

Alibaba Cloud E-MapReduce

Cluster Planning and Configurations

Issue: 20190912

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 content of the Alibaba Cloud website, including but not limited to works, products, images, archives, information, materials, website architecture, website graphic layout, and webpage design, 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 the Alibaba Cloud website, product programs, or content 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.

Generic conventions

Table -1: Style conventions

Style	Description	Example
	This warning information 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.
	This warning information 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 restore business.
	This indicates warning information, supplementary instructions, and other content that the user must understand.	 Notice: Take the necessary precautions to save exported data containing sensitive information.
	This indicates supplemental instructions, best practices, tips, and other content that is good to know for the user.	 Note: You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
Bold	It is used for buttons, menus, page names, and other UI elements.	Click OK.
Courier font	It is used for commands.	Run the <code>cd / d C :/ windows</code> command to enter the Windows system folder.
<i>Italics</i>	It is used for parameters and variables.	<code>bae log list --instanceid Instance_ID</code>
[] or [a b]	It indicates that it is an optional value, and only one item can be selected.	<code>ipconfig [-all -t]</code>

Style	Description	Example
<code>{}</code> or <code>{a b}</code>	It indicates that it is a required value, and only one item can be selected.	<code>switch {stand slave}</code>

Contents

Legal disclaimer.....	I
Generic conventions.....	I
1 Cluster planning.....	1
1.1 Role authorization.....	1
1.2 User management.....	4
1.3 Instance types.....	5
1.4 Components deployed on different instance nodes.....	6
1.5 Gateway clusters.....	12
1.6 ECS instances.....	13
1.7 Storage guide.....	14
1.8 D1 series.....	16
1.9 Enable access between classic networks and VPCs.....	17
1.10 Disaster recovery in E-MapReduce clusters.....	18
2 Configure clusters.....	19
2.1 Create a cluster.....	19
2.2 Cluster list and details.....	25
2.3 Expand a cluster.....	30
2.4 Release a cluster.....	32
2.5 Cluster renewal.....	33
2.6 Connect to a cluster using SSH.....	35
2.7 Service list.....	42
2.8 Cluster scripts.....	43
2.9 Access links and ports.....	45
2.10 Security groups.....	45
2.11 Create a Gateway.....	47
2.12 Manage cluster resources.....	54
2.13 Auto Scaling.....	55
2.13.1 Overview.....	55
2.13.2 Configure Auto Scaling by time.....	56
2.13.3 Preemptible instances in AutoScaling.....	58
2.13.4 Auto Scaling records.....	60
2.14 VPC.....	60
2.15 MetaService.....	62
2.16 Metadata management.....	64
2.16.1 Manage Hive metadata.....	64
2.16.2 Manage Kafka metadata.....	71
3 Modify configurations.....	73
3.1 Upgrade node configuration.....	73
3.2 Expand disks.....	76
3.3 Convert your billing method.....	83

3.4 Remove abnormal nodes.....	83
3.5 Multiple host groups.....	83
4 Third-party software.....	85
4.1 Bootstrap actions.....	85
4.2 Software configuration.....	87
4.3 Configure parameters.....	89

1 Cluster planning

1.1 Role authorization

When you activate the E-MapReduce service, a default system role named **AliyunEMRDefaultRole** must be granted to the E-MapReduce service account. Once assigned, E-MapReduce can call the relevant services (such as ECS and OSS), create clusters, save logs, and perform other related tasks.




Notice:

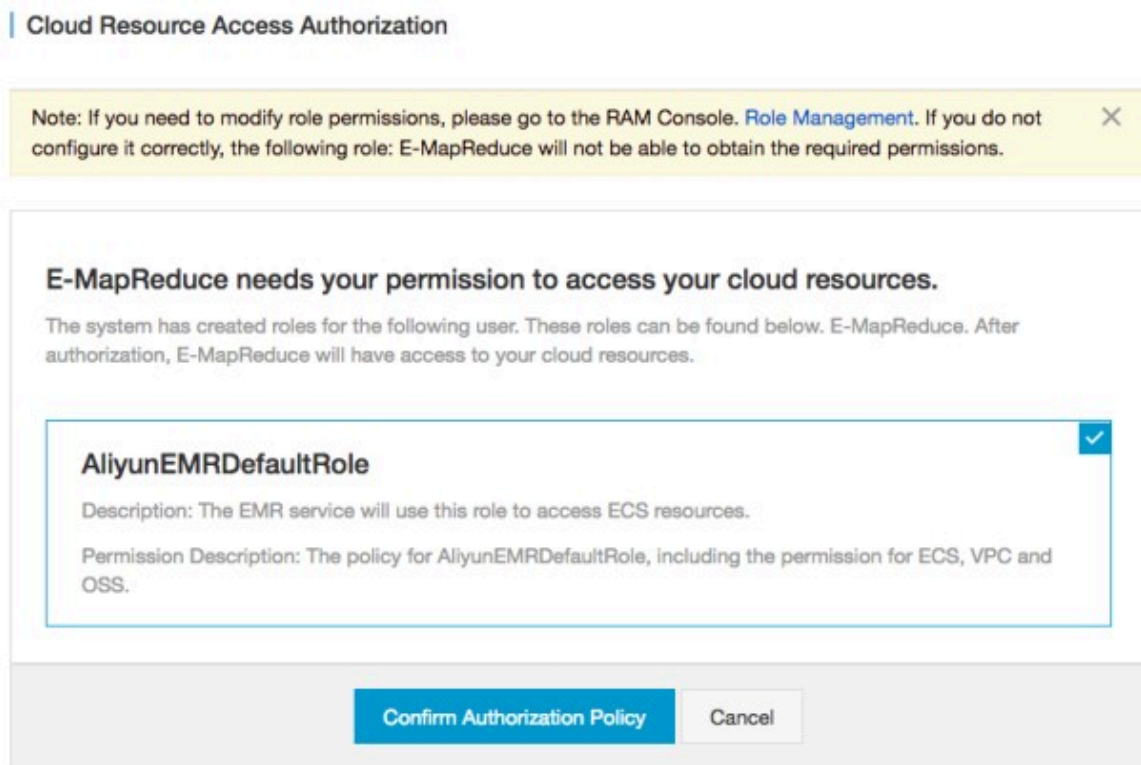
If you are activating E-MapReduce for the first time, you must authorize roles by using the primary account. Otherwise, the primary and user accounts cannot access or use E-MapReduce.

Role authorization process

1. When you create a cluster or an on-demand execution plan, if a default role is not authorized to the E-MapReduce service account, the following prompt is displayed. Click **Go to RAM for authorization** to authorize the role.



2. On the RAM authorization page, click **Confirm Authorization Policy** to authorize the default role **AliyunEMRDefaultRole** to the E-MapReduce service account.



3. Refresh the E-MapReduce console, and then perform relevant operations.
- If you want to view relevant detailed policy information of AliyunE-MapReduceDefaultRole, log on to the RAM console, or click [View Link](#).

Default role permissions

The default role AliyunEMRDefaultRole has the following permissions

- ECS related permissions:

Permission name (Action)	Description
ecs:CreateInstance	Create ECS instances
ecs:RenewInstance	Renew ECS instances.
ecs:DescribeRegions	Query ECS region information.
ecs:DescribeZones	Query zone information.
ecs:DescribeImages	Query image information.
ecs:CreateSecurityGroup	Create security groups.
ecs:AllocatePublicIpAddress	Allocate a public network IP address.

Permission name (Action)	Description
ecs:DeleteInstance	Delete machine instances.
ecs:StartInstance	Start machine instances.
ecs:StopInstance	Stop machine instances.
ecs:DescribeInstances	Query machine instances.
ecs:DescribeDisks	Query the machine's relevant disk information.
ecs:AuthorizeSecurityGroup	Set security group input rules.
ecs:AuthorizeSecurityGroupEgress	Set security group output rules.
ecs:DescribeSecurityGroupAttribute	Query security group details.
ecs:DescribeSecurityGroups	Query security group list information.

- OSS related permissions

Permission name (Action)	Description
oss: PutObject	Upload file or folder objects.
oss: GetObject	Get file or folder objects.
oss: ListObjects	Query file list information.

Grant permissions to a RAM user account

To ensure that user accounts can access the E-MapReduce service, you need to log on with your primary account to the [RAM console](#) and set AliyunEMRFullAccess or AliyunEMRDevelopAccess policies to grant user accounts access to E-MapReduce.



Note:

- The AliyunEMRFullAccess policy grants user accounts with full access permissions to E-MapReduce and E-MapReduce resources.
- The AliyunEMRDevelopAccess policy grants user accounts the E-MapReduce Developer permission, but, unlike AliyunEMRFullAccess, it does not grant users with other access permissions, such as the permissions for creating or releasing E-MapReduce clusters.

For information about RAM user accounts, policies, and roles, see [RAM](#).

1.2 User management

User management allows you to manage the accounts required to create services on specified clusters. E-MapReduce currently supports the creation of two types of accounts: Knox and Kerberos. This topic explains how to manage Knox accounts.

Create a RAM account

1. Log on to the [Alibaba Cloud E-MapReduce console](#) and go to the Cluster Management page.
2. Click Manage on the right side of the target cluster ID.
3. In the navigation panel on the left, click User Management.
4. In the upper-right corner of the page, click Create RAM User.

Add a Knox account

1. In the User Management page, select the account you want to add to a cluster, and then click Set Knox Account Password in the Actions column.
2. In the Add Knox User dialog box, enter a password to use for logon and click OK.
3. Refresh the User Management page. When Synchronized is displayed in the Knox Account column, you have successfully added the Knox account.

You can then sign in to Knox using the User Name and the password set in the preceding step.

For more information about how to use Knox, see [#unique_7](#).

Delete a Knox account

1. In the User Management panel, select the account you want to delete from a cluster, and then click Delete Knox Account in the Actions column.
2. Refresh the User Management page. When Unsynchronized is displayed in the Knox Account column, you have successfully deleted the Knox account.

Kerberos account

When this Kerberos mode is enabled, the software on the cluster runs in Kerberos security mode. For more information, please refer to [#unique_8](#).

FAQs

- Different clusters cannot share the same Knox account. For example, Knox account A that you added to cluster-1 cannot be used in cluster-2. If you want to use Knox account A in cluster-2, you must re-add account A to cluster-2.
- If the message An error occurred while synchronizing the status is displayed when you add a Knox account, click Retry to add it again.
- If you try to add an account multiple times but it fails each time, click Clusters and Services on the left side of the page to check if ApacheDS is stopped. If it is, start ApacheDS and go back to User Management to try again.

1.3 Instance types

There are three types of node instances in an E-MapReduce cluster: master, core, and task.

Different service processes are deployed on each instance type. For example, with Hadoop, the HDFS NameNode and YARN Resource Manager services are deployed on master instances, while the HDFS DataNode and YARN Node Manager services are deployed on core instances. For task instances, because they are only used in computing tasks, only YARN Node Manager is deployed, and not HDFS-related services.

When you create a cluster, you must determine the ECS specifications for each instance type. ECS instances of the same type must be in the same instance group. If you increase the number of hosts in a core or task instance group, you can scale the cluster up at a later date. This does not apply to master instance groups.



Note:

Task instances are supported in version 3.2.0 or later.

Master instance

The master instance is where the management and control components of the cluster service are deployed. You can connect to the master instance using SSH and check service statuses in the cluster through the software's Web UI.

If you want to perform a test or run a job, log on to the master instance and submit jobs directly at the command line. By default, only one master instance is used. However, if the cluster's high availability feature is enabled, two are used.

Core instance

Core instances, which are managed by master instances, store all of the data in the cluster. They also deploy computing services to perform computing tasks. If you need more data storage or are experiencing heavier workloads, you can scale core instances up at any time without impacting the operations of the cluster. For more information, please refer to [Local disks](#) and [Block storage](#).

Task instance

Task instances are responsible for computing and can quickly add computing power to a cluster. They can also scale up and down at any time without impacting the operations of the cluster. However, this instance type is optional, and if the core instance has enough computing power, task instances are not necessary. Depending on the fault tolerance (or retries) of the computing service, a reduction in the number of task instance nodes may cause MapReduce and Spark jobs to fail.

