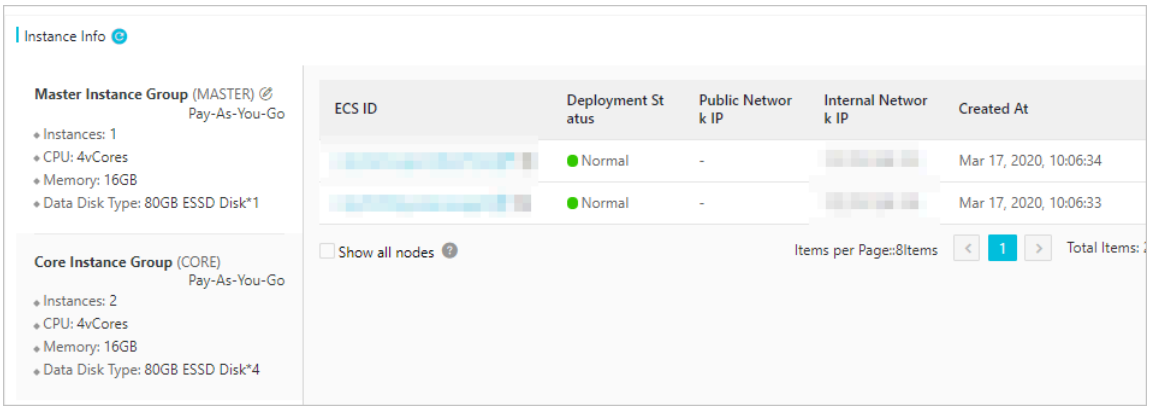

Parameter	Description
Instances	The current number of nodes and the actual number of nodes for which you applied. Theoretically, the two values must be the same. However, the current number is less than the actual number until creation is complete.
CPU	The number of CPU cores of a single node.
Memory	The memory capacity of a single node.

Parameter	Description
Data Disk Type	The data disk type and data disk capacity of a single node.
ECS instance list	The information about ECS instances in a master node group, including: - **ECS ID**: the ID of the purchased ECS instance - **Deployment Status**: including Initializing, Normal, and Scaling out - **Public Network IP**: the public IP address of the ECS instance - **Internal Network IP**: the private IP address of the ECS instance, which can be accessed by all nodes in the cluster - **Created At**: the creation time of the purchased ECS instance

2.3 Scale out a cluster

This topic describes how to scale out an E-MapReduce (EMR) cluster when its computing or storage resources are insufficient. Currently, both core nodes and task nodes can be added. The configurations of an added node is the same as those of an existing node by default.

Scale out a cluster

1. Log on to the Alibaba Cloud E-MapReduce console.
2. In the top navigation bar, select the region where your cluster resides.
3. Click the **Cluster Management** tab.
4. Find the target cluster and click **Details** in the Actions column.
5. On the **Cluster Overview** page that appears, select **Scale Out** from the **Change Configuration** drop-down list in the upper-right corner.

6. In the **Scale Out** dialog box that appears, click **CORE (Core Instance Group)** or **TASK (Task Instance Group)** and specify related parameters.

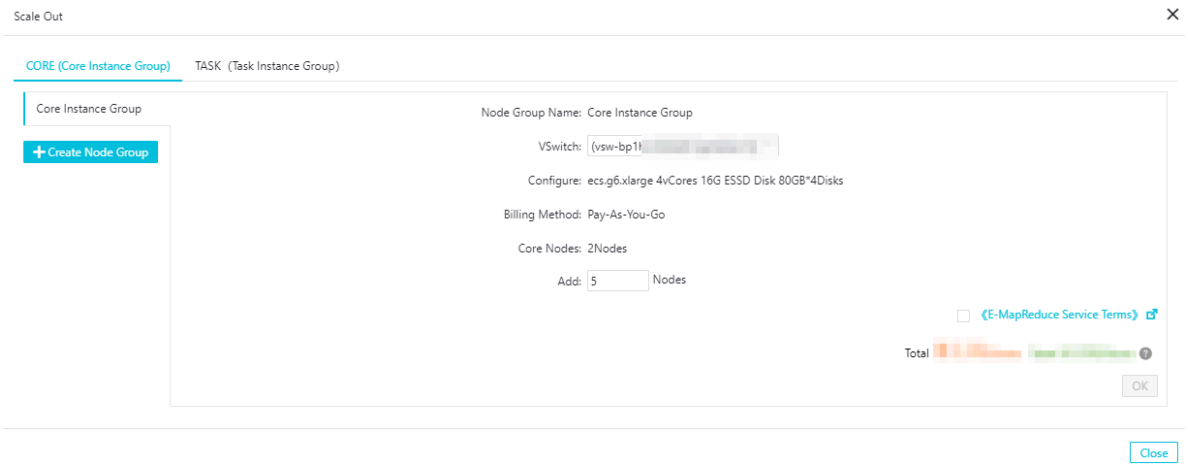

The following table uses CORE (Core Instance Group) as an example to describe the required parameters.

Parameter	Description
VSwitch	The VSwitch deployed in the cluster.
Configure	The configurations of the current instance 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 subscription period for the new node.
Core Nodes	The total number of your existing core nodes is displayed by default.
Add	The number of core nodes you want to add. Click the upward or downward arrow or enter a number in the Add field. **Note:** When you adjust the number, the total fee for scaling out a cluster is displayed in the lower-right corner of the Scale Out dialog box.
E-MapReduce Service Terms	Read and agree to the service terms.

7. Click **OK**.

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.

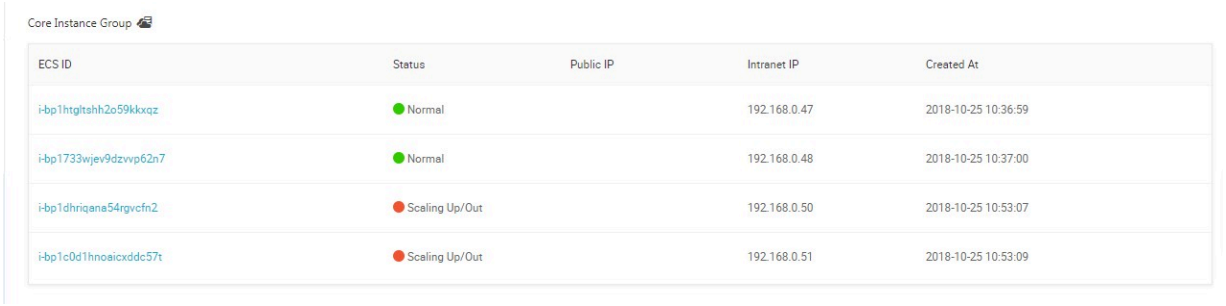
Core Instance Group

ECS ID	Status	Public IP	Intranet IP	Created At
i-bp1htgltshh2o59kkxqz	Normal		192.168.0.47	2018-10-25 10:36:59
i-bp1733wjev9dzvvp62n7	Normal		192.168.0.48	2018-10-25 10:37:00
i-bp1dhriqana54rgvcfn2	Scaling Up/Out		192.168.0.50	2018-10-25 10:53:07
i-bp1c0d1hnoaicxddc57t	Scaling Up/Out		192.168.0.51	2018-10-25 10:53:09

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. Run the following command to log on to the master node of the cluster:

   ```
   ssh root@ip.of.master
   ```

   Obtain the public IP address of the master node. For more information, see Cluster list and details.
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 Cluster list and details.
4. Run the following command to change the password of the root user:

   ```
   sudo passwd root
   ```

2.4 Scale in a cluster

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

Prerequisites

Currently, you can only reduce the number of task nodes to scale in an E-MapReduce cluster. The cluster you want to scale in must meet the following requirements:

- If the E-MapReduce version is V2.X.X, the version is later than V2.5.0. If the E-MapReduce version is V3.X.X, the version is 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. Log on to the Alibaba Cloud E-MapReduce console.
2. In the top navigation bar, select the region where your cluster resides. Select the resource group as required. By default, all resources of the account appear.
3. Click the **Cluster Management** tab.
4. On the **Cluster Management** page that appears, find the target cluster and click **Details** in the Actions column.
5. On the **Cluster Overview** page that appears, select **Scale In** from the **Change Configuration** drop-down list in the upper-right corner.
6. In the **Scale In** dialog box that appears, click the downward arrow for Reserved Nodes or enter a number in this field to set the number of task nodes to retain.
7. Click **Scale In**.

   In the **Confirm Scale-in Activity** dialog box that appears, click **OK**.

2.5 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

The cluster you want to release is in the Initializing, Running, or Idle state.

Background information

After you confirm the release of a cluster, the system handles the cluster as follows:

1. Forcibly terminates all jobs in the cluster.
2. Terminates and releases all ECS instances 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 five minutes. In most cases, the ECS instances can be released in seconds.

Procedure

1. Log on to the Alibaba Cloud E-MapReduce console.
2. Click the **Cluster Management** tab.
3. Find the target cluster and select **Release** from the **More** drop-down list in the Actions column.

   Alternatively, find the target cluster and click **Details** in the Actions column. On the **Cluster Overview** page that appears, select **Release** from the **Instance Status** drop-down list.
4. In the **Cluster Management-Release** message that appears, click **Release**. The system notifies you of the successful cluster release.

2.6 Cluster renewal

When your Subscription cluster is about to expire, you have to renew it in order to continue using E-MapReduce cluster services. Cluster renewal includes the renewal of E-MapReduce services and ECS instances.

Enter the renewal page

1. Log on to the Alibaba Cloud E-MapReduce console.
2. At the top of the page, click **Cluster Management**.
3. In the cluster list, select the cluster that you want to renew.
4. To the right the cluster, click **Renew** to enter the cluster renewal page.

Renewal page

As shown in the following figure, the renewal page contains a number of columns, which are detailed below.

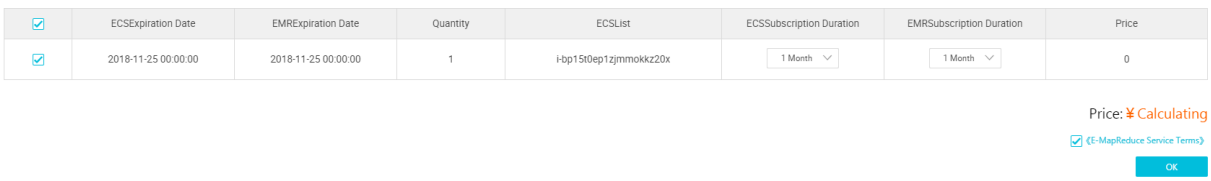

- **ECS Expiration Date**: The expiration date of an ECS instance.
- **EMR Expiration Date**: The expiration date of E-MapReduce services.
- **Quantity**: The number of machines for instance groups.
- **ECS List**: The ECS instance ID of the machine in the cluster.

- **ECS Subscription Duration**: The duration of your ECS subscription. You can select from one to nine months, or one, two, or three years.
- **EMR Subscription Duration**: The duration of your E-MapReduce subscription. We recommend that you keep it consistent with ECS.
- **Price**: The price of the renewal of E-MapReduce services and ECS instances.

Make a payment

> **Notice:**
> The fees consist of the combined costs of ECS renewal and E-MapReduce service products. If there are unpaid orders in the cluster list, you cannot expand or renew any clusters.

1. Click **OK** to view the prompt box for a successful order placement.
2. Click **Go to the payment page**. The payment page displays the total amount to be paid as well as the order details.
3. Click **Confirm payment**.
4. After you make the payment, click **Payment completed** to return to the cluster list page.

If an expired cluster is successfully renewed, its expiration date is updated to reflect the renewal. If an expired ECS instance is successfully renewed, its expiration date is usually updated around three to five minutes later.

If you confirm the renewal, but fail to pay for it, **Cancel order** and **Make the payment** are displayed on the right side of the cluster. Click **Make the payment** to complete the payment or **Cancel** to cancel the renewal.

2.7 Cluster tags

Tags are used to label clusters. You can use tags to identify and manage your cluster resources. You can bind tags during the creation of a cluster or on the basic information page of a created cluster. You can bind a maximum of 20 tags to a cluster.

Limits

> **Notice:**
> - Tags of ECS instances are also be updated when you update E-MapReduce (EMR) cluster tags.
> - However, EMR cluster tags are not updated when you update tags of ECS instances . To ensure the consistency between tags for ECS instances and EMR clusters, we

recommend that you do not modify 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.

- Each tag consists of a key-value pair.
- A tag key cannot be empty and can have a maximum of 64 characters. A tag value can be empty and can have a maximum of 128 characters.
- Tag keys and tag values cannot start with aliyun or acs: and cannot contain http:// or https://.
- A tag key must be unique. For example, **city:shanghai** is bound to a cluster. If **city:newyork** is then bound to the cluster, **city:shanghai** is automatically unbound.
- Tag information is not shared across regions. For example, tags created in the China (Hangzhou) region are not visible to the China (Shanghai) region.
- After you unbind a tag from a cluster, the system automatically deletes the tag if the tag is not bound to any other clusters.

Create or bind a tag

1. Log on to the Alibaba Cloud E-MapReduce console.
2. 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.
3. Click **Cluster Wizard** in the Clusters section.
4. In the **Advanced Settings** section of the Basic Settings step, click **+ Add** on the right side of **Tag**.

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

5. In the **Create/Bind Tag** dialog box that appears, select an existing tag key or enter a new one to set **Tag Key**. Tag Value is optional. You can select an existing tag value or enter a new value.