1.4 Components deployed on different instance nodes

This topic describes the deployment details for the components on different instance nodes in an EMR cluster.

When you create an EMR cluster, components are deployed on instance nodes based on the [#unique_13](#). For example, the NameNode component of Hadoop HDFS is deployed on the master instance node. For large-scale clusters, you can use bootstrap actions to complete custom deployment and allocation. For example, a stand-alone ZooKeeper cluster uses an Alibaba Cloud RDS instance as the Hive metastore. The following tables describe the service components that run on different instance nodes.

Hadoop cluster

Deployment details for service components that run on a Hadoop cluster

• Required services

Service name	Master instance node	Core instance node
HDFS	<ul style="list-style-type: none"> - KMS - SecondaryNameNode - HttpFS - HDFS Client - NameNode 	<ul style="list-style-type: none"> - DataNode - HDFS Client
YARN	<ul style="list-style-type: none"> - ResourceManager - App Timeline Server - JobHistory - WebAppProxyServer - Yarn Client 	<ul style="list-style-type: none"> - Yarn Client - NodeManager
Hive	<ul style="list-style-type: none"> - Hive MetaStore - HiveServer2 - Hive Client 	Hive Client
Spark jobs	<ul style="list-style-type: none"> - Spark Client - SparkHistory - ThriftServer 	Spark Client
Knox	Knox	N/A
Zeppelin	Zeppelin	N/A
Tez	<ul style="list-style-type: none"> - Tomcat - Tez Client 	Tez Client
Ganglia	<ul style="list-style-type: none"> - Gmond - Httpd - Gmetad - Ganglia Client 	<ul style="list-style-type: none"> - Gmond - Ganglia Client
Pig	Pig Client	Pig Client
Sqoop	Sqoop Client	Sqoop Client
Bigboot	<ul style="list-style-type: none"> - Bigboot Client - Bigboot Monitor 	<ul style="list-style-type: none"> - Bigboot Client - Bigboot Monitor
OpenLDAP	OpenLDAP	N/A
Hue	Hue	N/A

· Optional services

Service name	Master instance node	Core instance node
Flume	<ul style="list-style-type: none"> - Flume Agent - Flume Client 	<ul style="list-style-type: none"> - Flume Agent - Flume Client
Livy	Livy	N/A
Superset	Superset	N/A
Flink	<ul style="list-style-type: none"> - FlinkHistoryServer - Flink Client 	Flink Client
Ranger	<ul style="list-style-type: none"> - RangerPlugin - RangerAdmin - RangerUserSync 	RangerPlugin
Storm	<ul style="list-style-type: none"> - Storm Client - UI - Nimbus - Logviewer 	<ul style="list-style-type: none"> - Storm Client - Supervisor - Logviewer
Phoenix	Phoneix Client	Phoneix Client
Analytics Zoo	<ul style="list-style-type: none"> - Analytics-Zoo-Scala - Analytics-Zoo-Python 	<ul style="list-style-type: none"> - Analytics-Zoo-Scala - Analytics-Zoo-Python
SmartData	<ul style="list-style-type: none"> - JindoFS Namespace Service - JindoFS Storage Service - SmartData Client 	<ul style="list-style-type: none"> - JindoFS Storage Service - SmartData Client
Kudu	<ul style="list-style-type: none"> - Kudu Master - Kudu Client 	<ul style="list-style-type: none"> - Kudu Tserver - Kudu Master - Kudu Client
HBase	<ul style="list-style-type: none"> - HMaster - HBase Client - ThriftServer 	<ul style="list-style-type: none"> - HBase Client - HRegionServer
ZooKeeper	<ul style="list-style-type: none"> - ZooKeeper follower - ZooKeeper Client 	<ul style="list-style-type: none"> - ZooKeeper leader - ZooKeeper follower - ZooKeeper Client
Oozie	Oozie	N/A

Service name	Master instance node	Core instance node
Presto	<ul style="list-style-type: none"> - Presto Client - PrestoMaster 	<ul style="list-style-type: none"> - Presto Client - PrestoWorker
Impala	<ul style="list-style-type: none"> - Impala Runtime and Shell - Impala Catalog Server - Impala StateStore Server 	<ul style="list-style-type: none"> - Impala Runtime and Shell - Impala Daemon Server - Impala StateStore Server

Druid cluster

Deployment details for service components on a Druid cluster

- Required services

Service name	Master instance node	Core instance node
Druid	<ul style="list-style-type: none"> - Druid Client - Coordinator - Overlord - Broker 	<ul style="list-style-type: none"> - MiddleManager - Historical - Druid Client
HDFS	<ul style="list-style-type: none"> - KMS - SecondaryNameNode - HttpFS - HDFS Client - NameNode 	<ul style="list-style-type: none"> - DataNode - HDFS Client
Ganglia	<ul style="list-style-type: none"> - Gmond - Httpd - Gmetad - Ganglia Client 	<ul style="list-style-type: none"> - Gmond - Ganglia Client
ZooKeeper	<ul style="list-style-type: none"> - ZooKeeper follower - ZooKeeper Client 	<ul style="list-style-type: none"> - ZooKeeper leader - ZooKeeper follower - ZooKeeper Client
OpenLDAP	OpenLDAP	N/A

• Optional services

Service name	Master instance node	Core instance node
YARN	<ul style="list-style-type: none"> - ResourceManager - App Timeline Server - JobHistory - WebAppProxyServer - Yarn Client 	<ul style="list-style-type: none"> - Yarn Client - NodeManager
Superset	Superset	N/A

Kafka cluster

Deployment details for service components on a Kafka cluster

• Required services

Service name	Master instance node	Core instance node
Kafka-Manager	Kafka Manager	N/A
Kafka	<ul style="list-style-type: none"> - Kafka Client - KafkaMetadataMonitor - Kafka Rest Proxy - Kafka Broker controller - Kafka Schema Registry 	<ul style="list-style-type: none"> - Kafka Broker broker - Kafka Client
Ganglia	<ul style="list-style-type: none"> - Gmond - Httpd - Gmetad - Ganglia Client 	<ul style="list-style-type: none"> - Gmond - Ganglia Client
ZooKeeper	<ul style="list-style-type: none"> - ZooKeeper follower - ZooKeeper Client 	<ul style="list-style-type: none"> - ZooKeeper leader - ZooKeeper follower - ZooKeeper Client

• Optional services

Service name	Master instance node	Core instance node
Ranger	<ul style="list-style-type: none"> - RangerPlugin - RangerUserSync - RangerAdmin 	RangerPlugin
Knox	Knox	N/A

Service name	Master instance node	Core instance node
OpenLDAP	OpenLDAP	N/A

Zookeeper cluster

Deployment details for service components on a ZooKeeper cluster

Service name	Master instance node	Core instance node
Ganglia	N/A	<ul style="list-style-type: none"> • Gmond • Httpd • Gmetad • Ganglia Client
ZooKeeper	N/A	<ul style="list-style-type: none"> • ZooKeeper leader • ZooKeeper Client • ZooKeeper follower

Data Science cluster

Deployment details for service components on a Data Science cluster

- Required services

Service name	Master instance node	Core instance node
HDFS	<ul style="list-style-type: none"> - KMS - SecondaryNameNode - HttpFS - HDFS Client - NameNode 	<ul style="list-style-type: none"> - DataNode - HDFS Client
Jupyter	Jupyter	Jupyter
Ganglia	<ul style="list-style-type: none"> - Gmond - Httpd - Gmetad - Ganglia Client 	<ul style="list-style-type: none"> - Gmond - Ganglia Client
ZooKeeper	<ul style="list-style-type: none"> - ZooKeeper follower - ZooKeeper Client 	<ul style="list-style-type: none"> - ZooKeeper leader - ZooKeeper follower - ZooKeeper Client
Analytics Zoo	<ul style="list-style-type: none"> - Analytics-Zoo-Python - Analytics-Zoo-Scala 	<ul style="list-style-type: none"> - Analytics-Zoo-Python - Analytics-Zoo-Scala

Service name	Master instance node	Core instance node
ApacheDS	ApacheDS	N/A
Tensorflow on YARN	<ul style="list-style-type: none"> - TensorFlow-On-YARN - TensorFlow-On-YARN-Gateway - TensorFlow-On-YARN-History-Server 	<ul style="list-style-type: none"> - TensorFlow-On-YARN-Client - TensorFlow-On-YARN-Gateway
Zeppelin	Zeppelin	N/A
Spark jobs	<ul style="list-style-type: none"> - Spark Client - ThriftServer - SparkHistory 	Spark Client
YARN	<ul style="list-style-type: none"> - ResourceManager - App Timeline Server - JobHistory - WebAppProxyServer - Yarn Client 	<ul style="list-style-type: none"> - Yarn Client - NodeManager

• Optional services

Service name	Master instance node	Core instance node
Hue	Hue	RangerPlugin
Hive	<ul style="list-style-type: none"> - HiveServer2 - Hive MetaStore - Hive Client 	Hive Client
Tensorflow	Tensorflow	Tensorflow

1.5 Gateway clusters

A gateway cluster is an independent cluster that consists of multiple nodes of the same configuration.

When you create a gateway cluster, you can associate it with an existing Hadoop cluster. To facilitate cluster operations, it is recommended that you associate it with a cluster on which Hadoop (HDFS and YARN), Hive, Spark, Sqoop, Pig, or other clients have been deployed. It is an independent submission point and does not use up cluster resources, especially when you submit jobs on the Master node. This improves

the stability of the Master node. If there are too many jobs to submit, you can add nodes as and when you need them.

You can also create multiple gateway clusters for different users, allowing them to use their own environment to meet different service requirements.

1.6 ECS instances

This section contains information on the different ECS instance types.

ECS instance types supported by E-MapReduce

- General-purpose

This type uses cloud disks as storage. The ratio between vCPUs and memory is 1:4. For example, 32 cores and 128 GiB memory.

- Compute

This type uses cloud disks as storage and provides more computing resources. The ratio between vCPUs and memory is 1:2. For example, 32 cores and 64 GiB memory.

- Memory

This type uses cloud disks as storage and provides more memory resources. The ratio between vCPUs and memory is 1:8. For example, 32 cores and 256 GiB memory.

- Big data

This type uses local SATA disks as storage, which is highly cost-effective. If you want to store massive amounts of data (TB-level), it is recommended that you use this type.



Note:

Currently only core nodes supported by Hadoop, Data Science, and Druid clusters support big data instances. Zookeeper and Kafka clusters do not support core nodes.

- Ephemeral SSD

This type uses ephemeral SSDs as storage, which provides high local IOPS and throughput.

- Shared (entry level)

This type shares CPUs and is not stable enough for most scenarios. It is applicable for entry-level users, not enterprise customers.

- GPU

This type is a heterogeneous GPU-based model applicable in machine learning scenarios.

ECS instance types applicable in different scenarios

- Master instances

General-purpose and memory types are applicable in master instances, where data is directly stored on Alibaba Cloud's cloud disks. There are also three backups to guarantee high data reliability.

- Core instances

General-purpose, compute, and memory types are applicable for small data volumes (not TB-level) or when OSS is used as the primary data storage. When the amount of data is large (10 TB or more), it is recommended that you use the big data type, which is more cost-effective. Using an ephemeral disk makes it harder to ensure data reliability, but this can be maintained and guaranteed by the E-MapReduce platform.

- Task instances

All types except the big data type are suitable for task instances to give additional computing power to the cluster. Currently, the ephemeral SSD type is not supported, but will be added soon.

1.7 Storage guide

There are two types of disks on a node: the system disk, which is used to install operating systems, and the data disk, which is used to store data.

A node typically has one system disk by default, which must be a cloud disk. However, you can have more than one data disk (currently, up to sixteen on a single node). Each data disk can have different configurations, including having a different type or capacity. In E-MapReduce, a cluster's system disks are SSD cloud disks by default, and four are used by default. Considering current intranet bandwidth, this default configuration of four cloud disks is sufficient.

Cloud and ephemeral disks

Two types of disk are available for data storage.

- Cloud disks

Includes SSD, ultra, and basic cloud disks.

Cloud disks are not attached directly to the local computing node. Instead, they access a remote storage node through the network. Each piece of data has two real-time backups at the backend, meaning that there are three identical copies in total. When one is corrupted (due to disk damage), a backup is used automatically for recovery.

- Ephemeral disks

Includes ephemeral SATA disks in the big data type and ephemeral SSD disks used in the ephemeral SSD type.

Ephemeral disks are attached directly to the computing node and have a better performance than cloud disks. You cannot change the number of ephemeral disks. As with offline physical hosts, there is no data backup at the backend, meaning that upper-layer software is required to guarantee data reliability.

Usage scenarios

In E-MapReduce, when the hosting node is released, all of the data in the cloud and ephemeral disks is cleared. The disks can also not be kept independently and used again. Hadoop HDFS uses all data disks for data storage. Hadoop YARN uses all data disks as on-demand data storage for computing.

If you do not have massive amounts of data (below TB-level), you can use cloud disks, as the IOPS and throughput are smaller than local disks. In the event that you have large amounts of data, it is recommended that you use local disks whose data reliability is guaranteed by E-MapReduce. If you find the throughput to be insufficient, switch to ephemeral disks.

OSS

OSS can be used as HDFS in E-MapReduce, and you can have easy read and write access to OSS. All code that uses HDFS can also be easily modified to access data on OSS. Below you can find a number of examples:

Reading data from Spark

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

Changing the storage type from HDFS to OSS

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

This is the same for Map Reduce and Hive jobs.

HDFS commands process OSS data directly:

```
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 AK or endpoint. E-MapReduce automatically completes your information using the current cluster owner.

However, as OSS does not have high IOPS, it is not suitable for usage scenarios that require high IOPS, such as Spark Streaming or HBase.

1.8 D1 series

To meet the demand for storage in big data scenarios, Alibaba Cloud has launched the D1 series on the cloud.

Instead of using cloud disks in its data storage, the D1 series uses ephemeral disks , which solves the problem of high costs caused by keeping multiple copies of redundant data in cloud disks. Data also no longer needs to be transferred over the network, which improves disk throughput. Furthermore, with the D1 series, you can also take advantage of Hadoop's proximity computing.

Compared with cloud disks, the series greatly enhances storage performance while reducing prices. Indeed, the cost is almost the same as offline physical hosts.

Despite their advantages, however, ephemeral disks still cannot ensure data reliability, and upper-layer software is required to guarantee it. If a disk or node fails, operations and maintenance must be performed manually. Cloud disks, meanwhile, guarantee data reliability automatically, meaning that you do not need to worry about disk damage. Alibaba Cloud's default multi-disk backup policy is also helpful in this regard.

E-MapReduce + D1 solution

A complete set of automated O&M solutions, such as the D1 series, is now available for ephemeral disks in E-MapReduce. This allows Alibaba Cloud users to use ephemeral disks conveniently and reliably, without having to worry about the entire O&M process. Data reliability and service availability are guaranteed.

The main advantages are as follows:

- Highly reliable distribution of required nodes
- Ephemeral disk and node fault monitoring
- Automatic determination of data migration opportunities
- Automatic failed node migration and data balancing
- Automatic HDFS data detection
- Network topology optimization

With the automated O&M of the entire back-end management and control system, E-MapReduce helps you make better use of ephemeral disks and develop a cost-effective big data system.



Note:

If you want to set up a Hadoop cluster using the D1 series, submit a ticket. We will then be able to assist you in your operations.

1.9 Enable access between classic networks and VPCs

This section describes how to enable inter-access between ECS on classic networks and E-MapReduce clusters on VPC networks.

ClassicLink

Alibaba Cloud currently provides two types of cloudnetwork: classic and VPC. While some users still use classic networks, E-MapReduce clusters use VPCs.

To grant access between ECS on a classic network and an E-MapReduce cluster on a VPC network, Alibaba Cloud launches [ClassicLink](#). Follow these steps:

1. Create a vSwitch according to the CIDR block specified in ClassicLink.
2. To deploy a cluster that you have created, use the vSwitch of the CIDR block.
3. Connect the corresponding classic network node to the VPC in the ECS console.
4. Set the security group rules.

1.10 Disaster recovery in E-MapReduce clusters

This article will introduce disaster recovery of data and services 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

In this tutorial, you will learn how to create an Alibaba Cloud E-MapReduce (EMR) cluster.

Go to the EMR cluster creation page

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. Complete RAM authorization. For details, see [#unique_22](#).
3. Select a region for the cluster. The region cannot be changed once the cluster is created.
4. Click Create Cluster to go to the cluster creation page.

Create a cluster



Notice:

After you create an EMR cluster, the only thing that can be changed is its name.

To create a cluster, follow these three steps:

1. Software Settings

- **EMR version:** The main version of E-MapReduce represents a complete open source software environment and can be upgraded regularly based on upgrades made to the internal component software. If the software related to Hadoop is

upgraded, the main version of E-MapReduce is also upgraded. Clusters from an earlier version cannot be upgraded to a later version.

- **Cluster type:** Currently, E-MapReduce provides four cluster types.
 - Hadoop clusters, which provide the following semi-managed ecosystem components:
 - Hadoop, Hive, and Spark for large-scale offline distributed data storage and computing.
 - Spark Streaming, Flink, and Storm for stream processing.
 - Presto and Impala for running interactive analytics.
 - Oozie and Pig.
 - Druid clusters, which provide semi-managed, real-time interactive analysis services, query large amounts of data at millisecond latency, and support multiple data intake methods. When used with services such as EMR Hadoop, EMR Spark, OSS, and RDS, Druid clusters offer real-time query solutions.
 - Data Science clusters, which are mainly applicable in big data and AI scenarios, providing Hive and Spark offline big data, and TensorFlow model training.
 - Kafka clusters, which are semi-managed distributed message systems that feature high throughput and high scalability, providing a complete service monitoring system that can keep a stable running environment.
- **Required Services:** Displays a list of all software components under the selected cluster type, including their name and version number.
- **Optional Services:** You can select different components as required. The selected components start relevant service processes by default.



Note:

The more components you select, the higher the requirements are for configuration, as there may be insufficient resources to run these services.

- **Advanced Settings**
 - **Kerberos Mode:** In this mode, you can set the cluster's Kerberos authentication. This feature is unnecessary for clusters used by individual users and is turned off by default.
 - **Custom Software Settings:** You can specify a JSON file to modify software configuration. For more information about the procedure, see [Software configuration](#).

2. Hardware Settings

- **Billing Configuration**
 - **Billing Method:** You can select Pay-As-You-Go when testing the cluster. After all tests have been passed, you can create and use a Subscription-based cluster.
 - **Pay-As-You-Go**
 - **Subscription**
 - **Subscription Period:** You can get 15% off your total cost by choosing annual subscription.
 - **Auto Renewal:** Your subscription will be automatically renewed seven days before the expiration date. The subscription duration is one month.
- **Network Settings**
 - **Zone:** Select the zone where the cluster is to be located. If better network connectivity is required, we recommend selecting the same availability zone. However, this increases the risk of failure when creating a cluster, as the availability zone's storage may be insufficient. If you need a large number of nodes, please submit a ticket.
 - **Network type:** The Virtual Private Cloud (VPC) network is selected by default, which requires you to enter a VPC and a VSwitch. If you have not created a

network, go to the [VPC console](#) to create one. For more information about E-MapReduce VPC, see [#unique_24](#).

- VPC: Select the region of the VPC network.
- VSwitch: Select a zone for VSwitch under the corresponding VPC. If no VSwitch is available in this zone, you must create a new one.
- Security group name: A security group does not typically exist when you first create a cluster. To create a new security group, enter a name. If you already have a security group, you can select it here.
- Cluster configuration
 - High availability: When enabled, two master instances in the Hadoop cluster are used to ensure the availability of the Resource Manager and Name Node. HBase clusters support high availability by default.
 - Node type: The three types of node supported are as follows:
 - Master, which is mainly responsible for the deployment of control processes such as Resource Manager and Name Node.
 - Core, which is mainly responsible for the storage of all data in the cluster, and can be scaled up as required.
 - Task, which is the node used for computing. It does not store data and is used to adjust the computing capacity of the cluster.
 - Node configuration: Select different node types. Different types of nodes have different application scenarios.
 - Data disk type: The data disks used by a cluster node are either standard cloud disks, high-efficiency cloud disks, or SSD cloud disks. This varies between machine type and region. When the user selects different regions, disks that are supported by those regions are displayed in the drop-down list. By default, data disks are released when the cluster is released. The ephemeral disk type is set by default and cannot be changed.
 - Data disk volume: The recommended minimum cluster volume for a single machine is 40 G, and the maximum is 8000 G. The capacity of the ephemeral disk is set by default and cannot be changed.
 - Instance quantity: This indicates the number of instances of all required nodes. A cluster requires at least three instances. However, high availability clusters require at least four, and therefore add one master node.

3. Basic Settings

- Basic information

- Cluster name: The cluster name can contain Chinese characters, English letters (uppercase and lowercase), numbers, hyphens (-), and underscores (_), with a length of between 1-64 characters.
- Logon settings
 - Remote logon: It is turned on by default to enable security group port 22.
 - Key Pair: For the use of the key pair, please refer to [SSH Key Pair](#)
 - Password: Set the logon password at the master node. The logon password must contain English letters (both uppercase and lowercase letters), numbers, and special characters (!@#\$%^&*) with a length of between 8-30 characters.

- Advanced Settings

- Unified Metabases: Hive uses a unified metadatabase, which is independent of the cluster. The meta information will not be deleted after the cluster is released. We recommend that you disable this feature if you are an E-MapReduce beginner.
- Add Knox User: This account is used to access the Web interface of open-source big data software.
- Permission Settings: Typically, the default setting is used.
 - EMR role: This role authorizes E-MapReduce to use other Alibaba Cloud services, such as ECS and OSS.
 - ECS role: This role allows your programs running on the E-MapReduce computing nodes to access cloud services like OSS without providing the Alibaba Cloud AccessKey. E-MapReduce automatically applies for an on-demand AccessKey to authorize access. The AccessKey permission is controlled by this role.
- Bootstrap Actions (optional): Before Hadoop is enabled in the cluster, you can run the customized script. For more information, see [#unique_26](#)

Purchase lists and cluster costs

The cost of the cluster is displayed in the Configuration List pane. The price varies with the type of payment. For Subscription clusters, the total expense is shown. For Pay-As-You-Go clusters, the hourly cost is shown.

Confirm creation

Once you have entered all the necessary information, the Create button is highlighted. Click Create to create the cluster.



Note:

- If your cluster is Pay-As-You-Go, it is created immediately, and you are taken back to the Overview page. Here, your cluster is displayed with the status Initializing. It can take several minutes to finish creating the cluster. After the cluster is created, its status is switched to Idle.
- Subscription clusters are not created until the order is generated and paid.

Log on to the core node

To log on to the core node, perform the following steps:

1. Switch to the Hadoop account on the master node.