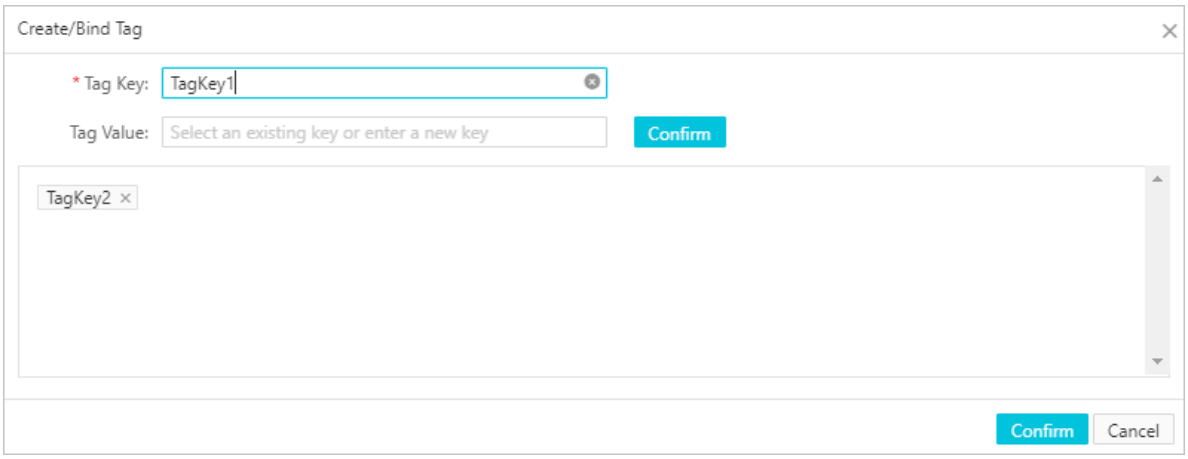

> **Note:**
> If you select an existing tag key and tag value, it indicates that you want to bind an existing tag. If you enter a new tag key, it indicates that you want to create a tag.

6. Click **Confirm**.

   Confirm the tag to be bound.

7. Click **Confirm**.

   Bind the tag.

Use tags to search for clusters

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

1. Log on to the Alibaba Cloud E-MapReduce console.
2. On the EMR homepage, click the **Cluster Management** tab.
3. On the **Cluster Management** tab that appears, click **Tag** and select a tag key. If you do not select a tag value, all clusters bound to the tag key are displayed.
4. Click **Search** next to **Tag**.

   Clusters that meet the search criterion are displayed in a list.

Manage cluster tags

Bind, unbind, or view tags of existing clusters.

1. Log on to the Alibaba Cloud E-MapReduce console.
2. On the EMR homepage, click the **Cluster Management** tab.

3. Click **Details** in the Actions column that corresponds to a cluster in the list.
4. On the **Cluster Overview** page, click **Edit Tag** in the **Cluster Info** section.
5. In the **Edit Tag** dialog box that appears, configure the following information:
   - Create a tag.
     a. Click **Add Tag** in the lower-left corner.
     b. Set **Tag Key** and **Tag Value**.
     c. Click **Confirm**.
     d. Click **Confirm**.
   - Bind an existing tag.
     a. Click **Existing Tags** in the lower-left corner.
     b. Select an existing **tag key** and the **tag value** that matches.
     c. Click **Confirm**.
   - Unbind a tag.

     Click next to a tag to unbind this tag.

> **Note:**
> The **tag value** is optional.

System tags

When you create an EMR cluster, two system tags are automatically bound to each ECS instance. You can perform the following operations to identify the cluster to which an ECS instance belongs and the role of the ECS instance in the cluster: Log on to the Alibaba Cloud ECS console. Click **Tags** in the left-side navigation pane. On the Tags page that appears, click **System Tags** in the left-side navigation tree on the Tags tab.

For example, the following information shows system tags for an ECS instance, which indicates that the ECS instance serves as a core node in the EMR cluster with the ID of C-A510C93EA117XXXX.

- acs:emr:clusterId=C-A510C93EA117XXXX
- acs:emr:hostGroupType=CORE

2.8 Use resource groups

You can group your resources based on usage, permissions, and regions. This achieves hierarchical resource management based on users and projects. 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

Note the following limits on resource groups:

- You can create and manage a resource group and authorize RAM users to access a resource group only in the Resource Management console. For more information, see #unique_41.
- 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.
- Resources can be moved only across resource groups under the same account.

   **Notice:**
  If you move a node resource to a different resource group, the resource group of the cluster to which the node belongs remains unchanged. 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.

Note the following points when you use resource groups:

- A resource group can contain resources in different regions. For example, Resource Group A contains both clusters in China (Hangzhou) and clusters in China (Shanghai).
- A resource group can contain different types of resources. For example, Resource Group A contains clusters, ECS instances, and projects.
- Clusters and projects under the same account can be associated across resource groups if they reside in the same region. For example, your account can access both Resource Group A and Resource Group B. Resource Group A contains a project that resides in China (Beijing). Resource Group B contains a cluster that also resides in China (Beijing). In this case, you can run a job of this project in this cluster.

- RAM users have access to resource groups as authorized. For example, if you authorize a RAM user to manage all Alibaba Cloud resources under your Alibaba Cloud account, all resource groups under your Alibaba Cloud account are accessible to the RAM user.

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 a default resource group. This section describes how to specify a resource group when you create a cluster and a project.

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

- Specify a resource group during cluster creation:

  1. Log on to the Alibaba Cloud E-MapReduce console.
  2. In the top navigation bar, select the region where you want to create a cluster.
  3. Click the Cluster Management tab. Then, click **Cluster Wizard** in the upper-right corner of the tab.
  4. In the **Advance Settings** section of the **Basic Settings** step, select a resource group from the **Select Resource Group** drop-down list.

     If you want to create a new resource group, click the **link** below the Select Resource Group drop-down list. For more information, see #unique_42.

     > **Note:**
     > For more information about how to create a cluster, see Create a cluster.

- Specify a resource group during project creation:

  1. Log on to the Alibaba Cloud E-MapReduce console.
  2. In the top navigation bar, select the region where you want to create a project.
  3. Click the **Data Platform** tab.
  4. Click **Create Project** in the upper-right corner of the tab.
  5. In the **Create Project** dialog box, specify **Project Name** and **Project Description** and select a resource group from the **Select Resource Group** drop-down list.

     If you want to create a new resource group, click the **link** below the Select Resource Group drop-down list. For more information, see #unique_42.

     Note:
     If you need to create a cluster to associate it with the project, see Create a cluster.
  6. Click **Create**.

     For more information about how to create a project, see #unique_43.

Scenarios

Resource groups can be used in the following scenarios:

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

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

Scenario 1: Add resources for different purposes to different resource groups

You can add clusters to be used in a test environment and those to be used in a production environment to different resource groups. During the test stage, 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 functions properly.

1. Create resource groups **Test Environment** and **Production Environment**.

   For more information, see #unique_42.

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

   For more information, see #unique_44.

3. Create clusters **TestEnv1** and **TestEnv2**.

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

4. Create clusters **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 log on to the .

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 their permissions by resource group

You can add clusters and projects of a department to an independent resource group and set an independent administrator for each resource group. This helps achieve separated management of users and permissions. In the following procedure, resource groups for a development department and a test department are used as an example.

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

   For more information, see #unique_42.

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

   For more information, see #unique_44.

3. Create cluster **ITCluster** and project **ITFlowProject** and specify the **Develop Dept** resource group for them.

   For more information about how to create a cluster, see Create a cluster. For more information about how to create a project, see #unique_43.

4. Create clusters **FinanceCluster1** and **FinanceCluster2**. Specify the **Test Dept** resource group for the clusters.

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

5. Use the administrator account of the Test Dept resource group to log on to the .
6. Select the **Test Dept** resource group from the top navigation bar.

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

2.9 Connect to a cluster using SSH

You can connect to the master node using SSH to view detailed settings and status of the jobs in the CLI. The public IP address of the master node is shown on the cluster overview page.

Environment variables on a cluster

Environment variables have been configured on the instances. Frequently-used environment variables are as follows.

- JAVA_HOME
- HADOOP_HOME
- HADOOP_CONF_DIR
- HADOOP_LOG_DIR
- YARN_LOG_DIR
- HIVE_HOME
- HIVE_CONF_DIR
- PIG_HOME
- PIG_CONF_DIR

You can reference the variables in the script. We recommend that you do not modify the values of the variables to avoid unexpected errors in the EMR cluster.

Connect to the master node

1. Run the following command to connect to the master node using SSH. Check the public IP address of the master node in the Host section on the Cluster list and details page.

   ```
   ssh root@ip.of.master
   ```

2. Enter the password you have set when creating the cluster.

Connect to the master node using SSH without a password

Typically, you need to connect to the cluster for management and operations. You can quickly connect to the master node by using SSH without a password. The master node in a cluster is assigned with a public IP address by default. Procedure:

1. Connect to the master node as the root user by using a password.
2. Switch to the Hadoop user or the HDFS user.

Use SSH on Linux

1. Send the private key file to your local system.

   ```
   sz ~/.ssh/id_rsa
   ```

2. Go back to your local machine and connect to the master node again.

   ```
   ssh -i private_key_path/id_rsa hadoop@120.26.221.130
   ```

   If you only have one private key file, you can store it under the ~/.ssh/ directory for immediate use without specifying the key path using the `-i` option every time.

Use SSH on Windows

You have multiple methods to connect to the master node using SSH without a password on Windows. The methods are shown as follows.

- Method 1: Use PuTTY

  1. Click Download PuTTY.
  2. Download PuTTYgen through this link.
  3. Start PuTTYgen and load your private key.

     Notice:
     Keep the private key safe. Generate a new key once the private key is stolen.

  4. Save the private key with the default configurations. A PuTTY private key file with the .ppk extension is generated.
  5. Run PuTTY and choose Session on the configuration page.
  6. Enter the public IP address of the node that you want to connect to with a logon username. For example, hadoop@MasterNodeIP.
  7. Select **Connetion** > **SSH** > **Auth** on the configuration page.
  8. Select the generated PPK file.
  9. Click **Open** to connect to the master node.

- Method 2: Use Cygwin or MinGW

  Cygwin and MinGW are easy-to-use tools for compiling Linux on Windows.

  You can refer to the operations described in the Use SSH on Linux section.

  We recommend that you use MinGW, which is the most light-weight method for using SSH on Windows. If you cannot load the official website, download Git for Windows and use the built-in Git BASH instead.

View the Web UIs of Hadoop, Spark, and Ganglia

Note:

Make sure that you have completed the preceding Connect to the master node using SSH without a password procedure.

For security reasons, the ports of Web UIs for monitoring Hadoop, Spark, and Ganglia in an EMR cluster are closed. If you want to access these Web UIs, you need to create SSH tunnels and enable port forwarding. The methods are shown as follows:

Notice:

The following operations are based on your local machine, not on a cluster node.

- Method A: Dynamic port forwarding
  - Use a private key:

    Create an SSH tunnel to allow communication between your local machine and a dynamic port on the master node in the EMR cluster.

    ```
    ssh -i /path/id_xxx -ND 8157 root@masterNodeIP
    ```

  - Use a username and password:

    ```
    ssh -ND 8157 root@masterNodeIP
    ```

  Replace 8157 with any port number on your local machine that has not been used.

  After dynamic port forwarding is complete, you can view a Web UI using the following methods.

  - Recommended method:

    We recommend that you use the Chrome browser. You can run the following command to access a Web UI.

    ```
    chrome --proxy-server="socks5://localhost:8157" --host-resolver-rules="MAP * 0.0.0.0 , EXCLUDE localhost" --user-data-dir=/tmp/
    ```

    For Windows, an example of the temporary file path is d:/tmppath. For Linux and OSX, the format of the temporary file path is like /tmp/.

    Chrome is installed on different paths based on the operating system. For more details, see the following table.

OS	Chrome path
Mac OS X	/Applications/Google Chrome.app/Contents/MacOS/Google Chrome
Linux	/usr/bin/google-chrome

OS	Chrome path
Windows	C:\Program Files (x86)\Google\Chrome\Application\chrome.exe

- Extension:

■ Use a Chrome extension to view the Web UIs.

1. Install the Chrome extension SwitchyOmega.
2. After installation is complete, click the icon of SwitchyOmega and click **Options** in the pop-up menu to perform configurations.
3. Click **New Profile**. In the Profile name field, enter "SSH tunnel." Click PAC Profile for the type of the profile.
4. In the PAC Script editor, enter the following code.

```
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';
}
```

5. In the left-side navigation pane, click **Apply changes** to complete the configurations.
6. Start a CLI. Choose one of the following methods and run the corresponding command.

```
//Method a: Use a private key
ssh -i /path/id_xxx -ND 8157 hadoop@masterNodeIP
```

```
//Method b: Use a username and password
ssh -ND 8157 hadoop@masterNodeIP
```

7. Click the SwitchyOmega icon in the Chrome menu. From the drop-down list, select SSH tunnel.
8. In the address bar, enter the IP address of the node and the port number to access the Web UI. The node refers to the one that you want to connect to using

the SSH command. Generally, it is a master node. Two frequently-used ports are port 8088 for YARN and port 50070 for HDFS.

By using this method, you can browse webpages and use Web UIs of clusters at the same time.

- After creating an SSH tunnel between your local machine and the master node of the EMR cluster, you need to configure a local agent for viewing the Web UIs of Hadoop, Spark, and Ganglia through browsers. Procedure:

  1. Assume that you use Chrome or Firefox. Click Download FoxyProxy Standard.
  2. After the installation is complete, restart your browser, open a text editor, and then enter the following code.

```
<? xml version="1.0" encoding="UTF-8"? >
<foxyproxy>
<proxies>
<proxy name="aliyun-emr-socks-proxy" id="2322596116" notes=""
fromSubscription="false" enabled="true" mode="manual" selectedTabIndex
="2" lastresort="false" animatedIcons="true" includeInCycle="true" color
="#0055E5" proxyDNS="true" noInternalIPs="false" autoconfMode="pac"
clearCacheBeforeUse="false" disableCache="false" clearCookiesBeforeUse="
false" rejectCookies="false">
<matches>
<match enabled="true" name="120.*" pattern="http://120.*" isRegEx="false"
isBlackList="false" isMultiLine="false" caseSensitive="false" fromSubscription
="false" ></match>
</matches>
<manualconf host="localhost" port="8157" socksversion="5" isSocks="true"
username="" password="" domain="" ></manualconf>
</proxy>
</proxies>
</foxyproxy>
```

     In this example:

     - Port `8157` on your local machine is used to establish an SSH tunnel to the master node of the cluster.
     - Replace `120.*` with the actual IP address of the master node.
  3. Click the **Foxyproxy** icon in the Chrome menu. From the drop-down list, select **Options**.
  4. Select **Import/Export**.
  5. Select the XML file you have created. Click **Open**.
  6. In the **Import FoxyProxy Setting** dialogue box, click **Add**.
  7. Click the **Foxyproxy** icon. From the drop-down list, select **Use Proxy aliyun-emr-socks-proxy for all URLs**.
  8. Enter localhost: 8088 in the address bar to view the Web UI of Hadoop.

- Method 2: Local port forwarding

  Notice:
  If you use this method to view a Web UI, errors will occur when you try to open the subpages.

  - Use a private key:

    ```
    ssh -i /path/id_rsa -N -L 8157:masterNodeIP:8088 hadoop@masterNodeIP
    ```

  - Use a username and password:

    ```
    ssh -N -L 8157:masterNodeIP:8088 hadoop@masterNodeIP
    ```

  Parameter description:

  - **path**: The path where the private key file is stored.
  - **masterNodeIP**: The IP address of the master node to be connected.
  - **8088**: The port that is used by the Resource Manager Web UI on the master node.

2.10 View the service list

You can view the running status of services, such as HDFS and YARN, on the Clusters and Services page.

Procedure

1. Log on to the Alibaba Cloud E-MapReduce console.
2. In the top navigation bar, select the region where your cluster resides. Select the resource group as required. By default, all resources of the account appear.
3. Click the **Cluster Management** tab.
4. In the cluster list, find the target cluster and click its cluster ID.

   On the page that appears, you can view all services in the cluster.

   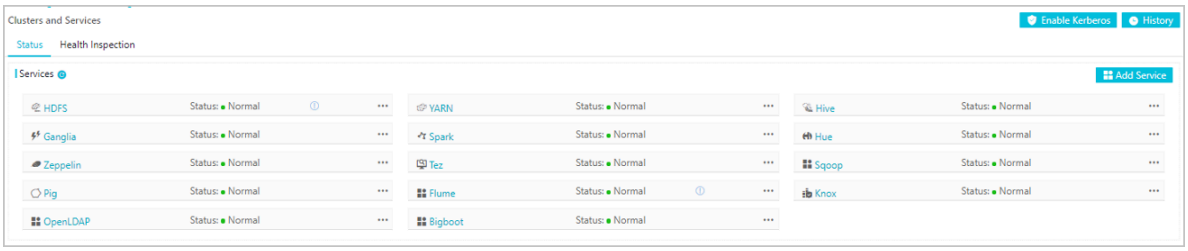

  Note:

Only the services you selected when you created the cluster appear in the Services section.

5. Click the name of a service in the Services section.

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

   The service status can be **Normal** or **Error**. If a service on a specific node is in the **Error** state, you can log on to the node from the master node to fix the error.

2.11 Cluster scripts

This topic describes how to select nodes in batches and run specific scripts on these nodes to meet your business requirements, such as installing third-party software and modifying the running environment of a cluster after you create it, especially a subscription cluster.

Prerequisites

- A cluster is in the Idle or Running state. Scripts cannot be executed on clusters in other states.
- Cluster scripts are developed or obtained and uploaded to Object Storage Service (OSS). For more information about the cluster scripts, see Examples.

Overview

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. For example:

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

Features

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 up to 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 executed, 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 path, the operational log is also stored in the osslogpath/clusterId/ip/cluster-scripts/clusterScriptId path. To create a cluster script, follow these steps:

1. Log on to the Alibaba Cloud E-MapReduce console.
2.
3. Click the **Cluster Management** tab.
4. Find the target cluster and click **Details** in the Actions column.
5. In the left-side navigation pane of the Cluster Overview page that appears, click **Cluster Scripts**.
6. On the **Cluster Scripts** page that appears, click **Create and Run** in the upper-right corner.
7. In the **Create Script** dialog box, specify **Name**, **Script**, and **Target Nodes**. (You can select multiple nodes.)

   

   Note:

   We recommend that you test the cluster script feature on a single node before you use the feature on the entire cluster.
8. Click **OK**.

   After a cluster script is created, it is displayed in the cluster script list and is in the running state.
9. 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.

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 oss://yourbucket/myfile.tar.gz file is downloaded to your local machine and decompressed to the /yourdir directory.

```
#! /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>
```

The osscmd tool is pre-installed on a specific node. You can use this tool to download the file.

Note:

The OSS address specified by using **host** 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.

Additional system software packages can also be installed by using YUM, for example, ld-linux.so.2.

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

The root account is used to run a specified script on a cluster by default. You can run the **su hadoop** command in the script to switch to the hadoop user.

2.12 Security groups

Currently, security groups created in the EMR console are required for creating an EMR cluster.

When creating an E-MapReduce cluster, you can open port 22 of the security group where the cluster resides (It is closed by default, open in **Create Cluster** > **Basic Configuration** > **Remote Logon**).We recommend that you add ECS instances to security groups based on their functions. For example, the security group for an EMR cluster is **EMR security group**. **user security group** is an existing security group that you have created. Instances of each security group are configured to open different ports based on the IP addresses allowed to access the instances.

Add an EMR cluster to an existing security group

Note:

- You need to add instances in a classic network to a classic network security group of the same region.
- You need to add instances in a VPC to a VPC security group of the same region.

1. Log on to the Alibaba Cloud E-MapReduce console.
2. Click the **Cluster Management** tab.
3. In the cluster list, locate the cluster that you want to add to the security group. Click **View Details** in the Actions column for the cluster to go to the Cluster Overview page.
4. On the Cluster Overview page, click the **security group ID** in the **Network** area.
5. In the left-side navigation pane, click **Instances in Security Group** to view the IDs and names of all instances.
6. Go to the Alibaba Cloud ECS console. In the left-side navigation pane, click **Instances**.
7. The ECS instances listed on the **Instances** page, whose names start with emr-xxx, are instances in an EMR cluster. Locate the instance that you want to add to a security group, click **Manage** in the **Actions** column for the instance.
8. Click **Security Groups**.
9. Click **Add to Security Group**.
10. Select the security group to which you want to add the instance. Click **Join multiple security groups**. From the Security Group drop-down list, select the security groups as needed. The security groups are displayed in the dotted box below.
11. Click **OK**.

Add an existing cluster to an EMR security group

You can refer to the operations described in the Add an EMR cluster to an existing security group section. The full process starts with locating the security group of the existing cluster and ends with adding the instances to the EMR security group. For instances that do not belong to any security groups, select the instances and choose More > Network and security > Add to Security Group to add multiple instances to the EMR security group at a time.

Security group rules

If an ECS instance belongs to multiple security groups, the instance follows the union of all security group rules. For example, the security group rule of the EMR security group is configured to open port 22 and the security group rules of the **user security group** are configured to open all ports. After you add the EMR cluster to the **user security group**, the instances of the EMR cluster open all ports. We recommend that you use security groups carefully.

2.13 Access open-source components

This topic describes how to configure security groups so that you can access open-source components in an E-MapReduce cluster. After a cluster is created, E-MapReduce binds several domain names to your cluster by default for you to access the following open-source components: YARN, HDFS, Spark, Ganglia, Hue, and Zeppelin.

Prerequisites

An Elastic IP Address (EIP) is assigned to the E-MapReduce cluster.

Grant access to a security group

If you use a component for the first time, follow these steps to grant access to the corresponding security group:

1. Obtain your current public IP address.

   For security purposes, we recommend that you only allow access from the current public IP address when you configure a security group policy. To obtain your current public IP address, access ip.taobao.com. You can view your public IP address in the lower-left corner.

2. Add a security group policy.

   a) Log on to the Alibaba Cloud E-MapReduce console.

   b) In the top navigation bar, select the region where your cluster resides. Select the resource group as required. By default, all resources of the account appear.

   c) Click the **Cluster Management** tab.

   d) On the **Cluster Management** page that appears, find the target cluster and click **Details** in the Actions column.

   e) In the **Network Info** section of the **Cluster Overview** page that appears, take note of the network type of the cluster and click the ID of the target security group.

   f) On the page that appears, click **Add Security Group Rule** in the upper-right corner and enable the required ports.

   Notice:

It is forbidden to set **Authorization Object** to **0.0.0.0/0**, which may cause security issues.

The following table lists the ports you need to enable for accessing the Web UIs of different components.

Component	Port to be enabled
YARN	8443
HDFS	**Note:** After Ranger is deployed in your cluster, you can access the Web UI of Ranger.
Spark History Server	
Ganglia	
Ranger	
Zeppelin	8080
Hue	8888

For example, you can follow these steps to enable port **8443**:

A. Click **Add Security Group Rule**.

B. 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, **VPC** is used as an example.
- Follow the principle of least privilege when you configure inbound and outbound rules for applications. Enable only the ports required by your applications.

g) View the policy on the **Inbound** tab.

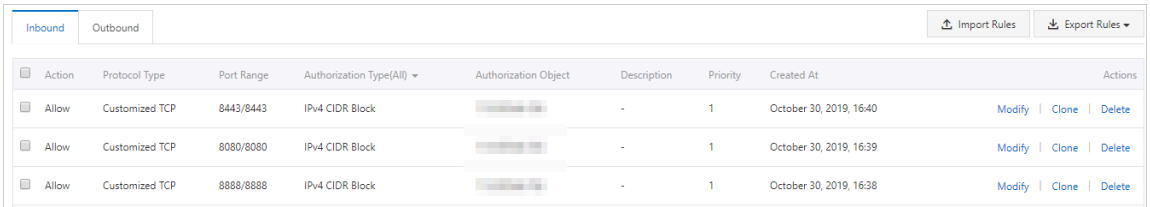

The network access is securely enabled and the network configuration is complete.

Access the Web UIs of open-source components

1. Log on to the Alibaba Cloud E-MapReduce console.
2. In the top navigation bar, select the region where your cluster resides. Select the resource group as required. By default, all resources of the account appear.
3. Click the **Cluster Management** tab.
4. On the **Cluster Management** page that appears, find the target cluster and click **Details** in the Actions column.
5. In the left-side navigation pane of the Cluster Overview page that appears, click **Connect Strings**.
6. On the **Connect Strings** page that appears, find the target component and click its link to access the Web UI of this component.

Note:

- After a cluster is created, E-MapReduce binds several domain names to your cluster by default for you to access the following open-source components: YARN, HDFS, Spark, Ganglia, Hue, and Zeppelin.
- 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 HDFS, YARN, Spark, and Ganglia. For more information about how to create a Knox account, see Manage users. For more information about how to use Knox, see Knox. Enter the Hue username and password to access the Web UI of Hue. For more information about how to use Hue, see #unique_49. 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 #unique_50.

2.14 Create a gateway cluster

A gateway cluster is composed of Elastic Compute Service (ECS) instances that reside in the same internal network as the E-MapReduce (EMR) cluster associated with the gateway cluster. You can use a gateway cluster to achieve load balancing and security isolation and to submit jobs to the EMR cluster.

Prerequisites

An EMR gateway cluster only supports Hadoop and Kafka clusters. Before you create a gateway cluster, make sure that you have created an EMR Hadoop or Kafka cluster.

1. Log on to the Alibaba Cloud E-MapReduce console.
2. Click the **Cluster Management** tab.
3. Click **Create Gateway** in the upper-right corner.
4. On the **Create Gateway** page that appears, specify the required parameters.

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 Network IP	Specifies whether an EIP address is associated with a gateway cluster.
Password/Key Pair	• **Password**: the password used to log on to the gateway cluster. The password must be 8 to 30 characters in length and contain uppercase letters, lowercase letters, digits, and special characters. The following special characters are supported: ! @ # $ % ^ & * • **Key Pair**: the name of the key pair used to log on to the gateway cluster. If no key pair is created, click **Create Key Pair** next to this field to go to the SSH Key Pairs page of the ECS console and create a key pair. Keep the .pem private key file safe. After a gateway cluster is created, 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.
Billing Method	• **Subscription**: The system charges you for a cluster only once per subscription period. The unit price of a subscription cluster is lower than that of a pay-as-you-go cluster of the same specification. Subscription clusters are offered discounts, and longer subscription periods offer larger discounts. • **Pay-As-You-Go**: The system charges you for a cluster based on the time the cluster is actually used in hours.