```
su    hadoop
```

2. Log on to the core node through SSH without a key.

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

3. Get root permissions through the sudo command.

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

Failure during cluster creation

If a cluster fails to be created, the message Cluster creation failed is displayed on the cluster list page. If you hover your cursor over the red exclamation point, the reason for the failure is displayed.

You do not need to perform any additional operations because the corresponding computing resources are not created. The cluster is automatically hidden after three days.


2.2 Cluster list and details

You can learn about the E-MapReduce cluster information from the cluster list or cluster details.

Cluster list

The Cluster Management page, displays basic information about all of your clusters.

On the Cluster Management page, information of clusters are displayed as follows:



The screenshot shows the E-MapReduce Cluster Management interface. At the top, there is a navigation bar with tabs: Overview, Cluster Management (selected), Data Platform, Alert, Operation Logs, Help, and Old EMR Scheduling. Below the navigation bar, there is a 'Cluster Management' section with a 'Refresh' button and two buttons: 'Create Gateway' and 'Create Cluster'. The main content is a table with the following columns: Cluster ID/Name, Cluster Type, Status, Created At, Runtime, Billing Method, and Actions. There are two rows of cluster data.

Cluster ID/Name	Cluster Type	Status	Created At	Runtime	Billing Method	Actions
E-MR-CLUSTER-NAME dplus_docs	HADOOP	Idle	2018-10-29 09:41:45	4 Hours 16 Minutes 21 Seconds	Pay-As-You-Go	Monitoring Manage View Details More
E-MR-CLUSTER-NAME dual o	HADOOP	Idle	2018-10-24 14:47:08	4 Days 23 Hours 10 Minutes 58 Seconds	Subscription Expiration Date 2018-11-25 00:00:00	Monitoring Manage View Details Review More

The items in the cluster list are as follows:

- **Cluster ID/Name:** The ID and name of a cluster. Move your cursor over a cluster's name to modify it.
- **Cluster Type:** The types include Hadoop, Druid, Data Science, and Kafka.
- **Status:** The status of a cluster. For more information, see [Cluster statuses](#). If a cluster experiences an abnormality, such as a creation failure, prompt information appears on the right. If you hover your cursor over it, you can view detailed error information. You can also sort the statuses by clicking Status.
- **Created At:** Time at which a cluster was created.
- **Runtime:** The time from the point of creation to the current time. Once the cluster is released, the timing is terminated.
- **Billing Method:** The billing method of the cluster.

- **Actions:** Operations that can be performed on clusters, including the following:
 - **Monitoring:** Monitors the CPU usage rate, memory capacity, and disk capacity of E-MapReduce clusters to help users monitor the running status of the cluster.
 - **Manage:** Enters the Clusters and Services panel.
 - **View details:** Enters the Cluster Overview panel and view detailed information after the cluster is created.
 - **More:**
 - **Scale Up/Out:** Expands the cluster.
 - **Release:** Releases a cluster. For more information, see [Release a cluster](#).
 - **Restart:** Restarts a cluster.

Cluster details

Cluster details display detailed information about a cluster.

Detailed information is provided in the following areas: cluster, software, network, and host:

• Cluster

Cluster		
Name: dtplus_docs ID: c-47f8c4f771120f0c01 Region: cn-hangzhou Start Time: 2018-10-25 10:36:51	Software Configuration: I/O Optimization: Yes High Availability: No Security Mode: Standard	Billing Method: Pay-As-You-Go Current Status: Idle Runtime: 3 Hours27 Minutes8 Seconds

- **Name:** The name of a cluster.
- **ID:** The instance ID of a cluster.
- **Region:** The region where a cluster is located.
- **Start Time:** The time at which a cluster is created.
- **Software Configuration:** Software configurations.
- **I/O Optimization:** Whether the I/O optimization setting is enabled.
- **High Availability:** Whether high-availability clusters are enabled.
- **Security Mode:** Software in clusters is started in Kerberos secure mode. For more information about Kerberos, see [Introduction to Kerberos](#).
- **Billing Method:** Cluster billing method.
- **Current Status:** For more information, see [Cluster statuses](#).
- **Runtime:** The time from the point of creation to the current time.
- **Unified Metabases:** Whether Unified Metabases is enabled.
- **Bootstrap:** The names, paths, and parameters of all configured bootstrap actions are listed here.
- **ECS Role:** When your program runs on an E-MapReduce compute node, you can access the related Alibaba Cloud services, such as OSS, without an AccessKey. E-MapReduce automatically requests a temporary AccessKey to authorize this access. The permission control of this temporary AccessKey is controlled by this role.

• Software

Software
EMR Version: EMR-3.13.0 Cluster Type: HADOOP Software: HDFS2.7.2 / YARN2.7.2 / Hive2.3.3 / Ganglia3.7.2 / Spark2.3.1 / HUE4.1.0 / Tez0.9.1 / Sqoop1.4.7 / Pig0.14.0 / ApacheDS2.0.0 / Knox0.13.0

- **Main Version:** The main version of E-MapReduce.
- **Cluster Type:** The selected cluster type.
- **Software:** All application programs installed are listed here with their versions, such as HDFS2.7.2, Hive 2.3.3, or Spark 2.3.1.


- **Network**

Network
Region ID: cn-hangzhou-f
Network Type: vpc
Security Group ID: sg-bp13guzhwhw3q355kqnl
VPC/VSwitch: vpc-bp13guzhwhw3q355kqnl / vsw-bp13guzhwhw3q355kqnl

- **Region ID:** The region where a cluster is located, such as cn-hangzhou-b, which is the same as ECS.
- **Network Type:** The network type of a cluster.
- **Security Group ID:** The ID of the security group that a cluster joined.
- **VPC/VSwitch:** The VPC and VSwitch IDs of a cluster.

- Host

- Master Instance Group (Master): Configurations of all master nodes.

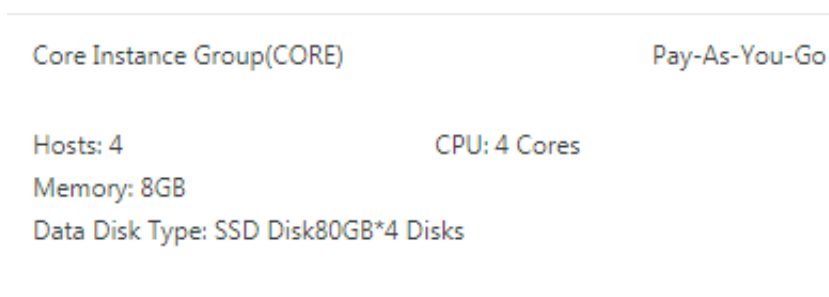


- Hosts: The number of current nodes. During the creation process, the number of current nodes is less than the number of nodes you applied for until the creation is complete.
- CPU: The number of cores in a node's CPU.
- Memory: Memory capacity of a node.
- Data Disk Type: Data disk type and capacity of a node.
- ECS ID: The ID of the ECS instances purchased.

Master Instance Group


ECS ID	Status	Public IP	Intranet IP	Created At
i-3p0fzq2h3j8l9zcvoy0	Normal	47.99.190.51	192.168.0.46	2018-10-25 10:36:58

- Status: Includes Creating, Normal, Expanding, and Released.
- Public IP : The public IP of master nodes.
- Intranet IP: The internal network IP of the machine that can be accessed by all nodes in the cluster.
- Created At: The creation time of the ECS instance purchased.
- Core Instance Group (Core) : Configurations of all core nodes.



- Hosts: The number of current nodes. This is the same as the number of the nodes you applied for.
- CPU: The number of cores in a node's CPU.

- **Memory:** Memory capacity of a node.
- **Data Disk Type:** Data disk type and capacity of a node.
- **ECS ID:** The ID of the ECS instances purchased.

Core Instance Group 

ECS ID	Status	Public IP	Intranet IP	Created At
i-4qp74tqjz4b43a08k4kqg	Normal		192.168.3.127	2018-10-25 10:36:59
i-4qp74tqjz4b43a08k4kqg	Normal		192.168.3.128	2018-10-25 10:37:00
i-4qp74tqjz4b43a08k4kqg	Normal		192.168.3.129	2018-10-25 10:53:07
i-4qp74tqjz4b43a08k4kqg	Normal		192.168.3.131	2018-10-25 10:53:09

- **Status:** Includes Creating, Normal, Expanding, and Released.
- **Intranet IP:** The internal network IP of the machine that can be accessed by all nodes in the cluster.
- **Created At:** The creation time of the ECS instance purchased.

2.3 Expand a cluster

If your cluster does not have enough resources, you can expand it horizontally. Only core and task nodes can be expanded. When expanding a cluster, configurations default to be consistent with the ECS instance purchased previously.

Enter expansion interface

1. On the Cluster Management page, click Details in the Actions column for the cluster that you want to upgrade node specification for.
2. On the Cluster Overview page, click Change Configuration. From the drop-down list, select Scale Out.

Expansion interface

CORE (Core Instance Group) TASK (Task Instance Group)

Configuration: ecs.n4.xlarge 4 Cores 8G SSD

Billing Method: Pay-As-You-Go

Current Core Instances: 2 Instances

New Instances: 2 Instances

VSwitch: es_test_switch




Note:

Although expansion is supported, reduction is not.

- **Configuration:** Displays the configurations of the current instance.
- **Billing Method:** Displays the payment method of the current cluster.
- **Current Core Instances:** Displays the total number of your current core nodes.
- **New Instances:** Enter the quantity of instances that you want to add.
- **VSwitch:** Displays the VSwitch of the current cluster.

Expansion status

In the following figure, the statuses of the cluster expansions are shown.

Core Instance Group 

ECS ID	Status	Public IP	Intranet IP	Created At
i-bp1htgltshh2o59kxqz	● Normal		192.168.0.47	2018-10-25 10:36:59
i-bp1733wjev9dzvp62n7	● Normal		192.168.0.48	2018-10-25 10:37:00
i-bp1dhrigans54rgvcfn2	● Scaling Up/Out		192.168.0.50	2018-10-25 10:53:07
i-bp1c0d1hnoaicxdds57t	● Scaling Up/Out		192.168.0.51	2018-10-25 10:53:09

To view the expansion status of a cluster, click Core Instance Group (CORE) in the Cluster Overview panel. Nodes that are being expanded are displayed as Scaling Up/Out. When the status of an ECS instance changes to Normal, ECS has been added to the cluster and can now provide services.

Change password

After you expand a cluster, you can log on to the expanded node with SSH and change your root password. To do so, follow these steps:

1. Log on the master host with SSH using the following command and obtain the public IP of the master cluster in the [#unique_32/unique_32_Connect_42_section_ntj_xsk_bhb](#) panel.

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

2. Switch to the hadoop user.

```
su hadoop
```

3. Log on to the expanded node and obtain the intranet IP of the expanded node in the [#unique_32/unique_32_Connect_42_section_ntj_xsk_bhb](#) panel.

```
ssh ip.of.worker
```

4. Change your root password using the following command.

```
sudo passwd root
```

2.4 Release a cluster

Pay-As-You-Go clusters can be released on the Cluster Management page. Only those with the following statuses can be released:

- Creating
- Running
- Idle

Common release

Before releasing a cluster, you are prompted to confirm the operation. Once the release is confirmed, the following occurs:

- All jobs in the cluster are forcibly terminated.
- If you have selected to save logs to OSS, all current job logs are saved to OSS. This may take several minutes.
- Clusters are released. Depending on their size, this may take seconds or minutes. ECS clusters are billed before they are released.

**Warning:**

To save money, make sure to release a cluster before a new billable hour starts.

Forcible release

If you no longer need logs and want to immediately terminate running clusters, perform a forcible release. Logs are not collected and clusters are released directly.

Cluster release failure

Due to system errors or other causes, clusters may fail to be released. If a failure occurs, E-MapReduce starts background protection until the cluster is finally released .

2.5 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.

<input checked="" type="checkbox"/>	ECSExpiration Date	EMRExpiration Date	Quantity	ECSList	ECSSubscription Duration	EMRSubscription Duration	Price
<input checked="" type="checkbox"/>	2018-11-25 00:00:00	2018-11-25 00:00:00	1	i-bp15l0ep12jnmokkz20x	1 Month ▾	1 Month ▾	0

Price: ¥ Calculating

☒ (E-MapReduce Service Terms)

OK

- **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.

2.6 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 [#unique_32](#) 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 Connection > 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 hadoop@masterNode IP
```

- Use a username and password:

```
ssh -ND 8157 hadoop@masterNode IP
```

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 @
masterNode IP
```

```
// Method b : Use a username and password
```

```
ssh - ND 8157 hadoop @ masterNode IP
```

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 = "" fromSubscr iption = " false "
enabled = " true " mode = " manual " selectedTa bIndex
= " 2 " lastresort = " false " animatedIc ons = " true "
includeInC ycle = " true " color = "# 0055E5 " proxyDNS = "
true " noInternal IPs = " false " autoconfMo de = " pac
" clearCache BeforeUse = " false " disableCac he = " false
" clearCooki esBeforeUs e = " false " rejectCook ies = "
false ">
< matches >
< match enabled = " true " name = " 120 .*" pattern = " http
:// 120 .*" isRegEx = " false " isBlackLis t = " false
" isMultiLin e = " false " caseSensit ive = " false "
fromSubscr iption = " false " ></ match >
</ matches >
< manualconf host = " localhost " port = " 8157 "
socksversi on = " 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 : masterNode IP :
8088 hadoop @ masterNode IP
```

- Use a username and password:

```
ssh -N -L 8157 : masterNode IP : 8088 hadoop @
masterNode IP
```


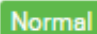

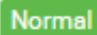
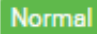

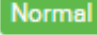

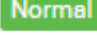

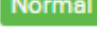

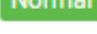
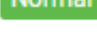
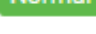
Parameter description:

- `path` : The path where the private key file is stored.
- `masterNode IP` : 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.