Parameter	Description
Associated Cluster	The Hadoop cluster associated with the gateway cluster. The gateway cluster submits jobs to this Hadoop cluster. The gateway cluster is automatically configured with the same Hadoop environment as the Hadoop cluster.
Zone	The zone where the associated cluster resides.
Network Type	The network type of the associated cluster.
VPC	The Virtual Private Cloud (VPC) to which the associated cluster belongs.
VSwitch	The VSwitch you want the gateway cluster to use. Select the VSwitch that corresponds to the zone and VPC.
Security Group Name	The name of the security group to which the associated cluster belongs.
Gateway Instance	The available instance types in the current zone. For more information, see Instance families. • **System Disk Type**: the type of the system disk you want the gateway cluster to use. System disk types include ultra disk, standard SSD, and ESSD. The system disk types that are available for you to create a gateway cluster depend on the selected region and instance type. The system disk is released when the relevant cluster is released. • **Disk Size**: the size of the system disk. Unit: GB. Valid values: 40 to 2048. Default value: 300. • **Data Disk Type**: the type of data disks you want the gateway cluster to use. Data disk types include ultra disk, standard SSD, and ESSD. The data disk types that are available for you to create a gateway cluster depend on the selected region and instance type. Data disks are released when the relevant cluster is released. • **Disk Size**: the size of a data disk. Unit: GB. Valid values: 200 to 4000. Default value: 300. • **Count**: the number of data disks. Valid values: 1 to 10.

5. Select E-MapReduce Service Terms and click **Create**.

   The created gateway cluster appears in the cluster list and its state changes from **Initializing** to **Idle** a few minutes later.

2.15 Manage cluster resources

2.15.1 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.

2.15.2 Use Capacity Scheduler

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

Enable Capacity Scheduler

Notice:

After you turn on Enable Resource Queue, you can no longer configure data 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.

1. Log on to the Alibaba Cloud E-MapReduce console.
2. In the top navigation bar, select the region where your cluster resides. Select the resource group as required. By default, all resources of the account appear.
3. Click the **Cluster Management** tab.
4. On the **Cluster Management** page that appears, find the target cluster and click **Details** in the Actions column.
5. In the left-side navigation pane of the Cluster Overview page that appears, click **Cluster Resources**.
6. On the **Cluster Resources** page that appears, turn on **Enable Resource Queue**.
7. Click **Save**.

   Capacity Scheduler is enabled.

Configure Capacity Scheduler

After Enable Resource Queue is turned on, follow these steps to configure Capacity Scheduler:

1. On the **Cluster Resources** page, click **Queue Settings** in the upper-right corner.
2. On the **Queue Settings** tab, find the target queue and click **Edit** in the Actions column.

   Select **Create Child Queue** from the **More** drop-down list in the Actions column to create a child 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 under the root queue.

   Notice:

   - The total value of Capacity for all child queues at the same level under the same parent queue must be 100. For example, two child queues default and department are available under the root queue. The total value of Capacity set for these two child queues must be 100. Child queues market and dev are available under the department queue. The total value of Capacity for these two child queues must also be 100.
   - If you do not specify a queue during the running of applications, jobs are submitted to the default queue.

- If you create level-2 queues under root, you can click **Deploy** without the need to restart ResourceManager.
- If you create or set level-3 queues under root, you need to restart ResourceManager.
- You must restart ResourceManager after you modify the name of a queue.

Switch the scheduler type

To switch the scheduler type after Enable Resource Queue is turned on, follow these steps:

1. On the **Cluster Resources** page, click **Select Scheduler** in the upper-right corner.
2. Select the required scheduler.
3. Click **Save**.
4. Select **RestartResourceManager** from the **Actions** drop-down list in the upper-right corner.
5. In the **Cluster Activities** dialog box that appears, specify related parameters and click **OK**.
6. In the **Confirm** dialog box that appears, click **OK**.

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

Submit a job

- If you do not specify a queue during the running of applications, jobs are submitted to the default queue.
- You must specify a child queue. Tasks cannot be submitted to the parent queue.
- You must use **mapreduce.job.queuename** 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`
```

2.15.3 Use Fair Scheduler

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

Enable Fair Scheduler

 **Notice:**

After you turn on Enable Resource Queue, you can no longer configure data 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.

1. Log on to the Alibaba Cloud E-MapReduce console.
2. In the top navigation bar, select the region where your cluster resides. Select the resource group as required. By default, all resources of the account appear.
3. Click the **Cluster Management** tab.
4. On the **Cluster Management** page that appears, find the target cluster and click **Details** in the Actions column.
5. In the left-side navigation pane of the Cluster Overview page that appears, click **Cluster Resources**.
6. On the **Cluster Resources** page that appears, turn on **Enable Resource Queue**.
7. Select **Fair Scheduler**.
8. Click **Save**.

Configure Fair Scheduler

After Enable Resource Queue is turned on, follow these steps to configure Fair Scheduler:

1. On the **Cluster Resources** page, click **Queue Settings** in the upper-right corner.
2. On the **Queue Settings** tab, find the target queue and click **Edit** in the Actions column.

   Select **Create Child Queue** from the **More** drop-down list in the Actions column to create a child 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 under the root queue.

   Notice:
   Parameters of the parent queue have higher priority than those of the child queue 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 mandatory. A queue name cannot contain periods (.).

- **Weight** is mandatory. The scheduler evenly allocates resources to queues if the weight threshold is reached. Weights take effect on level-2, level-3, and deeper-level queues. For example, a parent queue has a child queue with a weight of 2 and another with a weight of 1. If both of the two child queues have running tasks and the resource allocation ratio reaches 2:1, resources are evenly allocated.
- **Maximum Resources** indicates the maximum number of vCores and the maximum memory space that can be allocated to queues. The value of this parameter must be greater than that of **Minimum Resources** but less than the resource scale that the YARN service can provide. If the value of Maximum Resources is greater than the resource scale, the resource scale takes effect for queues. Assume that the number of vCores that the YARN service can provide is 16, but the value of vCore for Maximum Resources is 20. The 16 vCores are allocated to the queues.
- If you do not specify a queue during the running of applications, 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 validate the modification. After you delete a level-2 queue, select **RestartResourceManager** from the **Actions** drop-down list in the upper-right corner to validate the modification.

Switch the scheduler type

To switch the scheduler type after Enable Resource Queue is turned on, follow these steps:

1. On the **Cluster Resources** page, click **Select Scheduler** in the upper-right corner.
2. Select the required scheduler.
3. Click **Save**.
4. Select **RestartResourceManager** from the **Actions** drop-down list in the upper-right corner.
5. In the **Cluster Activities** dialog box that appears, specify related parameters and click **OK**.
6. In the **Confirm** dialog box that appears, click **OK**.

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

Submit a job

You must use **mapreduce.job.queuename** 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`
```

2.16 Auto Scaling

2.16.1 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. For more information about preemptible instances, see Overview.

2.16.2 Enable and configure auto scaling

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

Prerequisites

The E-MapReduce cluster is a Hadoop cluster. Currently, only Hadoop clusters support auto scaling.

Procedure

1. Log on to the Alibaba Cloud E-MapReduce console.
2. In the top navigation bar, select the region where your cluster resides. Select the resource group as required. By default, all resources of the account appear.
3. Click the **Cluster Management** tab.
4. On the **Cluster Management** page that appears, find the target cluster and click **Details** in the Actions column.
5. In the left-side navigation pane of the Cluster Overview page that appears, click **Auto Scaling**.
6. Click **Enable Auto Scaling**.
7. On the **Auto Scaling Settings** page that appears, configure the scaling policy.

   The following table describes the rule types of the scaling policy.

Type	Description
Scale By Time	Computing resources are scaled based on time. We recommend that you select this option if your business needs change periodically. For more information, see Configure AutoScaling by time.
Spot Instance	A preemptible instance is more cost-effective than a pay-as-you-go instance. However, Alibaba Cloud may release your preemptible instances at any time based on changes in the supply and demand of resources or the market price. For more information, see Create a preemptible instance.

8. After the scaling policy is configured, click **Save**.

   After auto scaling is configured and triggered, you can view the scaling history and the result of each scaling activity. For more information, see View the scaling history.

2.16.3 Configure AutoScaling by time

If the computing capability of a Hadoop cluster sees significant peaks and troughs over a specified period of time, you can set a fixed time frame within which a certain number of task nodes supplement the computing capability. This not only ensures that jobs are completed, but it also saves you money.

The expansion nodes are billed in Pay-As-You-Go mode. However, for the same computing capability, the price ratio between Pay-As-You-Go and Subscription modes is around 3:1 . Therefore, it isnecessary to design a ratio for both modesbased on the time needed. For example, if there are 8 peak business hours a day, the price for Pay-As-You-Go is essentiall

y the same as that for Subscription. If peak hours last longer than 8 hours, the Subscription mode is more cost-effective than Pay-As-You-Go.

Configure the number of scaling instances

- Maximum number of nodes: The maximum number of nodes that can be expanded . Once this number is reached, even if the AutoScaling rule is met, expansion and contraction will stop. Currently, you can set up to 1,000 task nodes.
- Minimum number of nodes: The minimum number of nodes that can be expanded. If the number of task nodes set in the AutoScaling rule is less than this minimum number of nodes, the cluster scalesbased on the minimum number of nodes at the first execution.

  For example, if the AutoScaling rule is set to expand 1 node at 00:00:00 every day and the minimum number of nodes is 3, then the system expands 3 nodes at the 00:00:00 on the first day.

Configure scaling rules

AutoScaling rules include expansion rules and contraction rules. When AutoScaling is disabled, all rules are cleared. If it is enabled again, the scaling rules need to be reconfigured.

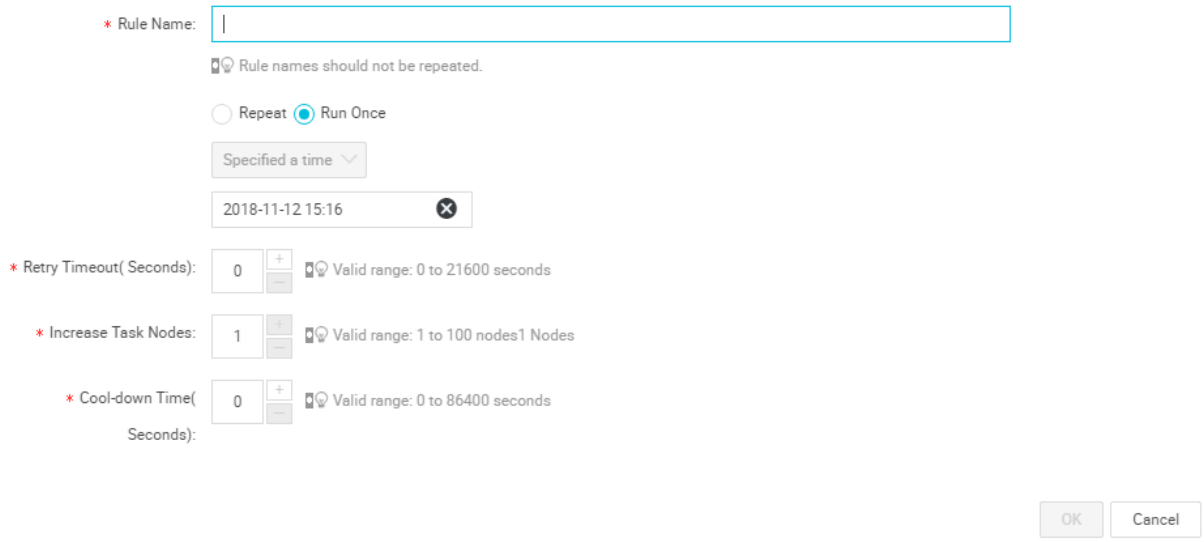

- **Rule Name**: In a cluster, the scaling rule names (including expansion rules and contraction rules) cannot be repeated.

- Execution cycle:
  - **Run Once**: The cluster performs a scaling operation at a specified time.
  - **Repeat**: You can choose to perform a scaling operation at a specific time every day, every week, or every month.
- **Retry Timeout**: When the specified time is reached, scaling cannot be performed. By setting the retry timeout, the system can detect that scaling can be performed every 30 seconds in the time range. Once the condition is met, scaling is performed. The range is 0 to 21600 seconds.

  It is assumed that if expansion A needs to be performed in the specified time period, but expansion B is being performed or the cluster is in **cool-down time**, expansion A cannot be performed. During the retry timeout you set, the system detectsthat expansion A can be performed every 30 seconds. Once the conditions are met, the cluster immediately performs scaling.
- **Increase Task Nodes**: The number of task nodes to be increased or decreased by the cluster each time the rule is triggered.
- **Cool-down Time**: The interval between a scaling operation being completed and the same operation being performed again. Scaling operationsare not performed during cool-down time.

Configure scaling instance specifications

When Auto Scaling is enabled, you can specify the hardware specifications for the scaling nodes. The specifications cannot be modified after the configuration is saved. If you need to modify them, you can disable AutoScaling and then enable it again.

- When you select specifications for vCPU and memory, the system automatically matches the instances that meet the criteria and displays them in the instance list below. For the cluster to be able scale according to the selected instance specifications, you need to add an optional instance to the list on the right.
- To avoid scaling failures caused by insufficient ECS instance storage, you can choose up to 3 ECS instance types.
- Regardless of whether you choose an efficient cloud disk or a SSD cloud disk, the data disk is set to a minimum of 40G.

2.16.4 Create a preemptible instance

Preemptible instances in E-MapReduce (EMR) are suitable for scenarios where there are no requirements on the execution of big data jobs, but the cost of computing resources is a

major concern. With the auto scaling feature, you can purchase preemptible instances to increase the computing resources of clusters.

Enable auto scaling

To enable auto scaling and configure scaling rules, follow these steps:

1. Log on to the Alibaba Cloud E-MapReduce console.
2. Click the **Cluster Management** tab.
3. Find the target cluster and click **Details** in the Actions column.
4. In the left-side navigation pane of the Cluster Overview page that appears, click **Auto Scaling**.
5. On the Auto Scaling page that appears, click **Enable Auto Scaling** in the upper-right corner.
6. Configure scaling rules. For more information, see Configure auto scaling rules by time.
7. In the Scaling Settings section, select **Spot Instance** for Billing Method.

Configure preemptible instances

Note:

A preemptible instance is more cost-effective than a pay-as-you-go instance, but Alibaba Cloud may release your preemptible instances at any time based on changes in the supply and demand of resources or market exchange values.

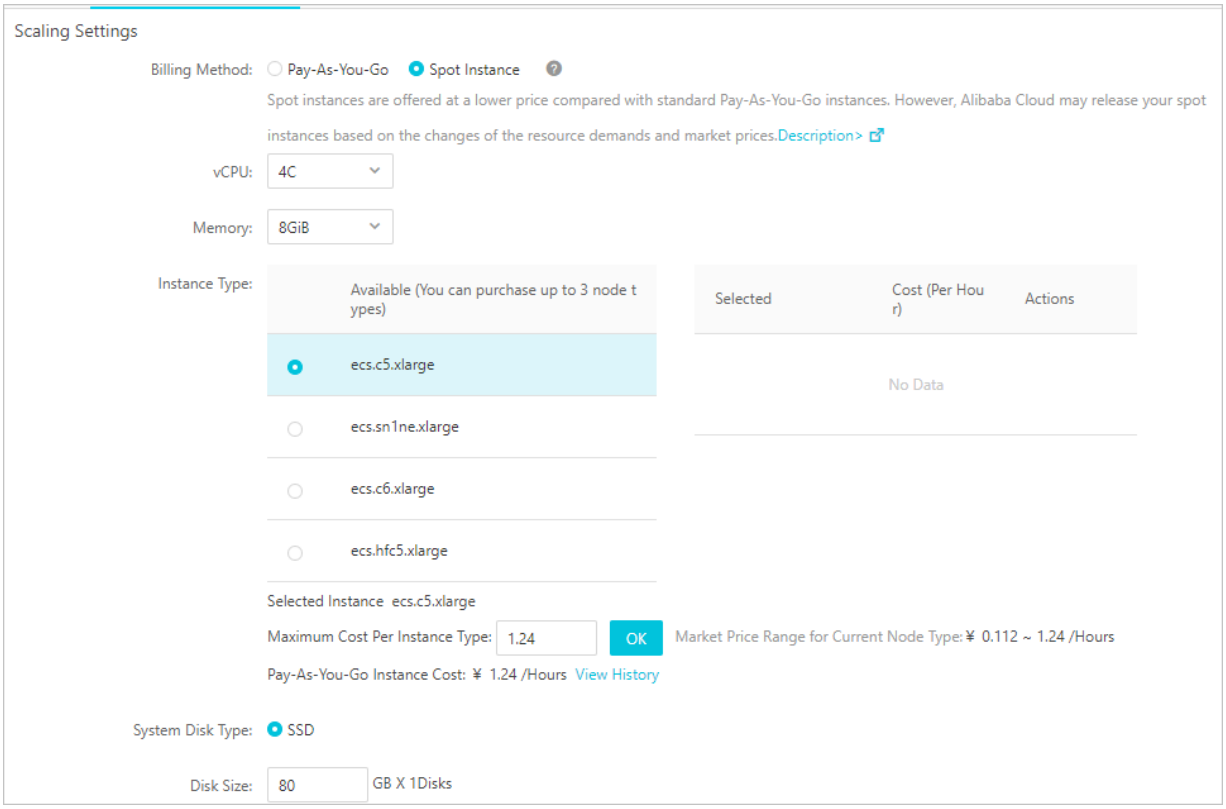

To configure preemptible instances, follow these steps:

1. Specify vCPU and Memory.
2. Select instance types. EMR automatically selects all the instance types that meet specifications. You can select up to three instance types to ensure that you can purchase preemptible instances as often as you require.
3. Specify Maximum Cost Per Instance Type for each instance type and click **OK**. The selected instance types appear in the Selected column. If you want to change the price of a selected instance type, select the instance type in the Available list and modify Maximum Cost Per Instance Type (which is calculated on an hourly basis). Your instance starts to run if your bid is higher than the current market price. The system charges you the market price for the final instance types you use.
4. Adjust the data disk size based on your needs. The system disk is used to deploy the operating system and basic EMR services. You do not need to adjust the disk size.
5. Click **Save**. The final configuration price is calculated as follows: Maximum instance bid + Price of the system disk + Total price of data disks.

For more information about preemptible instances, see #unique_63.

2.16.5 Disable auto scaling

This topic describes how to disable auto scaling when your business volume stabilizes.

Context

After auto scaling is disabled, all task nodes added by using auto scaling are released. However, the data of core nodes stored in HDFS is not affected. If a cluster has insufficient computing resources after you disable auto scaling, we recommend that you scale out the cluster to increase computing resources. For more information, see Scale out a cluster.

Procedure

1. Log on to the Alibaba Cloud E-MapReduce console.
2. Click the **Cluster Management** tab.
3. On the **Cluster Management** page that appears, find the target cluster and click **Details** in the Actions column.
4. In the left-side navigation pane of the Cluster Overview page that appears, click **Auto Scaling**.
5. On the **Auto Scaling** page, click **Disable Auto Scaling** in the upper-right corner.
6. In the **Auto Scaling Disabled** message that appears, click **OK**.

2.16.6 View the scaling history

This topic describes how to view the scaling history after scaling activities are triggered.

Procedure

1. Log on to the Alibaba Cloud E-MapReduce console.
2. In the top navigation bar, select the region where your cluster resides. Select the resource group as required. By default, all resources of the account appear.
3. Click the **Cluster Management** tab.
4. On the **Cluster Management** page that appears, find the target cluster and click **Details** in the Actions column.
5. In the left-side navigation pane of the Cluster Overview page that appears, click **Auto Scaling**.

6. Click **History**.

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

   - **executing**: The scaling activity is in progress.
   - **success**: All instances specified in the scaling rule are added or removed in the scaling activity.
   - **partOfSuccess**: In the scaling activity, some instances specified in the scaling rule are added or removed. The other instances fail to be added or removed because of insufficient disk quota or Elastic Compute Service (ECS) resources.
   - **failure**: No instance specified in the scaling rule is added or removed in the scaling activity.
   - **reject**: The scaling activity is rejected because the number of instances will exceed the upper or lower limit of the number of instances in the scaling group after scaling.

2.17 Configure VPCs

Virtual Private Cloud (VPC) helps you build an isolated network environment. You can customize Classless Inter-Domain Routing (CIDR) blocks, create one or more subnets, and configure route tables and gateways for VPCs. You can create E-MapReduce (EMR) clusters in different VPCs and use Express Connect to enable the VPCs to communicate with each other.

For more information about VPC, see What is a VPC. In addition, VPCs and Internet data centers (IDCs) can communicate with each other by using Express Connect. VPCs in different regions or under different accounts can communicate with each other by using Cloud Enterprise Network (CEN).

Create an EMR cluster in a VPC

You can create an EMR cluster in a classic network or a VPC. To create a cluster in a VPC, you must configure the following information:

- VPC: Select the VPC to which the EMR cluster belongs. If no VPCs are created, create one in the VPC console. You can create a maximum of two VPCs under an account. To create more than two VPCs, submit a ticket.
- VSwitch: Select the VSwitch over which Elastic Compute Service (ECS) instances in the EMR cluster can communicate with each other. If no VSwitches are created, log on to

the VPC console and create a VSwitch on the **VSwitches** page. A VSwitch belongs to a specific zone. You must make sure that the VSwitch is in the zone where the EMR cluster resides.

- Security group: Select the existing security group to which the EMR cluster belongs. For security purposes, only the security groups created in EMR are available in the drop-down list. To create a security group, you can directly enter a name in the Security Group Name field.

Connect EMR clusters that belong to different VPCs under the same account

This example creates a Hive cluster and an HBase cluster that belong to different VPCs and uses CEN to enable the Hive cluster to access the HBase cluster.

1. Create clusters.

   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.

   Create a CEN instance to connect the two clusters. For more information, see #unique_68. Make sure that the CEN instance resides in the same region as the clusters.

3. Use SSH to log on to the HBase cluster and run the following command in the HBase shell:

   ```
   hbase(main):001:0> create 'testfromHbase','cf'
   ```

4. Use SSH to log on to the Hive cluster and follow these steps:

   a. Append the following line to the hosts file:

   ```
   $zk_ip emr-cluster //$zk_ip indicates the IP address of the ZooKeeper node in the HBase cluster.
   ```

   b. Run the following commands in the Hive shell to connect to the HBase cluster:

   ```
   hive> set hbase.zookeeper.quorum=172. *. *.111,172. *. *.112,172. *. *.113;
   hive> CREATE EXTERNAL TABLE IF NOT EXISTS testfromHive (rowkey STRING, pageviews Int, bytes STRING) STORED BY 'org.apache.hadoop.hive.hbase.HBaseStorageHandler' WITH SERDEPROPERTIES ('hbase.columns.mapping' = ':key, cf:c1,cf:c2') TBLPROPERTIES ('hbase.table.name' = 'testfromHbase');
   ```

   c. If the **java.net.SocketTimeoutException** error appears, which indicates that the connection failed, add rules in 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.

Authorization policy	Protocol type	Port range	Authorization type	Authorization object
Allow	TCP	2181/2181	Address field access	192.168.1.0/16
Allow	TCP	22/22	Address field access	0.0.0.0/0
Allow	TCP	16000/16000	Address field access	192.168.1.0/16
Allow	TCP	16020/16020	Address field access	192.168.1.0/16

By default, only port 22 is enabled for security groups created in EMR. However, the Hive cluster cannot access the HBase cluster at this port. Therefore, you must enable all ports required for the Hive cluster to access the HBase cluster.

2.18 MetaService

E-MapReduce supports MetaService. MetaService allows you to access Alibaba Cloud resources from E-MapReduce clusters without providing the AccessKey.

Default application role

You can grant the default application role AliyunEmrEcsDefaultRole to E-MapReduce when creating an E-MapReduce cluster. Then, your E-MapReduce jobs can access Alibaba Cloud resources without explicitly providing the AccessKey. By default, the following permissions are granted to AliyunEmrEcsDefaultRole:

```
{
  "Version": "1",
  "Statement": [
    {
      "Action": [
        "oss:GetObject",
        "oss:ListObjects",
        "oss:PutObject",
        "oss:DeleteObject",
        "oss:ListBuckets",
        "oss:AbortMultipartUpload"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }
  ]
}
```

That is, MetaService-based jobs can access only Object Storage Service (OSS) data by default. To allow MetaService-based jobs to access other Alibaba Cloud resources, such as Log Service data, you must grant the required permissions to AliyunEmrEcsDefaultRole. You

can grant the default application role to E-MapReduce and configure the permissions of the role in the Resource Access Management (RAM) console.

> **Notice:**
> Currently, MetaService allows you to access only the data of OSS, Log Service, and Message Service (MNS) without providing the AccessKey. Edit or delete the default application role with caution. If the role is edited or deleted by mistake, you may fail to create clusters or run jobs.

Custom application role

The default application role can meet most business requirements. You can directly use it or edit it as required. E-MapReduce also allows you to use a custom application role. That is, when creating a cluster, you can use the default application role or select a custom application role. For more information about how to create a role and grant it to a service, see RAM documentation.

Access MetaService

MetaService is a Hypertext Transfer Protocol (HTTP). You can access MetaService to obtain metadata information. For example, you can obtain the region where the current cluster resides by running the **curl http://localhost:10011/cluster-region** command.

Currently, 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

Use MetaService

Your jobs can use MetaService to access Alibaba Cloud resources without providing the AccessKey, which brings the following benefits:

- Reduces the risk of an AccessKey leakage. The use of a RAM role can minimize the security risk. You can grant only the required permissions to a role. This minimizes the permissions that are granted.

- Improves user experience. MetaService is especially useful when you access OSS resources because it shortens the OSS path that you need to enter.

The following examples show how to use MetaService.

```
I. Hadoop
   Previously, we used: hadoop fs -ls oss://ZaH******As1s:Ba23N**************sdaBj2@
bucket.oss-cn-hangzhou-internal.aliyuncs.com/a/b/c
   Now, we use: hadoop fs -ls oss://bucket/a/b/c
II. Hive
   Previously, we 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';
   Now, we use:
      CREATE EXTERNAL TABLE test_table(id INT, name string)
      ROW FORMAT DELIMITED
      FIELDS TERMINATED BY '/t'
      LOCATION 'oss://bucket/a/b/c';
III. Spark
   Previously, we used: val data = sc.textFile("oss://ZaH******As1s:Ba23N**************
sdaBj2@bucket.oss-cn-hangzhou-internal.aliyuncs.com/a/b/c")
   Now, we use: val data = sc.textFile("oss://bucket/a/b/c")
```

3 Modify configurations

3.1 Upgrade node configuration

You may need to upgrade the specification of the CPU and memory of the master node when using an EMR cluster.