2.7 Service list

HDFS, YARN, and other services are listed on the Clusters and Services tab of the cluster management page. You can view the running statuses of these services.

The service list shows the following information.

Services 			
	HDFS		Actions ▼
	YARN		Actions ▼
	Hive		Actions ▼
	Ganglia		Actions ▼
	Spark		Actions ▼
	Hue		Actions ▼
	Tez		Actions ▼
	Sqoop		Actions ▼
	Pig		Actions ▼
	HAProxy		Actions ▼
	ApacheDS		Actions ▼
	Knox		Actions ▼
	EmrFlow		Actions ▼

A service's running status is only showed for clusters that are in the Idle or Running status. Services that are not checked when creating a cluster, such as Storm, are not listed.

Click a service to view the corresponding tabs, including Status, Component Topology, Configuration, and Configuration Change History. The status can either be

Normal or Error. If the status of a service on a node is Error, you can use the master node to log on to the node and check the service process.

2.8 Cluster scripts

To modify a cluster's running environment, you can install third-party software in the cluster. This mostly applies to Subscription clusters. After a cluster is created, the cluster script feature allows you to select nodes in batches and run your specified script to meet your own requirements.

Role of a cluster script

A cluster script is similar to a bootstrap action. After creating a cluster, you can install packages that are unavailable to your cluster before, such as:

- YUM, which is used to install software already provided.
- Public software packages from the public network.
- Software that allows you to read your data from OSS.
- Services like Flink or Impala, which require a more complex script.

We strongly recommend that you test the cluster script on a node first. After the script has been verified, you can perform operations on the whole cluster.

Create and run a cluster script

1. A cluster script can run on an idle cluster or a running cluster. On the Cluster Management page, click Details next to the cluster you want to run a script on.
2. On the menu on the left, click Cluster Scripts to enter the cluster script execution interface.
3. In the upper-right corner, click Create and Run to enter the creation interface.
4. Enter configurations in the script creation pane, select a node for execution, and click OK.

To update the running status of the cluster script on each node, click Refresh. To display the running status of the cluster script on each node, click View Details.

Cluster scripts are applicable for long-standing clusters and can only run on available clusters that are idle or running. To initialize on-demand clusters, perform a bootstrap action.

The cluster script feature downloads a script from the OSS and runs it on a specified node. If the returned value is 0, the execution has failed. If the execution fails, you

can log on to the node to check the running log. The running log for each node is located at `/var/log/cluster-scripts/clusterScr iptId`. If the cluster is configured with an OSS log directory, the running log is also uploaded to `osslogpath / clusterId / ip / cluster - scripts / clusterScr iptId`.

By default, the root account is used to run the specified script. In the script, you can use `su hadoop` to switch to a Hadoop account.

A cluster script can successfully run on some nodes, but fail on others. For example, restarting a node can lead to a failure in script operation. After resolving the error, you can run the cluster script again. After a cluster is expanded, you can specify the expanded node for separate execution of the cluster script.

Only one cluster script can run on a cluster at a time. For each cluster, you can retain up to ten cluster script records. If you have ten records and want to create a new cluster script, you must first delete the previous records.

Script example

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

```
#!/ bin / bash
osscommand -- 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 >
```

OSSCMD is pre-installed on the node and can be called directly to download the file.



Note:

The OSS host address can be an intranet address, Internet address, or VPC network address. If a classic network is used, you must specify an intranet address. If the network is located in Hangzhou, the intranet address is `oss-cn-hangzhou-internal.aliyuncs.com`. If a VPC network is used, you must specify a domain name that can be accessed from the VPC intranet. If the network is located in Hangzhou, the domain name is `vpc100-oss-cn-hangzhou.aliyuncs.com`.

Additional system software packages can be installed on the script using YUM. For example, `ld-linux.so.2`:

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

2.9 Access links and ports

This section serves as a quick portal for components.

When a cluster is created, several domain names are bound to it by default so that you can access the following open source components:

- HDFS
- YARN
- Spark
- Hue
- Ganglia

In the Cluster Management page, click **Manage** or **View Details**, and find the **Access Links and Ports** tab where you can see links for these components.

By default, there is no user name or password required for access, which means that the access request cannot pass the HTTP authentication. Click **Set User Name and Password** to set a user name and password to access your component interface.

Only one user name and one password can be used. A new user name and password will always replace previous ones.



Note:

Currently, this function is only supported by version 2.3 and later.

2.10 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.

2.11 Create a Gateway

A Gateway is an ECS instance in the same intranet as the E-MapReduce cluster. You can use a Gateway to achieve load balancing and security isolation, or submit jobs to the E-MapReduce cluster.

You can use the following methods to create a Gateway.

- (Recommended) Use the [Alibaba Cloud E-MapReduce console](#).
- Create a Gateway manually.

Create a Gateway on the E-MapReduce console

E-MapReduce Gateway only supports E-MapReduce Hadoop clusters. Before creating a Gateway, make sure that you have created an E-MapReduce Hadoop cluster. To create a Gateway, perform the following steps.

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. Click Create Gateway.
3. Perform configurations on the Create Gateway page.
 - **Cluster Name:** the name of the Gateway. The length can be 1 to 64 characters. It only allows Chinese characters, letters, numbers, hyphens (-), and underscores (_).
 - **Password/Key Pair:**
 - **Password:** the password used to log on to the Gateway. The password must be 8 to 30 characters in length and can contain uppercase letters, lowercase letters, numbers, and special characters including exclamation marks (!), at

signs (@), number signs (#), dollar signs (\$), percent signs (%), ampersands (&), and asterisks (*).

- **Key Pair:** the name of the key pair used to log on to the Gateway. If no key pair has been created, click Create Key Pair to go to the ECS console for creating a key pair. Keep the private key file (PEM format) safe. After the Gateway is created, the public key is bounded to the ECS instance on which the Gateway

is deployed. When you log on to the Gateway by using SSH, enter the private key provided by the private key file.

- **Billing Method**
 - **Subscription:** one-time payment for a period of time. You get a lower price by choosing Subscription as the billing method. A big discount is offered for a three-year Subscription.
 - **Pay-As-You-Go:** payment is made according to the number of hours actually used and is calculated on an hourly basis.
- **Associated Cluster:** the cluster to which the Gateway can submit jobs. The Gateway is automatically configured with the same Hadoop environment as the cluster.
- **Zone:** the zone in which the associated cluster is.
- **Network Type:** the network type used by the associated cluster.
- **VPC:** the VPC where the associated cluster is in.
- **VSwitch:** a VSwitch in the VPC.
- **Security Group Name:** the security group to which the associated cluster belongs.
- **Assign Public Network IP:** indicates whether the Gateway is assigned with a public network IP address.
- **Gateway Instance:** the available ECS instance specifications in the region.
- **System Disk Type:** the system disk type that the Gateway node uses. System disk types include SSD and Ultra Disk. The system disk is released with the release of the cluster by default.
- **Disk Size:** the minimum is 40 GB and the maximum is 500 GB. The default value is 300 GB.
- **Data Disk Type:** the data disk type that the Gateway uses. Data disk types include SSD and Ultra Disk. The data disk is released with the release of the cluster by default.
- **Disk Size:** the minimum is 200 GB and the maximum is 4,000 GB. The default value is 300 seconds.
- **Count:** the number of data disks. The minimum is 1 and the maximum is 10.

4. Click Create to save the configurations.

The new Gateway is displayed in the cluster list and Idle appears in the Status column.

Create a Gateway manually

- Network environment

Make sure that the Gateway node is in the security group of the corresponding E-MapReduce cluster and the Gateway node can access the E-MapReduce cluster. For more information, see [Create a security group](#).

- Software environment

- System environment: we recommend that you use CentOS 7.2 or later versions.
- Java environment: JDK 1.7 or later versions. We recommend that you use OpenJDK 1.8.0.

- Procedure

- E-MapReduce V2.7 and later versions or V3.2 and later versions

We recommend that you use the E-MapReduce console to create a Gateway when you use these versions of EMR.

If you want to create a Gateway manually, create a script file, enter the following commands, and then run it on the Gateway node. Run the `sh deploy . sh < master_ip > master_password_file command`.

■ `deploy.sh`: the script name. The commands are as follows.

■ `master_ip`: the IP address of the master node in the cluster. Make sure that the IP address can be accessed.

■ `master_password_file`: the file that stores the password of the master node.

```
#!/usr/bin/bash
if [ $# != 2 ]
then
    echo " Usage : $ 0  master_ip  master_password_file "
    exit 1 ;
fi
masterip=$1
masterpassword_file=$2
if ! type sshpass >/dev/null 2>&1 ; then
    yum install -y sshpass
fi
if ! type java >/dev/null 2>&1 ; then
    yum install -y java-1.8.0-openjdk
fi
mkdir -p /opt/apps
```

```

mkdir -p /etc/ecm
echo "Start to copy package from $masterip to
local gateway (/opt/apps)"
echo "- copying hadoop - 2.7.2"
sshpass -f $masterpwdfile scp -r -o 'StrictHost
KeyChecking no' root@$masterip:/usr/lib/hadoop-
current/opt/apps/
echo "- copying hive - 2.0.1"
sshpass -f $masterpwdfile scp -r root@$masterip:/
usr/lib/hive-current/opt/apps/
echo "- copying spark - 2.1.1"
sshpass -f $masterpwdfile scp -r root@$masterip:/
usr/lib/spark-current/opt/apps/
echo "Start to link /usr/lib/${app}-current to
/opt/apps/${app}"
if [ -L /usr/lib/hadoop-current ]
then
    unlink /usr/lib/hadoop-current
fi
ln -s /opt/apps/hadoop-current /usr/lib/hadoop-
current
if [ -L /usr/lib/hive-current ]
then
    unlink /usr/lib/hive-current
fi
ln -s /opt/apps/hive-current /usr/lib/hive-
current
if [ -L /usr/lib/spark-current ]
then
    unlink /usr/lib/spark-current
fi
ln -s /opt/apps/spark-current /usr/lib/spark-
current
echo "Start to copy conf from $masterip to local
gateway (/etc/ecm)"
sshpass -f $masterpwdfile scp -r root@$masterip:/
etc/ecm/hadoop-conf /etc/ecm/hadoop-conf
sshpass -f $masterpwdfile scp -r root@$masterip:/
etc/ecm/hive-conf /etc/ecm/hive-conf
sshpass -f $masterpwdfile scp -r root@$masterip:/
etc/ecm/spark-conf /etc/ecm/spark-conf
echo "Start to copy environment from $masterip to
local gateway (/etc/profile.d)"
sshpass -f $masterpwdfile scp root@$masterip:/etc/
profile.d/hdfs.sh /etc/profile.d/
sshpass -f $masterpwdfile scp root@$masterip:/etc/
profile.d/yarn.sh /etc/profile.d/
sshpass -f $masterpwdfile scp root@$masterip:/etc/
profile.d/hive.sh /etc/profile.d/
sshpass -f $masterpwdfile scp root@$masterip:/etc/
profile.d/spark.sh /etc/profile.d/
if [ -L /usr/lib/jvm/java ]
then
    unlink /usr/lib/jvm/java
fi
echo "" >>/etc/profile.d/hdfs.sh
echo export JAVA_HOME=/usr/lib/jvm/jre-1.8.0
>>/etc/profile.d/hdfs.sh
echo "Start to copy host info from $masterip to
local gateway (/etc/hosts)"
sshpass -f $masterpwdfile scp root@$masterip:/etc/
hosts /etc/hosts_bak
cat /etc/hosts_bak | grep emr | grep cluster >>/etc
/hosts

```

```
if ! id hadoop >& / dev / null
then
    useradd hadoop
fi
```

- E-MapReduce versions earlier than V2.7 or V3.2

Create a script file, enter the following commands, and then run it on the Gateway node. Run the `sh deploy . sh < masteri_ip > master_password_file` command.