Note:

- Only Subscription EMR clusters support configuration upgrades.
- Local disk-based instances do not support configuration upgrades. You can only improve cluster performance by increasing the number of nodes.
- Non-local disk-based instances support configuration upgrades and do not support specification downgrades.

Procedure

1. On the **Cluster Management** page, click **View Details** in the Actions column for the cluster that you want to upgrade node specification for.
2. On the Cluster Overview page, click **Resource Allocation**. From the drop-down list, select **Upgrade**.
3. Modify the node configuration as needed.
4. Click **OK** and wait until the order is generated.
5. Pay the order.
6. Go to the Cluster Overview page, refresh the page, and confirm that the node configuration has been upgraded. For example, CPU: 4 cores and Memory: 16 GB. A message appears in the instance group section as shown in the following figure.

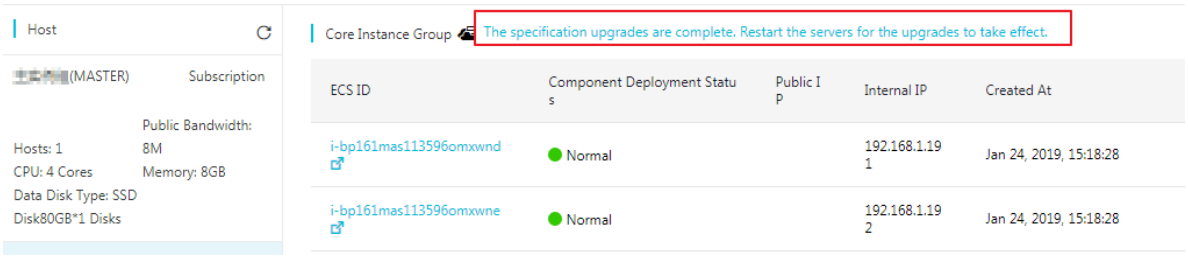

7. Click **The specification upgrades are complete. Restart the servers for the upgrades to take effect.** A dialog box appears as shown in the following figure.

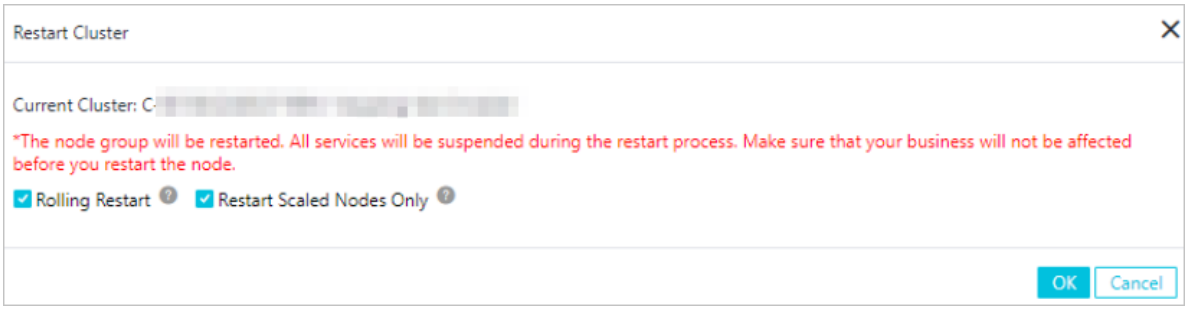

8. Click **OK**.

Notice:

- When you restart a cluster, the ECS instances of the cluster are restarted. During restart, the big data services deployed on ECS instances are not available. Make sure that the cluster restart does not impact your business.
- Rolling Start
  - If you select Rolling Start, ECS instances restart one by one after the previous ECS instance completes restarting and all big data services deployed on this instance are restored. Restarting an instance takes about five minutes.
  - If you clear Rolling Start, all ECS instances restart at the same time.
- Only Restart Updated Nodes
  - An updated node (a core node or the master node) refers to a node of which the configuration has been updated.
  - If you select Only Restart Updated Nodes, nodes that have not been upgraded do not restart. For example, if only the configurations of nodes in the core group are updated, then only the ECS instances in the core group are restarted. The ECS instances in the master group are not restarted.
  - If you clear Only Restart Updated Nodes, all nodes are restarted. That is to say, all instances of the cluster are restarted.

9. In the progress of restarting, the **Instance Group Restarting** prompt appears next to the instance group name as shown in the following figure.

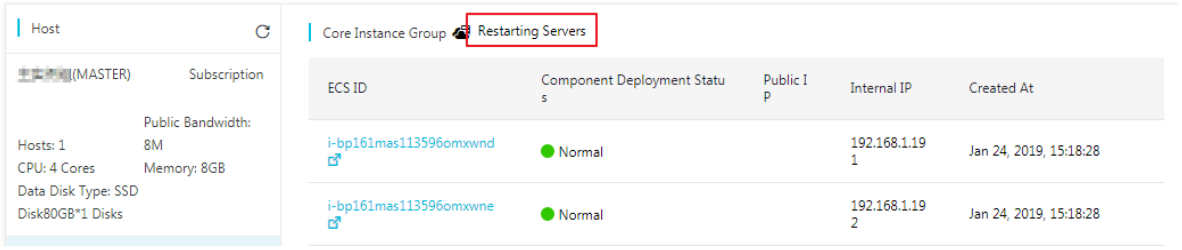

10. When the prompt disappears, the configurations are complete. You can log on to the cluster for verification.
11. If you only update the CPU configurations, you can skip this step. If you update the memory configurations, you need to modify the cluster service configurations for YARN to use new resources. If you update both, you need to modify both configurations. Procedure:

    a. On the **Clusters and Services** page, click **YARN**.

    b. Modify the CPU configurations

    Go to the Service page, click **Configure** and locate the **yarn.nodemanager.resource.cpu-vcores** configuration item. For computing-intensive scenarios, we recommend that you set the ratio to 1:1. For mixed scenarios, we recommend that you set the ratio to less than 1:2. For example, a 32 vCPU-based compute node that is intended for computing-intensive scenarios requires large amounts of configurations for regular expressions. You can set as follows. `yarn.nodemanager.resource.cpu-vcores=32`. If the compute node is intended for mixed scenarios, you can set the value of **yarn.nodemanager.resource.cpu-vcores** to a number between 32 and 64.

    c. Modify the memory configurations

    Go to the Service page, click the **Configure** tab and locate the **yarn.nodemanager.resource.memory-mb** configuration item. Set the value of the configuration item to 0.8 times the machine memory. The unit is megabytes. For example, if the memory

size is 32 GB, the setting is as follows. yarn.nodemanager.resource.memory-mb=26214

d. Click **Save** in the upper-right corner to save the configurations.

e. Click **Actions** and select **Configure All Components** from the drop-down list.

f. Click **History**. Wait until **Successful** appears in the Status column.

g. Click **Actions** and select **RESTART All Components** from the drop-down list.

h. Click **History**. When **Successful** appears in the Status column, you can uses new resources in YARN.

3.2 Expand disks

E-MapReduce (EMR) allows you to expand disks for EMR clusters. You can expand both system and data disks.

Overview

According to the version of EMR that you are using and the type of disks that you need to expand, choose one of the following methods to expand disks:

- Data disk
  - EMR V3.11.0 and later: You can expand data disks in the EMR console. The current EMR console does not support shrinking data disks.
  - Versions earlier than EMR V3.11.0: You can expand data disks in the Elastic Compute Service (ECS) console.
- System disk
  - All EMR versions: You can expand the system disk in the ECS console.

Expand data disks (EMR V3.11.0 and later)

The current EMR console only supports expanding data disks. For more information about how to expand the system disk, see Expand the system disk.

Notice:
Before you expand a data disk, make sure that your account has sufficient balance. Fees for data disk expansion are automatically deducted from your account. If your account does not have sufficient balance, the expansion process is terminated.

1. Log on to the Alibaba Cloud EMR console and then go to the **Cluster Management** page.

2. On the **Cluster Management** page, click the ID of the target cluster in the **Cluster ID/ Name** column. You can also perform this task on the **Overview** page.
3. In the left-side navigation pane, choose **Cluster Overview**, and then choose **Change Configuration** > **Disk Expansion** in the upper-right corner of the page.
4. In the **Disk Expansion** dialog box, set the parameters.

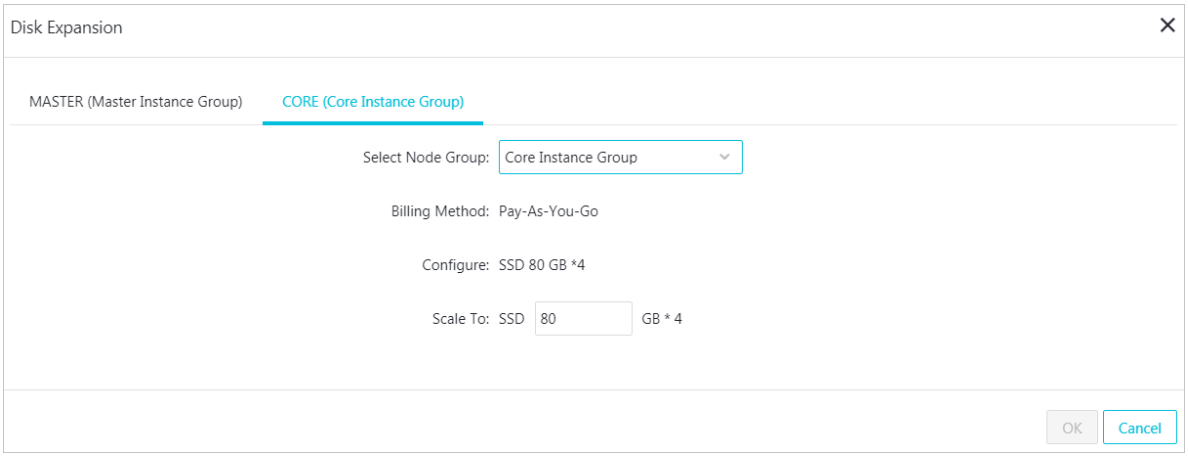

The required parameters are described as follows:

Parameter	Description
Instance Group	Instance groups that support disk expansion include: • **MASTER (Master instance group)** • **CORE (Core instance group)**
Select Node Group	Select a node group from the **Select Node Group** list.
Billing Method	Indicates the billing method that the selected EMR cluster is using.
Configure	The size of each data disk for the selected instance group .
Scale To	The size of the storage space that you want to expand each data disk to.

5. After you set the parameters, click **OK**.

   After you expand the disks, the **The disk is expanded. Restart the node group for the change to take effect.** message is displayed at the bottom of the **Cluster Overview** page for the instance group.

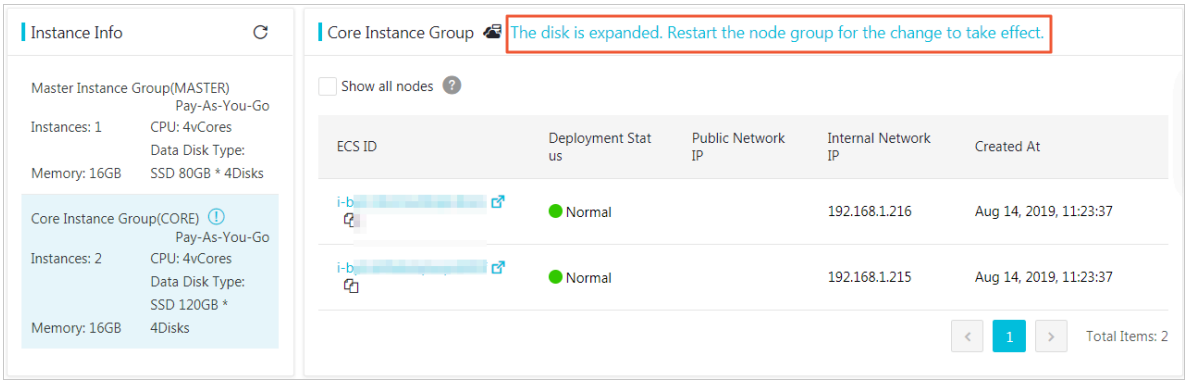

Notice:

You must restart the cluster for the disk expansion to take effect. This operation restarts the ECS instances in the cluster. During the restart process, the big data services are unavailable. Make sure that the restart operation does not adversely affect your businesses before you restart the cluster.

6. Click **The disk is expanded. Restart the node group for the change to take effect.** to configure the restart settings.

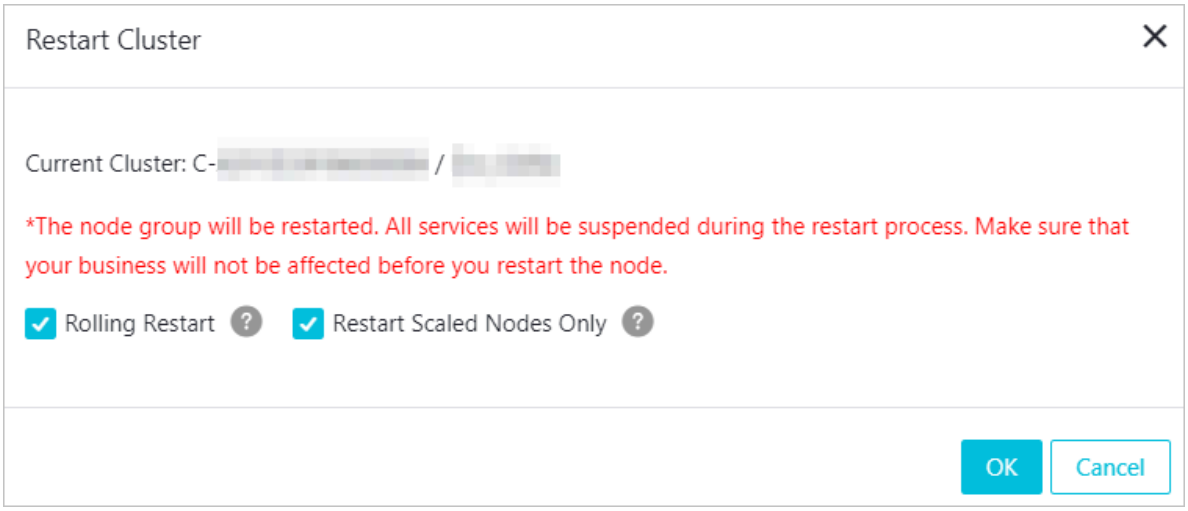

The cluster restart settings are described as follows:

Parameter	Description
Rolling Start	**Rolling restart**: • Select the **Rolling Restart** check box: The ECS instances in the cluster are restarted in sequence. EMR must complete restarting an ECS instance and recovering the big data services on the ECS instance before it can restart another ECS instance. It takes about five minutes to restart an ECS instance. • Clear the **Rolling Restart** check box: EMR restarts all ECS instances at the same time.
Restart Scaled Nodes Only	A scaled node is a node that has its disks or configuration upgraded, such as a master or core node. **Restart scaled nodes only**: • Select the **Restart Scaled Nodes Only** check box: Only scaled nodes are restarted. Other nodes are not restarted. For example, if you have expanded the disks of the nodes in a **core instance group**, then only the ECS instances in the **core instance group** are restarted. • Clear the **Restart Scaled Nodes Only** check box: All nodes are restarted. This means that all ECS instances in the cluster are restarted.

7. After you set the parameters, click **OK** to restart the cluster.

   During the restart process, the instance group displays **Restarting server group...**.
   If the node group is successfully restarted, the **Restarting server group...** message disappears. You can go to the Cluster Overview page to verify the disk expansion result.

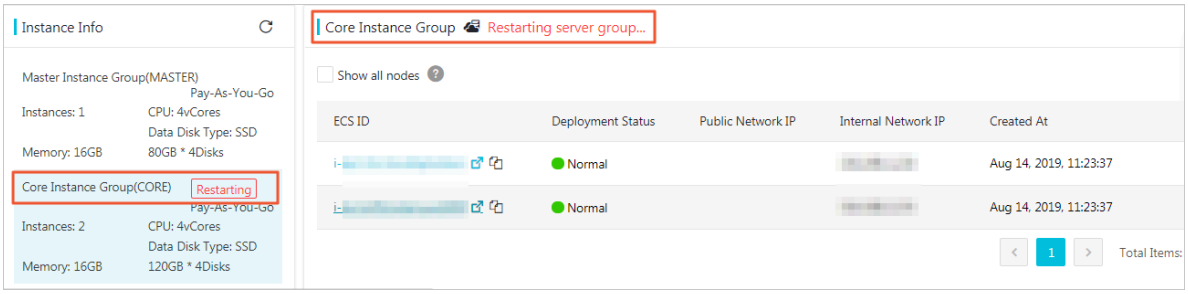

Expand data disks (Versions earlier than EMR V3.11.0)

> **Notice:**
> You must expand the disks of all nodes in a cluster to make sure that the nodes have the same disk size.

1. Log on to the Alibaba Cloud EMR console and then go to the **Cluster Management** page.
2. On the **Cluster Management** page, click the ID of the target cluster in the **Cluster ID/Name** column. You can also perform this task on the **Overview** page.
3. Select the instance group that you want to expand disks for in the **Instance Info** area at the bottom of the **Cluster Overview** page. Copy the **ECS ID** of an ECS instance shown on the right side.

   ECS ID example: **i-bp1bsithym5hh9h93xxx**.
4. Click the **ECS ID** of the ECS instance to log on to the ECS console.

   After you log on to the ECS console, information about the selected ECS instance is displayed.
5. In the left-side navigation pane, choose **Disks**. You can then expand the data disks for the ECS instance on the right side of the page. For more information, see #unique_72 or #unique_73.

   > **Note:**
   > Your cluster may not be able to meet all the requirements for expanding disks online. In this case, you can expand disks offline. Currently you cannot expand multiple disks at the same time. You can only expand one disk at a time.

6. After the disks are expanded, you have to expand the partitions and file systems of the disks. For more information, see #unique_74.

   Note:

   - During the partition and file system expansion process, if an error occurs when the system runs the **umount** command, disable the **YARN** and **HDFS** services on the EMR cluster to resolve this issue.
   - For the partition Disk1, the system may fail to run the **umount** command if the **ilogtail** service is writing log data at the same time. To resolve this issue, run the **sudo pgrep ilogtail | sudo xargs kill** command to disable the **ilogtail** service. After the partitions and file systems are expanded, restart the instance and recover the **ilogtail** service.

7. After you complete the preceding tasks, log on to the ECS instance through SSH and then run the **df -h** command to verify the disk expansion result.
8. Reference the preceding steps to expand other data disks for the ECS instance.

Expand the system disk

Notice:
It is a complex procedure to expand the system disk. Do not expand the system disk unless it is necessary. Clusters other than HA clusters become unavailable during the system disk expansion process.

1. Log on to the Alibaba Cloud EMR console and then go to the **Cluster Management** page.
2. On the **Cluster Management** page, click the ID of the target cluster in the **Cluster ID/ Name** column. You can also perform this task on the **Overview** page.
3. Select the instance group that you want to expand the system disk for in the **Instance Info** area at the bottom of the **Cluster Overview** page. Copy the **ECS ID** of an ECS instance shown on the right side.

   ECS ID example: **i-bp1bsithym5hh9h93xxx**.
4. Click the **ECS ID** of the ECS instance to log on to the ECS console.

   After you log on to the ECS console, information about the selected ECS instance is displayed.

5. In the left-side navigation pane, select **Disks**. You can then expand the system disk for the ECS instance on the right side of the page. Each ECS instance has only one system disk. For more information, see #unique_72 or #unique_73.

   Note:
   Your cluster may not be able to meet all the requirements for expanding disks online. In this case, you can expand disks offline.

6. After the disks are expanded, you have to expand the partitions and file systems of the disks. For more information, see #unique_75.

   Note:
   - During the partition and file system expansion process, if an error occurs when the system runs the **umount** command, disable the **YARN** and **HDFS** services on the EMR cluster to resolve this issue.
   - For the partition Disk1, the system may fail to run the **umount** command if the **ilogtail** service is writing log data at the same time. To resolve this issue, run the **sudo pgrep ilogtail | sudo xargs kill** command to disable the **ilogtail** service. After the partitions and file systems are expanded, restart the instance and recover the **ilogtail** service.

7. After you complete the preceding tasks, log on to the ECS instance through SSH and then run the **df -h** command to verify the disk expansion result.

   Note:
   After the system disk is expanded, the ECS instance may have the following issues:
   - The ECS instance will make some changes to the system disk. This may change the /etc/hosts file on the ECS instance. You need to fix the file after you expand the system disk.
   - The SSH logon configuration becomes invalid. You need to manually fix the configuration. Other services on the ECS instance are not affected.

8. Reference the preceding steps to expand the system disk for other ECS instances.

3.3 Remove abnormal nodes

If an ECS node in an E-MapReduce cluster is abnormal, and if you do not need this node and want to remove it, you can do so.

To remove an abnormal node, follow these steps:

1. Log on to the Alibaba Cloud E-MapReduce console.
2. At the top of the page, click **Cluster Management**.
3. Click the ID link of the cluster that includes the node that you want to remove.
4. In the left-side navigation pane, click **Hosts**.
5. Find the instance that you want to remove, and in the action column, click **Remove**. ECS instances can only be removed when they are in a stopped or released status.
6. Click **OK** to remove the instance.

3.4 Multiple host groups

When you use an EMR Hadoop cluster, you need to specify the instance types for the instance groups based on the real scenarios. For example:

- You need some memory instances (CPU: Mem = 1 vCore: 8 GiB) to perform offline big data processing and some compute instances (CPU: Mem=1 vCore: 2 GiB) to train models .
- You want to assign certain departments with memory instances (CPU: Mem = 1 vCore: 8 GiB) and compute instances (CPU: Mem=1 vCore: 2 GiB) based on their requirements.

You can use Task instances to create multiple host groups and select different configurations for the host groups to meet the preceding requirements. Procedure:

1. Log on to the Alibaba Cloud E-MapReduce console. Click **Cluster Management** to go to the cluster management page.
2. Click **Clusters and Services** in the Actions column for the cluster that you want to create host groups for.
3. In the left-side navigation pane, click **Cluster Overview**.
4. On the Cluster Overview page, click **Resource Allocation**. From the Resource Allocation drop-down list, select **Scale Up/Out**.
5. Click the TASK (Task Instance Group) tab, click **Add Host Group**.

6. In the Host Group Name field, enter a name for the host group. Select the specifications as needed. A host group name is required to be unique for each cluster.

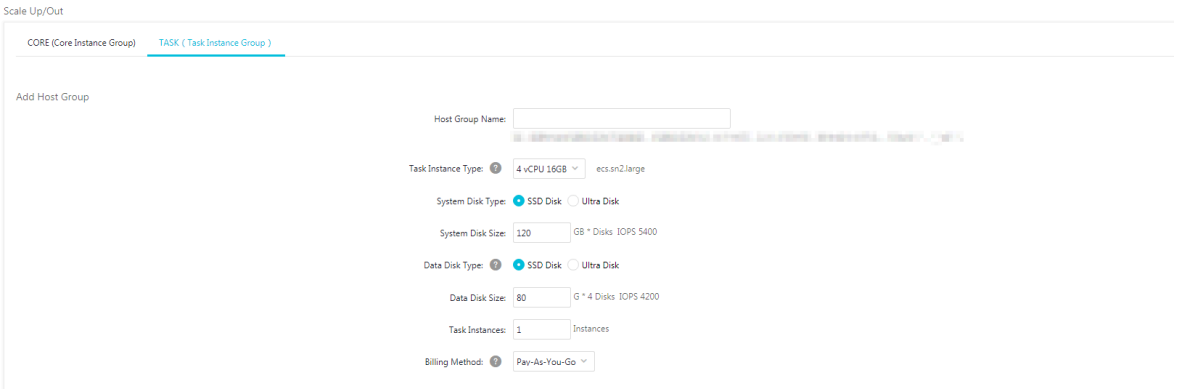

Notice:

You can create a maximum of 10 host groups to categorize the task instances of an EMR Hadoop cluster.

4 Third-party software

4.1 Bootstrap actions

Bootstrap actions are used to run a customized script before the cluster starts Hadoop. The customized script is used to install your required third-party software or change the cluster operating environment.

Function

With bootstrap actions, you can install and perform a variety of things to your cluster that are not currently supported by clusters, such as:

- Install software with Yum.
- Directly download open software from a public network.
- Read your data from OSS.
- Install and operate a service, such as Flink or Impala. (The script for this is more complex .)

We strongly recommend that you test the bootstrap action with a Pay-As-You-Go cluster and create a Subscription cluster only after the test is successful.

Procedure

1. Log on to the Alibaba Cloud E-MapReduce console.
2. Select a region. The cluster associated with this region is listed.
3. Click **Create Cluster** to enter the cluster creation page.
4. At the bottom of the basic configuration page, click **Add** to enter the operation page.
5. Enter the configuration items.

You can add up to 16 bootstrap actions to be performed during cluster initialization in a specified sequence. By default, your customized script is run with the root account. Use **su hadoop** in the script to switch to a Hadoop account.

It is possible for a bootstrap action to fail, but this does not affect the creation of a cluster. After a cluster is created successfully, you can view all abnormalities in the bootstrap/ software configuration column in the cluster details page. In the event of an abnormality, you can log on to all nodes to view the operation logs in the /var/log/bootstrap-actions directory.

Bootstrap action type

Bootstrap actions are categorized into customized bootstrap actions and operating-condition bootstrap action. The main difference is that the operating-condition bootstrap action can only operate your operation in nodes that meet the requirements.

Customized bootstrap action

For customized bootstrap actions, the position of the bootstrap action name and the execution script in OSS must be specified and, where necessary, optional parameters filled in. During cluster initialization, all nodes download the specified OSS scripts and runs them directly or after adding the optional parameters.

You can specify the files that need to be downloaded from OSS in the script. In the following example, the oss://yourbucket/myfile.tar.gz file is downloaded locally and extracted to the /yourdir directory:

```
#! #!/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>
```

The osscmd is preinstalled on the node and can be invoked directly to download the file.