■ *deploy.sh*: the script name. The commands are as follows.

■ *masteri_ip*: the IP address of the master node in the cluster. Make sure that the IP address can be accessed.

■ *master_password_file*: the file that stores the password of the master node.

```
! / usr / bin / bash
if [ $# != 2 ]
then
    echo " Usage : $ 0 master_ip master_password_file "
    exit 1 ;
fi
masterip=$ 1
masterpassword_file=$ 2
if ! type sshpass >/ dev / null 2 >& 1 ; then
    yum install - y sshpass
fi
if ! type java >/ dev / null 2 >& 1 ; then
    yum install - y java - 1 . 8 . 0 - openjdk
fi
mkdir - p / opt / apps
mkdir - p / etc / emr
echo " Start to copy package from $ masterip to local gateway (/ opt / apps )"
echo " - copying hadoop - 2 . 7 . 2 "
sshpass - f $ masterpassword_file scp - r - o ' StrictHostKeyCheckin g no ' root @$ masterip : / usr / lib / hadoop - current / opt / apps /
echo " - copying hive - 2 . 0 . 1 "
sshpass - f $ masterpassword_file scp - r root @$ masterip : / usr / lib / hive - current / opt / apps /
echo " - copying spark - 2 . 1 . 1 "
sshpass - f $ masterpassword_file scp - r root @$ masterip : / usr / lib / spark - current / opt / apps /
echo " Start to link / usr / lib /\${ app }- current to / opt / apps /\${ app }"
if [ - L / usr / lib / hadoop - current ]
then
    Unlink / usr / lib / hadoop - Current
fi
ln - s / opt / apps / hadoop - current / usr / lib / hadoop - current
if [ - L / usr / lib / hive - current ]
then
    unlink / usr / lib / hive - current
fi
ln - s / opt / apps / hive - current / usr / lib / hive - current
```

```

if [ -L /usr/lib/spark-current ]
then
    unlink /usr/lib/spark-current
fi
ln -s /opt/apps/spark-current/usr/lib/spark-Current
echo "Start to copy conf from $masterip to local gateway (/etc/emr)"
sshpass -f $masterpwdfile scp -r root@$masterip:/etc/emr/hadoop-conf /etc/emr/hadoop-conf
sshpass -f $masterpwdfile scp -r root@$masterip:/etc/emr/hive-conf /etc/emr/hive-conf
sshpass -f $masterpwdfile scp -r root@$masterip:/etc/emr/spark-conf /etc/emr/spark-conf
Echo "start to copy environment from $masterip to local Gateway (/etc/profile.d)"
sshpass -f $masterpwdfile scp root@$masterip:/etc/profile.d/hadoop.sh /etc/profile.d/
if [ -L /usr/lib/jvm/java ]
then
    unlink /usr/lib/jvm/java
fi
ln -s /usr/lib/jvm/java-1.8.0-openjdk-1.8.0.131-3.b12.el7_3.x86_64/jre /usr/lib/jvm/java
echo "Start to copy host info from $masterip to local gateway (/etc/hosts)"
sshpass -f $masterpwdfile scp root@$masterip:/etc/hosts /etc/hosts_bak
cat /etc/hosts_bak | grep emr | grep cluster >>/etc/hosts
if ! id hadoop >& /dev/null
then
    useradd hadoop
fi

```

- Test

- Hive

```

[hadoop@iZ23bc05hr vZ ~]$ hive
hive> show databases ;
OK
default
Time taken : 1.124 seconds , Fetched : 1 row(s)
hive> create database school ;
OK
Time taken : 0.362 seconds
hive>

```

- Run the Hadoop job

```

[hadoop@iZ23bc05hr vZ ~]$ hadoop jar /usr/lib/hadoop-current/share/hadoop/mapreduce/hadoop-mapreduce-examples-2.6.0.jar pi 10 10
Number of Maps = 10
Samples per Map = 10
Wrote input for Map # 0
Wrote input for Map # 1
Wrote input for Map # 2
Wrote input for Map # 3
Wrote input for Map # 4

```

```

Wrote input for Map # 5
Wrote input for Map # 6
Wrote input for Map # 7
Wrote input for Map # 8
Wrote input for Map # 9
File Input Format Counters
  Bytes Read = 1180
File Output Format Counters
  Bytes Written = 97
Job Finished in 29.798 seconds
Estimated value of Pi is 3.2000000000 0000000000

```

2.12 Manage cluster resources

EMR (E-MapReduce) cluster resource management applies to scenarios that involve large-scale clusters with many tenants.

With EMR cluster resource management, the following features are supported.

- Different resource queues are used based on the users. This helps to isolate resources.
- Resource queues are elastic, which improves the efficiency of the cluster.



Note:

- Currently, you can choose [Capacity Scheduler](#) or [Fair Scheduler](#) for managing EMR cluster resources.
- Currently, you can only manage resources of EMR Hadoop clusters.

Enable resource queues

1. Log on to the [Alibaba Cloud E-MapReduce console](#). Click Cluster Management to go to the Cluster Management page.
2. Click Manage in the Actions column.
3. In the left-side navigation pane, click Cluster Resources.
4. Turn on Enable Resource Queue.



Notice:

After turning on Enable Resource Queue, on the Configure tab page of YARN, capacity-scheduler and fair-scheduler sections are frozen. On the Cluster Resources page, the configurations are synchronized. You need to turn off Enable Resource Queue on the Cluster Resources page to configure clusters by using XML again on the Configure tab page of YARN.

Configure resource queues

- After turning on Enable Resource Queue, choose fair scheduler or capacity scheduler as the scheduler type.
- Click Queue Settings to go to the Queue Settings page.
- Click Set Scheduler Defaults to set the default configurations of the scheduler.

Resource queues taking effect

After configuring resources queues, click Check Configuration to preview the configured XML file and click Deploy to deploy the configurations.



Note:

- After switching the scheduler type (for example, from Fair Scheduler to Capacity Scheduler) or modifying parameters related to resource preemption in Fair Scheduler, you need to restart all YARN components on the Clusters and Services page before clicking Deploy.
- After creating resource queues or setting related parameters, you only need to click Deploy.
- You can view the deployment progress of configurations on the Clusters and Services page.

2.13 Auto Scaling

2.13.1 Overview

This topic describes how to enable and disable the Auto Scaling feature.

You can use EMR Auto Scaling to reduce costs and improve execution efficiency for the following scenarios.

- You have a temporary need for adding nodes to enhance computing power.
- You need to add nodes based on some cluster metrics to ensure that important jobs are completed on schedule.



Note:

- Auto Scaling only supports scaling in and scaling out a cluster by adding or removing task nodes.
- Auto Scaling only supports Subscription and Pay-As-You-Go Hadoop clusters.

- We recommend that you choose Scale by Time as the rule type if you can specify the time to scale a cluster.
- We recommend that you choose Scale by Rule as the rule type if you cannot specify the time to scale a cluster and need to add and remove computing resources based on the specified YARN metrics.

Enable Auto scaling

1. Log on to the [Alibaba Cloud E-MapReduce console](#). Click the Cluster Management tab to go to the cluster management page.
2. Click Clusters and Services in the Actions column for the cluster that you want to perform auto scaling for.
3. In the left-side navigation pane, click Scaling.
4. In the top right corner of the page, click Enable Auto Scaling.

When using Auto Scaling through an Alibaba Cloud account for the first time, you need to assign an Auto Scaling default role to the EMR service.

5. Click OK on the authorization page of Auto Scaling.

Disable Auto Scaling

Click Disable Auto Scaling. The task nodes that have been added through Auto Scaling are to be released. HDFS data that is stored on core nodes is not affected.

Restrictions

You can only choose one scaling rule (Scale by Time or Scale by Rule) for Auto Scaling .

If you switch the scaling type, the original scaling rules are reserved but invalid and not to be triggered. Nodes that have been added for scaling out are not to be removed unless the rule for scaling in is triggered.

2.13.2 Configure Auto Scaling 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 is necessary to design a ratio for both modes based on the time needed. For example, if there are 8 peak business hours a day, the price for Pay-As-You-Go is essentially 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 Auto Scaling 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 Auto Scaling rule is less than this minimum number of nodes, the cluster scales based on the minimum number of nodes at the first execution.

For example, if the Auto Scaling 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

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

* Rule Name:

Rule names should not be repeated.

☐ Repeat ☒ Run Once

Specified a time

2018-11-12 15:16

* Retry Timeout(Seconds): Valid range: 0 to 21600 seconds

* Increase Task Nodes: Valid range: 1 to 100 nodes1 Nodes

* Cool-down Time(Valid range: 0 to 86400 seconds
Seconds):

- **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 detects that 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 operations are 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 Auto Scaling 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.

2.13.3 Preemptible instances in AutoScaling

E-MapReduce **preemptible instances** are suitable for scenarios where there is no requirement for the successful execution of big data jobs and where the cost

of computing resources plays an important role. By using Auto Scaling, you can purchase preemptible instances to increase the computing resources of your clusters.

Enable Auto Scaling

To enable Auto Scaling and set scaling rules, complete the following steps:

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. Click Cluster Management.
3. Find the cluster you want to add a preemptible instance to and click Manage.
4. In the navigation pane on the left, click **Scaling**.
5. Click **Enable Auto Scaling**.
6. Configure scaling rules. For more information, see [Configure Auto Scaling by time](#).
7. In the scaling configuration area, select **Preemptible instance**.

Configure a preemptible instance



Note:

Preemptible instances are more cost-effective than Pay-As-You-Go instances. However, Alibaba Cloud may release your preemptive instances at any time based on changes in demand resources or market rates.

To configure your preemptible instance, complete the following steps:

1. Select the vCPU and memory for your instance.
2. Select instance types. You can select up to three instance types. E-MapReduce filters out all other instance types to ensure that you purchase a preemptible instance that meets your requirements.
3. After you select the instance types, click the maximum price of each type and click OK. The instance types appear in the selected instances list. If you want to modify the price of a selected instance type, select the target one in the selected instances list and change the price (by hour). Your instance will run when your bid is higher than the current market rate. Your final instance type is billed at the market rate.
4. The system disk is used for deploying basic services such as the OS and EMR, which are set by default. You can set the data disk type and size according to your requirements.
5. The final configuration price includes the maximum bid price, system disk price, and data disk price. Click **Save**.

For more information about preemptible instances, see [FAQs](#).

2.13.4 Auto Scaling records

After completing an Auto Scaling operation, click the Scaling Records tab at the top of the Scaling page to see the records of the operation and the number of nodes.

Auto Scaling statuses comprise the following four types:

- **Running:** The operation is being implemented.
- **Success:** All of the specified nodes involved in the scaling rule have been added to or removed from the cluster.
- **Partial success:** Some nodes were successfully added to or removed from the cluster, but others failed due to the disk or ECS instance storage.
- **Failure:** No node was added to or removed from the cluster.

2.14 VPC

Virtual Private Cloud (VPC) helps you build an isolated network environment, including customizing the IP address range, network segment, routing table, and gateway.

For more information, see [What is VPC](#). VPC can be interconnected with physical IDC equipment rooms using [Express Connect](#).