Notice:
OSS address host has an intranet address, an Internet address, and a VPC network address. If you use a classic network, you need to specify the intranet address. For Hangzhou, this is oss-cn-hangzhou-internal.aliyuncs.com. If you use a VPC network, you need to specify the domain name that the VPC intranet can access. For Hangzhou, this is vpc100-oss-cn-hangzhou.aliyuncs.com.

Bootstrap actions can install additional system software packages through Yum. The following example shows the installation of ld-linux.so. 2:

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

Operating-condition bootstrap action

The execution script of an operating-condition bootstrap action is predefined. This means that you only need to specify its name and optional parameters. Operating-condition bootstrap actions must provide optional parameters, including operating conditions and commands (separated by spaces). The operating conditions support

instance.isMaster=true/false, which specifies that it only operates on master or non-master nodes. The following example shows that the optional parameters of an operating-condition bootstrap action are specified to only create a directory on a master node:

```
instance.isMaster=true mkdir -p /tmp/abc
```

If you need to specify multiple operating commands, you can divide them by using a semicolon (;). For example: `instance.isMaster=true mkdir -p /tmp/abc;mkdir -p /tmp/def`.

4.2 Software configuration

Hadoop, Hive, and Pig require large amounts of configurations. You can use the software configuration feature to modify the software configurations. Currently, you can only perform software configurations when creating a cluster. For example, you want to increase the number of HDFS server handler threads (dfs.namenode.handler.count) from 10 to 50. Or you want to decrease the HDFS block size from 128 MB (default size) to 64 MB.

Procedure

1. Log on to the Alibaba Cloud E-MapReduce console.
2. In the navigation bar, select a region for creating a cluster.
3. Click **Create Cluster** to go to the Custom Purchase page.
4. In the Software Configuration step, you can view all services and the corresponding versions. You can turn on **Enable Custom Setting** and enter JSON text in the editor to add configurations or to overwrite the default configurations. A JSON file example is shown as follows.

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

Note:

- The argument of FileName does not include the file extension.

- The argument of ServiceName is required to be capitalized.

ConfigKey refers to the name of a configuration item and ConfigValue refers to the value of the configuration item.

The configuration files of the services are shown as follows.

- Hadoop

  Filename:

  - core-site.xml
  - log4j.properties
  - Hdfs-site.xml
  - mapred-site.xml
  - yarn-site.xml
  - httpsfs-site.xml
  - capacity-scheduler.xml
  - hadoop-env.sh
  - httpfs-env.sh
  - mapred-env.sh
  - yarn-env.sh
- Pig

  Filename:

  - pig.properties
  - log4j.properties
- Hive

  Filename:

  - hive-env.sh
  - hive-site.xml
  - hive-exec-log4j.properties
  - hive-log4j.properties

5. After the settings are complete, click **Next: Hardware Configuration**.

4.3 Configure parameters

EMR supports configuring parameters for components including HDFS, Yarn, Spark, Kafka, and Druid.

Configure parameters for existing components

Perform the following steps to modify the parameters for existing components.

1. Log on to the Alibaba Cloud E-MapReduce console.
2. Click the **Cluster Management** tab to jump to the Cluster Management page.
3. Click a cluster ID to jump to the Clusters and Services page as needed.
4. In the Services list, click the component you want to modify parameters. For example, **HDFS**.

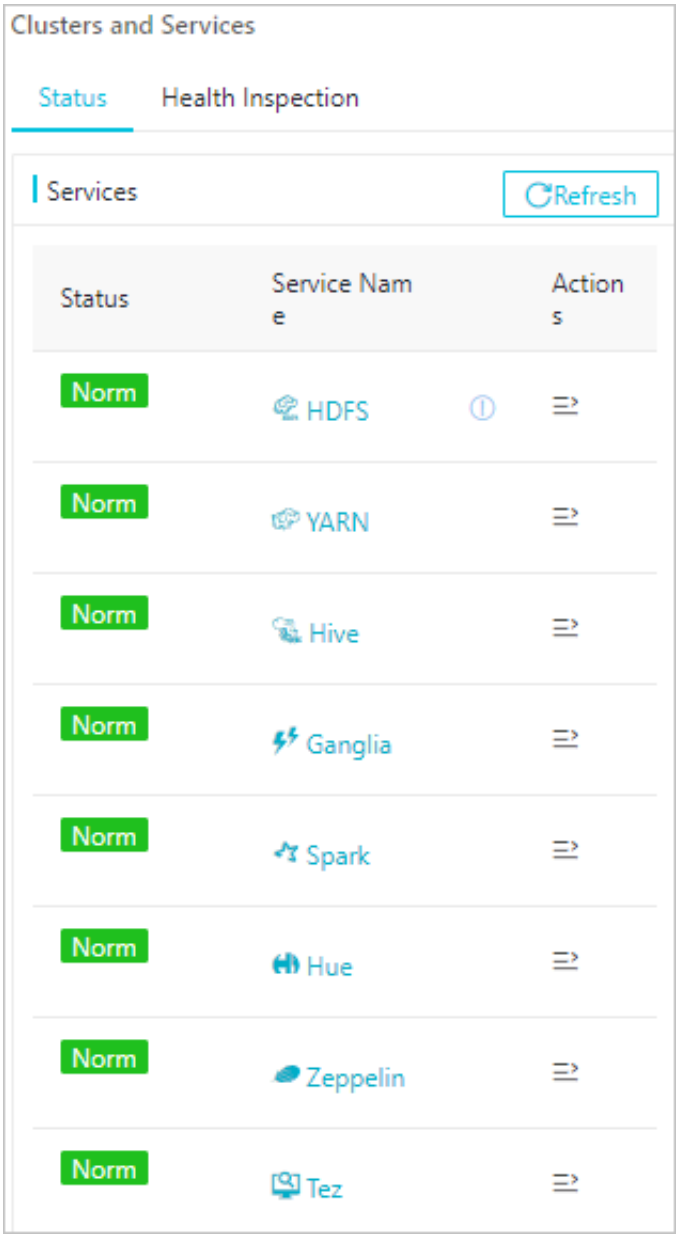

5. Click the **Configurations** tab to jump to the Service Configuration page.

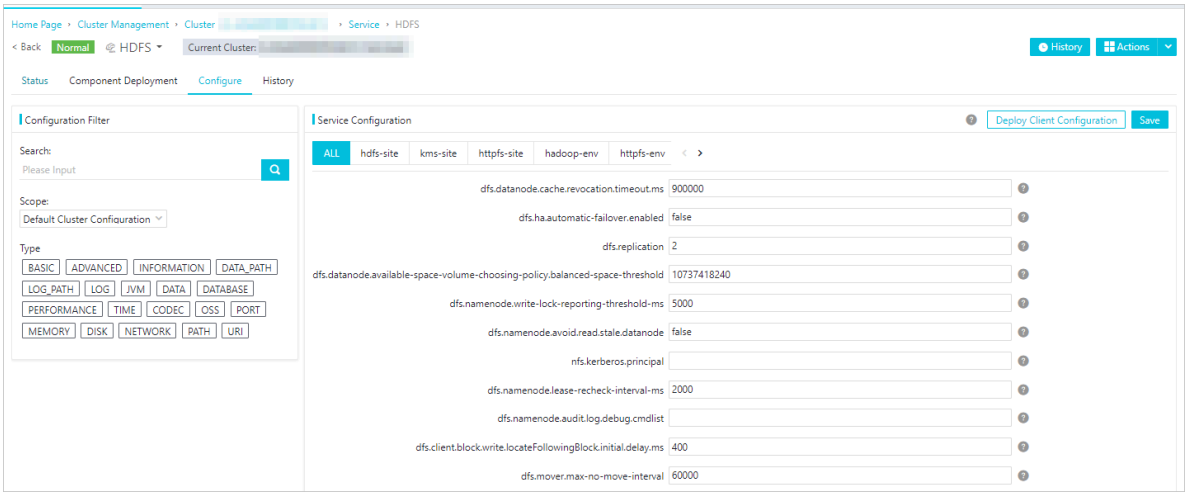

6. You can select configuration types and configuration files in the **Quick Configuration** area.
7. Locate the parameter you want to modify and set the value of the parameter. If you cannot find the parameter you want to modify, click **Custom Configuration** to add a configuration item.

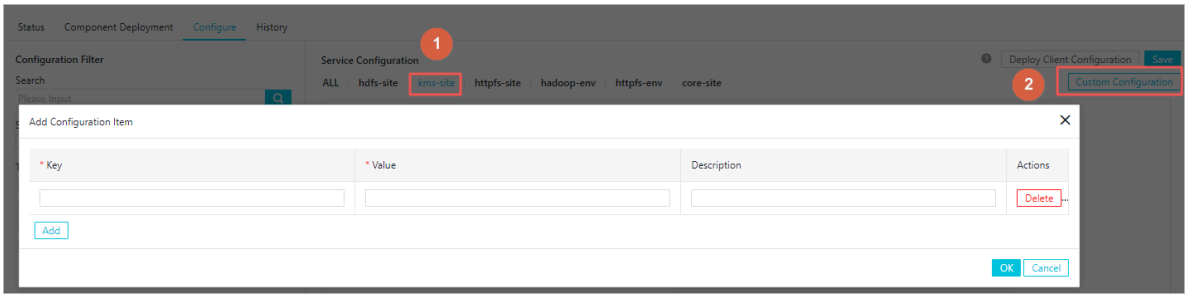

- Key: parameter name
- Value: parameter value
- Description: parameter description

8. Click **Save**.

Refer to the Deploy configurations section for making your parameter modification take effect.

Deploy configurations

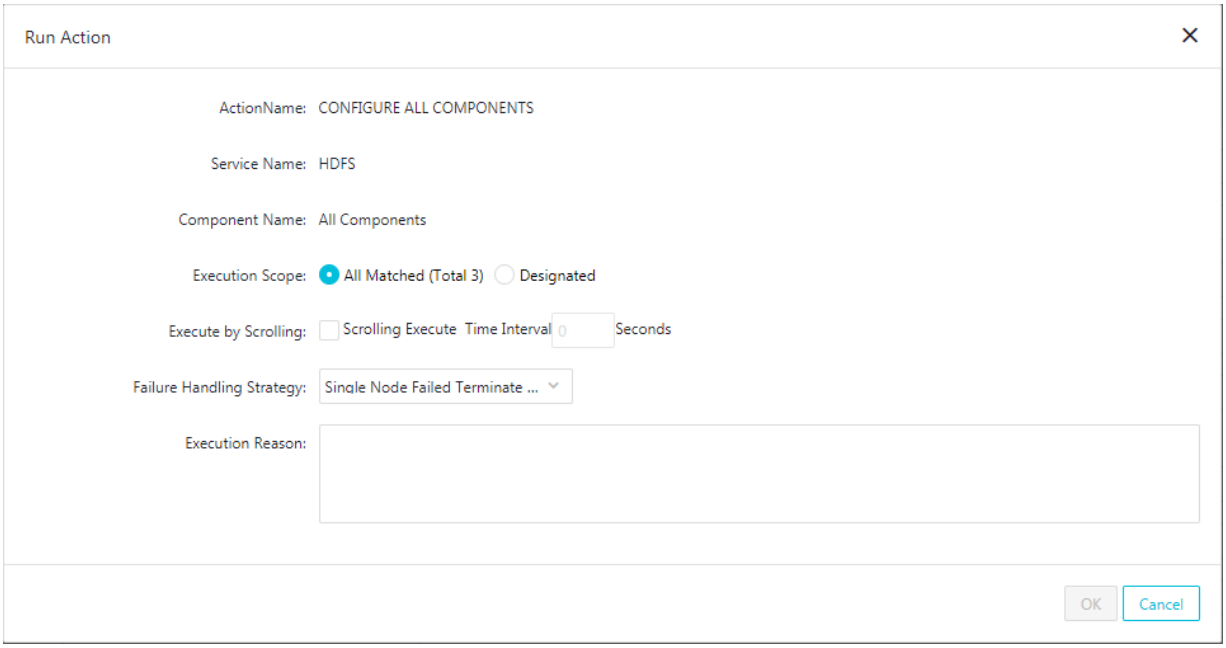

- Deploy client configurations
  1. After modifying a client configuration parameter, click **Deploy Client Config** and the **Run Action** dialog box is displayed.
  2. Click All Matched or Designated for the execution scope.
  3. Enter the execution reason.
  4. Click **OK** to make the modification take effect.
  5. You can click **View Operation Logs** to view the status and progress of the execution.
- Deploy server configurations
  1. After modifying a server configuration parameter, restart the corresponding service.
  2. Click **Deploy Client Config** and the **Run Action** dialog box pops up.
  3. Click the All Matched radio button or the Designated radio button for the execution scope.
  4. Select the Scrolling Execute check box if you need to perform a rolling restart.
  5. Enter the execution reason.
  6. Click **OK** to start the rolling execution.
  7. You can click **View Operation Logs** to view the status and progress of the execution.

Roll back parameter configurations

Perform the following steps to roll back parameter configurations.

1. Log on to the Alibaba Cloud E-MapReduce console.

2. Click the **Cluster Management** tab to jump to the Cluster Management page.
3. Click the cluster ID that you want to jump to the Clusters and Services page.
4. In the Services list, click the component that you want to modify parameters. For example, **HDFS**.
5. On the HDFS page, click the **Configuration Change History** tab to jump to the Configuration Change History tab page.
6. Click **roll back** in the Actions column for a record to roll back the corresponding parameter configurations.

See the Deloy configurations section for making a configuration rollback take effect.