Create a VPC cluster

When you create a cluster in E-MapReduce, you can select from two types of network: classic and VPC. If you select VPC, complete the following operations:

- **Subordinate VPC:** Select a VPC where the current E-MapReduce cluster is located. If you have not yet created a VPC, log on to the [VPC console](#) and create one.
- **VSwitch:** An ECS instance in the E-MapReduce cluster communicates through a VSwitch. If you have not yet created a VSwitch, log on to the [VPC console](#) and create one. Because a VSwitch has the properties of an availability zone, when you create a cluster in E-MapReduce, the VSwitch you create must also belong to the availability zone selected.
- **Security group:** The security group the cluster belongs to. Currently, only the security group of a VPC can be used, not the security group of a classic network. To ensure security, a security group created outside of E-MapReduce cannot be selected. Enter a security group name to create a security group.

Example

The following example shows how to enable Hive to access HBase clusters in E-MapReduce in different VPCs.

1. Create clusters.

Create two clusters in E-MapReduce. Hive cluster C1 is located in VPC1, whereas HBase cluster C2 is located in VPC2. Both clusters are located in the cn-hangzhou region.

2. Configure the high-speed channel.

For more information, see [Establish an intranet connection between VPCs under the same account](#). Select the same region.

3. Log on to the HBase cluster through SSH. Create a table through HBase Shell.

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

4. Log on to Hive through SSH.

- a. Modify the hosts and add the following line:

```
$ zk_ip emr-cluster // $ zk_ip is the zk node IP of Hbase cluster.
```

- b. Access HBase through Hive Shell.

```
hive > set hbase.zookeeper.quorum = 172.16.126.111, 172.16.126.112, 172.16.126.113;
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');
```

At this point, the `java.net.SocketTimeoutException` exception is reported. This is because the security group where the HBase cluster's ECS is located limits access to E-MapReduce at the related port. By default, security groups created by E-MapReduce only open port 22. Therefore, a security group rule must be

added to the HBase cluster's security group so as to open a port for the Hive cluster, as shown in the following figure.

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

2.15 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 " )
```

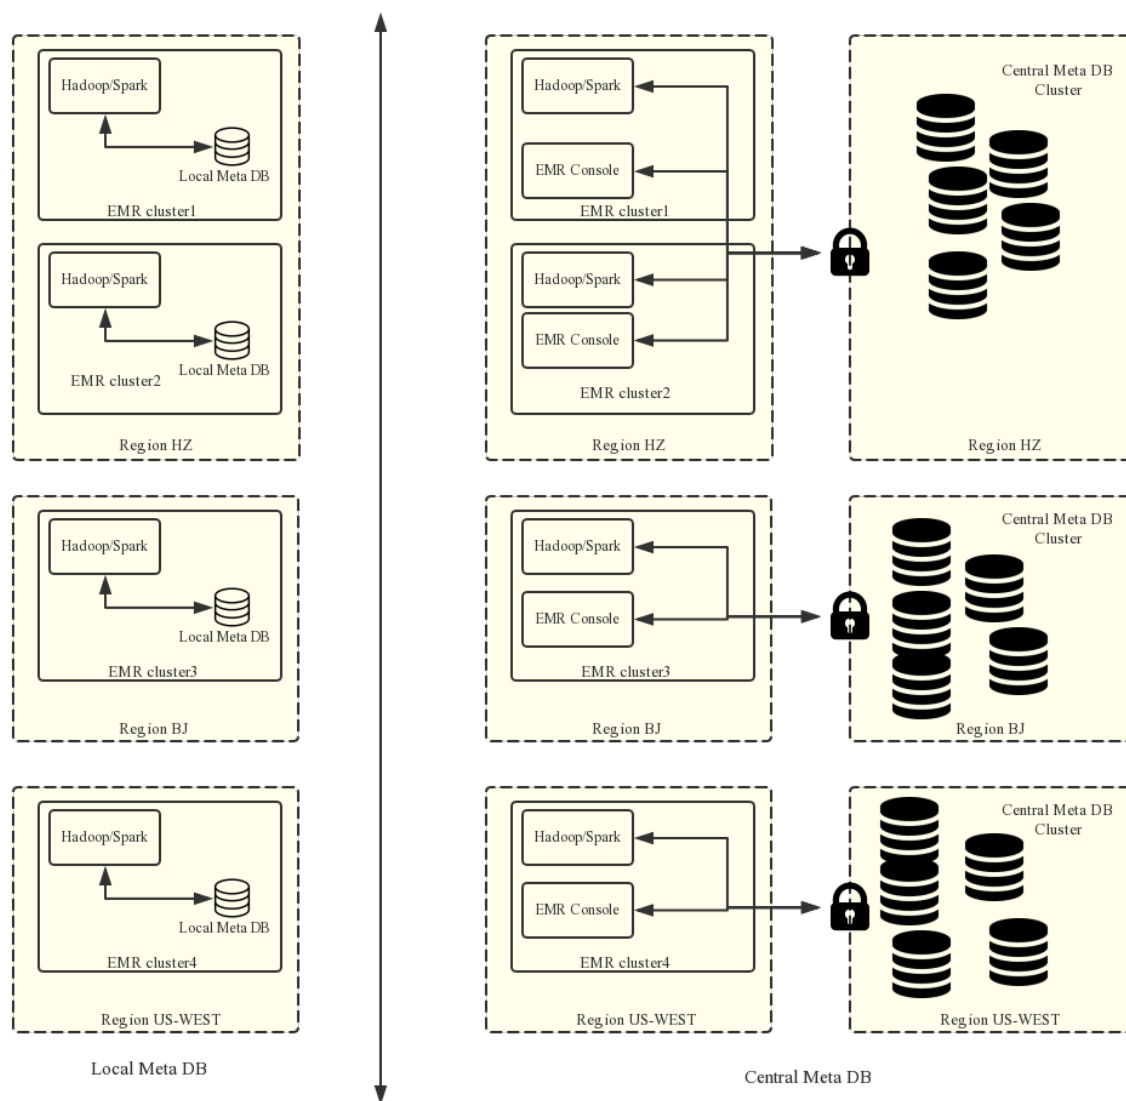
2.16 Metadata management

2.16.1 Manage Hive metadata

E-MapReduce has supported unified metadata management since V2.4.0. Before V2.4.0, clusters use local MySQL databases to store the metadata for Hive. E-

MapReduce V2.4.0 and later versions will support unified and high-reliability Hive metadatabases.

Overview



With unified metadata management, the following features are supported.

- Provides persistent metadata storage.

Before, metadata is stored in MySQL databases that are deployed on clusters.

Metadata is deleted when clusters are released. You can release a Pay-As-You-Go cluster after using it. For retaining the metadata, you need to log on to the cluster and export the metadata manually. Unified metadata management helps you avoid the issue.

- You can separate computing and storage.

E-MapReduce supports storing data in Alibaba Cloud OSS, which reduces the usage cost especially when the data volume is large. Meanwhile, E-MapReduce clusters are mainly used as computing resources and can be released anytime after use. Since data is stored in OSS, the metadata migration issue does not exist.

- Implement data sharing

With the central metadatabase, if all data is stored in OSS, all clusters can directly access data without migrating or restructuring metadata. This enables E-MapReduce clusters to provide different services while still ensuring direct data sharing.



Notice:

Before unified metadata management is supported, metadata is stored in the local MySQL database of each cluster and is lost when the cluster is released. With unified metadata management, releasing clusters does not clean up metadata. Before you delete the data in OSS or in the HDFS of a cluster or you release a cluster, make sure that the corresponding metadata is already deleted. That means the tables and database that store the data have been dropped. This prevents dirty metadata in the database.

Create a cluster using unified metadata

- Use webpages

On the Basic Settings page, turn on Unified Metabases.

- Use APIs

See the description of the [CreateClusterV2](#) operation. Set the value of the `useLocalMetaDb` parameter to false.

Operate metadata

- Prerequisites

The EMR console unified metadata management page only supports clusters that use unified metadata.

- Operate databases

You can search for databases by name. Click a database name to view all tables stored in the database.



Notice:

Before deleting a database, you need to delete all tables that are stored in the database.

- Operate tables

Table Name	Database Name	Location	Created At	Actions
call_center	tpcds_text_5	hdfs://emr-cluster/tmp/tpc...	Mar 28, 2019, 09:39:10	Search Details Delete
catalog_page	tpcds_text_5	hdfs://emr-cluster/tmp/tpc...	Mar 28, 2019, 09:39:11	Search Details Delete
catalog_returns	tpcds_text_5	hdfs://emr-cluster/tmp/tpc...	Mar 28, 2019, 09:39:11	Search Details Delete

Create a table: currently, only creating external tables and partition tables is supported. The separators can be common characters such as commas and spaces, special characters such as TAB, ^A, ^B, and ^C, and custom separators.

- Table details

Table Details

Basic Information

Advanced Information

Partition Information

Partition Name	Created At	Last Updated At
dt=20190501/country=China	May 31, 2019, 15:11:20	May 31, 2019, 15:11:20
dt=20190510/country=China	May 31, 2019, 15:11:30	May 31, 2019, 15:11:30

<1>

Total Items: 2

View partitions: For partition tables, you can go to Table Details > Partitions to view the partition lists. A maximum of 10,000 partitions can be displayed.



Notice:

- Databases and tables can be created and deleted only in E-MapReduce clusters.
- The HDFS is the internal file system of each cluster and does not support cross-cluster communication without special network settings. Therefore, the E-

MapReduce table management function only supports creating databases and tables based on the OSS file system.

- The location of a database or table must be in a directory under the OSS bucket , rather than the OSS bucket.

- **Metadatabase information**

Metadata management currently only supports managing metadata of clusters with "Unified Metabases" enabled.						
Metabase Name	Connect String	Capacity Usage	Total Capacity	Maximum Queries per Hour	Maximum Updates per Hour	
Hive Unified MetaStore	rm-bp1xmx4hx50hgv9q7o.m...	6.29%	512.0MB	14400000	7200000	

On the metadatabase information page, you can view the usage and restrictions of the current RDS instance. Submit a ticket if you need to modify the information.

Migration

- Migrate data from an EMR unified metadatabase to a user-created RDS instance
 1. Purchase an RDS instance and make sure the RDS instance can connect to the master node of the cluster. We recommend that you add the RDS instance to the security group of the ECS instances that are in the EMR cluster. By doing this, you can use a private address to connect to the RDS instance.



Note:

For security reasons, you need to add the IP addresses of the EMR nodes that are allowed to access the RDS instances to the whitelist.

2. Create a database in the RDS instance and name it hivemeta. Create a user and grant the user the read/write permission to hivemeta.
3. Export the unified metadata. We recommend that you stop the Hive metastore service before exporting the data to avoid the changes of metadata during the export progress.
4. Log on to the [Alibaba Cloud E-MapReduce console](#). Click the Cluster Management tab to go to the Cluster Management page.
- 5.
6. Click the cluster ID to go to the Clusters and Services page. On the Services list, click Hive > Configure to go to the Hive configuration page. Locate the values of the following configuration items. For earlier versions of EMR clusters that do

not support configuration management, locate the configuration items in the `$HIVE_CONF_DIR/hive-site.xml` file.

```
javax . jdo . option . Connection Username // the database
username ;
javax . jdo . option . Connection Password // the database
password ;
javax . jdo . option . Connection URL // the database
connection URL and database name ;
```

7. Run the following command to log on to the master node of the cluster using [SSH](#).

```
mysqldump -t hivemeta -h <unified metadataba se
connection URL> -u <unified metadataba se username>
-p > /tmp/metastore.sql
```

Enter the value of `javax.jdo.option.ConnectionPassword` as the password.

8. In the `usr / local / emr / emr - agent / run / meta_db_in fo . json` file that is stored on the master node of the cluster, set the value of `use_local_ meta_db` to false. Set the connection URL, username, and password of the RDS instance for the metadatabase.
9. On the Hive configuration page, set the connection URL, username, and password of the RDS instance for the metadatabase. For earlier versions of EMR clusters that do not support configuration management, locate the configuration items in the `$ HIVE_CONF_ DIR / hive - site . xml` file.
10. On the master node, set the connection URL, username, and password of the RDS instance for the metadatabase in the `hive - site . xml` file. Run the `init schema` command.

```
cd /usr/lib/hive-current/bin
./schematool -initSchema -dbType mysql
```

11. Import the exported metadata to the RDS instance. Use the CLI to log on to MySQL.

```
mysql -h {The url of rds} -u {RDS username} -p
```

In the MySQL CLI, run the `source /tmp/metastore.sql` file.

12. On the Clusters and Services page, click **RESTART ALL COMPONENTS** to restart all Hive components. Run the `hive cli` command on the master node to check whether the???

FAQ

- **Wrong FS:** `oss :// yourbucket / xxx / xxx / xxx` or **Wrong FS:** `hdfs :// yourhost : 9000 / xxx / xxx / xxx`

This error occurs when the table data on OSS is deleted while the metadata of the table is not. The table schema persists, while the actual data does not exist or is moved to another location. In this case, you can change the table location to an existing path and delete the table again.

```
alter table test set location 'oss :// your_bucke t /
your_folde r '
```

This can be completed on the E-MapReduce interactive console.

**Note:**

`oss :// your_bucke t / your_folde r` is required to an existing OSS path.

- The message “`java.lang.IllegalArgumentException: java.net.UnknownHostException: xxxxxxxx`” is displayed when the Hive database is deleted.

This error occurs because the Hive database is created in the HDFS of a cluster and it is not cleaned up when the cluster is released. As a result, its data in the HDFS of the released cluster cannot be accessed after a new cluster is created. Therefore, when releasing a cluster, remember to clean up the databases and tables that are manually created in the HDFS of the cluster.

Solutions:

Log on to the master node of the cluster using the command line, and find the address, username, and password for accessing the Hive metadatabase in \$

```
HIVE_CONF_ DIR / hive - site . xml .
```

```
javax . jdo . option . Connection Username // the database
username ;
javax . jdo . option . Connection Password // the database
password ;
```



```
javax .jdo . option . Connection URL // the database
connection URL and database name ;
```

```
<property>
  <name>javax.jdo.option.ConnectionUserName</name>
  <value>${ConnectionUserName}</value>
  <description>Username to use against metastore database</description>
</property>
<property>
  <name>javax.jdo.option.ConnectionPassword</name>
  <value>${ConnectionPassword}</value>
  <description>password to use against metastore database</description>
</property>
<property>
  <name>javax.jdo.option.ConnectionURL</name>
  <value>jdbc:mysql://${DBConnectionURL}/${DBName}?createDatabaseIfNotExist=true&amp;characterEncoding=UTF-8</value>
  <description>JDBC connect string for a JDBC metastore</description>
</property>
```

Connect to the Hive metadatabase on the master node.

```
mysql -h ${DBConnectionURL} -u ${ConnectionUserName} -
p [ Press Enter ]
[ Enter the password ]${ConnectionPassword}
```

After logging on to the Hive metadatabase, change its location to an existing OSS path in the region:

```
mysql> show databases;
+-----+
| Database |
+-----+
| information_schema |
| xxxxxxxxx77ac43c3bd0efae77e0bf1947d45fb4c896fb99 |
+-----+

mysql> use xxxxxxxxx77ac43c3bd0efae77e0bf1947d45fb4c896fb99;

mysql> select * from dbs;
+-----+-----+-----+-----+-----+-----+
| DB_ID | DESC | DB_LOCATION_URI | NAME | OWNER_NAME | OWNER_TYPE |
+-----+-----+-----+-----+-----+-----+
| 1 | Default Hive database | oss://mybucket/hive/warehouse | default | public | ROLE |
| 6 | NULL | hdfs://dirty-hostname/warehouse | dirty_db | NULL | USER |
+-----+-----+-----+-----+-----+-----+

mysql> update dbs set DB_LOCATION_URI = 'oss://your-bucket/your-db-folder' where DB_ID = 6;
```

2.16.2 Manage Kafka metadata

You can manage Kafka metadata in E-MapReduce V3.17.0 and later versions. Kafka metadata is automatically created for new Kafka clusters.

Currently, Kafka metadata management supports the following features.

- Creates and deletes topics.
- Views the description of topics.
- Modifies the partition number and configurations of topics.



Note:

Due to the callback methods, updates of topic lists and information are delayed.

Add a topic

1. Go to the Kafka data management page and click Add Topic in the upper-right corner.
2. Specify the topic name, the number of partitions, and the associated cluster. In the Advanced Settings section, click the + icon to add a configuration item.
3. After creating the topic, view the execution result in the Tasks list.

Delete a topic

You can search for topics by cluster ID and topic name. Note: you can only select active clusters as filters. For released clusters, only the latest update is displayed on the Topic list. You can operate active topics by clicking Delete or Update Configuration. View the execution results in the Tasks list.



Notice:

You cannot delete or update built-in topics such as `__consumer_offsets`.

View the description of a topic

Click Details for a topic to view the summary information, data distribution, and consumer groups of the topic.

- **Data distribution:** displays the status of partitions based on the number of partitions.
- **Consumer group:** an overview of all consumer groups related to the topic. Click the name of a consumer group to view the corresponding offsets and lags.

Modify the partition number and configurations of a topic

- **Modify multiple partitions at one time**
 1. Go to the Kafka Data Management page, click Batch Edit Partitions in the upper-right corner.
 2. In the dialog box that appears, click Select Topic and select multiple topics from a cluster and click OK.
 3. Set the values of the Partitions after Adjustment fields and click OK.
 4. View the execution status in the Tasks dialog box. If the operation is successful, you can view the changes of topic details in the console.

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.

Host

(MASTER) Subscription

Public Bandwidth:
8M

Hosts: 1
CPU: 4 Cores
Memory: 8GB


Data Disk Type: SSD
Disk: 80GB * 1 Disks

Core Instance Group

The specification upgrades are complete. Restart the servers for the upgrades to take effect.

ECS ID	Component Deployment Status	Public IP	Internal IP	Created At
i-bp16lmas113596omxwnd	Normal		192.168.1.191	Jan 24, 2019, 15:18:28
i-bp16lmas113596omxwne	Normal		192.168.1.192	Jan 24, 2019, 15:18:28

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.



ECS ID	Component Deployment Status	Public IP	Internal IP	Created At
i-bp161mas113596omxwnd	Normal		192.168.1.191	Jan 24, 2019, 15:18:28
i-bp161mas113596omxwne	Normal		192.168.1.192	Jan 24, 2019, 15:18:28

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 . nodemanage r . 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 . nodemanage r . resource . cpu - vcores = 32`. If the compute node is intended for mixed scenarios, you can set the value of `yarn . nodemanage r . 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 . nodemanage r . 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 . nodemanage r . 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 use 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: <ul style="list-style-type: none">· 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.



The screenshot displays the 'Core Instance Group' configuration in the Alibaba Cloud EMR console. A red box highlights a message: "The disk is expanded. Restart the node group for the change to take effect." The table below shows two ECS instances in a 'Normal' state.

ECS ID	Deployment Status	Public Network IP	Internal Network IP	Created At
i-b...	Normal		192.168.1.216	Aug 14, 2019, 11:23:37
i-b...	Normal		192.168.1.215	Aug 14, 2019, 11:23:37



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.

Restart Cluster
✕

Current Cluster: C- /

*The node group will be restarted. All services will be suspended during the restart process. Make sure that your business will not be affected before you restart the node.

☒ Rolling Restart ?
☒ Restart Scaled Nodes Only ?


OK
Cancel

The cluster restart settings are described as follows:

Parameter	Description
Rolling Start	<p>Rolling restart:</p> <ul style="list-style-type: none"> · 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	<p>A scaled node is a node that has its disks or configuration upgraded, such as a master or core node. Restart scaled nodes only:</p> <ul style="list-style-type: none"> · 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.

- 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.



ECS ID	Deployment Status	Public Network IP	Internal Network IP	Created At
i-bp1bsithym5hh9h93xxx	Normal			Aug 14, 2019, 11:23:37
i-bp1bsithym5hh9h93xxx	Normal			Aug 14, 2019, 11:23:37

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.

- Log on to the [Alibaba Cloud EMR console](#) and then go to the Cluster Management page.
- 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.
- 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.

- 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.

- 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_62](#) or [#unique_63](#).



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_64](#).



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_62](#) or [#unique_63](#).



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_65](#).



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 Convert your billing method

There are two billing methods available for E-MapReduce: Pay-As-You-Go and Subscription. We recommend that you use Pay-As-You-Go, as it is more cost-effective. If you are sure that you want to use E-MapReduce for a longer period of time, you can change to the Subscription method.

Convert Pay-As-You-Go to Subscription

You can select the Pay-As-You-Go method at the beginning during your E-MapReduce trial and at any time after that. In the Cluster Overview panel, click the Switch to Subscription button to switch your payment method from Pay-As-You-Go to Subscription. The payment method for the entire cluster will be converted.

Convert Subscription to Pay-As-You-Go

You cannot currently convert from Subscription to Pay-As-You-Go.

3.4 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.

3.5 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.

Scale Up/Out



Scale Up/Out

CORE (Core Instance Group) **TASK (Task Instance Group)**

Add Host Group

Host Group Name:

Task Instance Type: 4 vCPU 16GB ecs.sn2.large

System Disk Type: ☒ SSD Disk ☐ Ultra Disk

System Disk Size: 120 GB * Disk IOPS 5400

Data Disk Type: ☒ SSD Disk ☐ Ultra Disk

Data Disk Size: 80 GB * 4 Disks IOPS 4200

Task Instances: 1 Instances

Billing Method: ☒ Pay-As-You-Go

**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 . nodemanage r . 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

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.

Clusters and Services

Status			
Services			
Normal	HDFS	!	Actions ▼
Normal	YARN		Actions ▼
Normal	Hive		Actions ▼
Normal	Ganglia		Actions ▼
Normal	Spark		Actions ▼
Normal	Hue		Actions ▼
Normal	Tez		Actions ▼
Normal	Sqoop		Actions ▼
Normal	Pig		Actions ▼
Normal	HAProxy		Actions ▼
Normal	ApacheDS		Actions ▼
Normal	Knox		Actions ▼

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

Quick Configuration

Configuration Type: BASIC | ADVANCED | INFORMATION | DATA_PATH | LOG_PATH

LOG | JVM | DATA | DATABASE | PERFORMANCE | TIME

CODEC | OSS | PORT | MEMORY | DISK | NETWORK | PATH

URI

Configuration File: core-site | hdfs-site | hadoop-env | httpfs-site | kms-site

httpfs-env

Service Configuration: Default Configuration

Deploy Client Config Save

Custom Configuration

hdfs-site

dfs.datanode.cache.revocation.timeout.ms 900000

dfs.namenode.resource.check.interval 5000

dfs.client.https.need-auth false

dfs.replication 2

dfs.datanode.available-space-volume-choosing-policy.balanced-space-threshold 10737418240

dfs.namenode.avoid.read.stale.datanode false

nfs.kerberos.principal

dfs.hosts.exclude /etc/ecm/hadoop-conf/dfs.exclude

dfs.namenode.replication.min 1

dfs.datanode.directoryscan.threads 1

dfs.datanode.directoryscan.interval 21600

dfs.namenode.fs-limits.min-block-size 1048576

dfs.namenode.acls.enabled false

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.

Add Configuration Item

* Key:

* Value:

Description:


OK Cancel

- 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



The image shows a 'Run Action' dialog box with a title bar and a close button (X). The dialog contains the following fields and controls:

- ActionName: CONFIGURE ALL COMPONENTS
- Service Name: HDFS
- Component Name: All Components
- Execution Scope: ☒ All Matched (Total 3) ☐ Designated
- Execute by Scrolling: ☐ Scrolling Execute Time Interval Seconds
- Failure Handling Strategy: Single Node Failed Terminate ... (dropdown menu)
- Execution Reason: (text area)
- Buttons: OK, Cancel

- 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 [Deploy configurations](#) section for making a configuration rollback take effect.