

Alibaba Cloud E-MapReduce

Cluster Planning and Configurations

Issue: 20190417

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>swich {stand slave}</code>

Contents

Legal disclaimer.....	I
Generic conventions.....	I
1 Cluster planning.....	1
1.1 User management.....	1
1.2 Instance types.....	2
1.3 Gateway clusters.....	3
1.4 ECS instances.....	3
1.5 Storage guide.....	5
1.6 D1 series.....	7
1.7 Enable access between classic networks and VPCs.....	8
1.8 Disaster recovery.....	8
1.8.1 Disaster recovery in E-MapReduce clusters.....	8
2 Configure clusters.....	10
2.1 Create a cluster.....	10
2.2 Cluster list and details.....	16
2.3 Cluster details.....	20
2.4 Expand a cluster.....	23
2.5 Release a cluster.....	25
2.6 Cluster renewal.....	26
2.7 Service list.....	28
2.8 Cluster scripts.....	29
2.9 Access links and ports.....	31
2.10 Security groups.....	31
2.11 Create a gateway.....	33
2.12 Auto Scaling.....	39
2.12.1 Overview.....	39
2.12.2 Configure Auto Scaling by time.....	40
2.12.3 Scale By Rule.....	42
2.12.4 Preemptible instances in Auto Scaling.....	45
2.12.5 Auto Scaling records.....	46
2.13 VPC.....	46
2.14 MetaService.....	48
2.15 Metadata management.....	50
2.15.1 Table management.....	51
3 Modify configurations.....	56
3.1 Node upgrade.....	56
3.2 Disk expansion.....	59
3.3 Convert your billing method.....	62
3.4 Remove abnormal nodes.....	63
3.5 Multiple host groups.....	63

4 Third-party software.....	65
4.1 Bootstrap actions.....	65
4.2 Software configuration.....	67
5 Workflow development.....	70
5.1 Manage a workflow project.....	70
5.2 Job operations.....	72
5.3 Ad hoc queries.....	74
5.4 Manage a workflow.....	77
5.5 Jobs.....	79
5.5.1 Configure a Hadoop MapReduce job.....	79
5.5.2 Configure a Hive job.....	81
5.5.3 Configure a Pig job.....	82
5.5.4 Configure a Spark job.....	85
5.5.5 Configure a Spark SQL.....	87
5.5.6 Configure a Shell job.....	88
5.5.7 Configure a Sqoop job.....	89
5.5.8 Job operations.....	90
5.5.9 Time and date variables.....	91
5.6 Old EMR Scheduling (Soon will be unavailable).....	92
5.6.1 Notebooks.....	92
5.6.1.1 Introduction.....	92
5.6.1.2 Operations.....	99
5.6.1.3 Examples.....	108
5.6.1.3.1 Query bank employee information.....	108
5.6.1.3.2 Video playback data.....	108
5.6.2 Execution plans.....	110
5.6.2.1 Create an execution plan.....	110
5.6.2.2 Manage an execution plan.....	114
5.6.2.3 Execution plan list.....	115
5.6.2.4 View job results and logs.....	116
5.6.2.5 Parallel execution of multiple execution plans.....	118

1 Cluster planning

1.1 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 [Step 2](#).

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.

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.

1.2 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 ResourceManager services are deployed on master instances, while the HDFS DataNode and YARN NodeManager services are deployed on core instances. For task instances, because they are only used in computing tasks, only YARN NodeManager 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.

1.3 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.

1.4 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.

1.5 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 recommend 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.

1.6 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.

1.7 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 cloud network: 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.

1.8 Disaster recovery

1.8.1 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 [Role authorization](#).
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. Configure the software.
 - **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.

- **High security 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.
- **Enable custom setting:** Before you start a cluster, you can specify a JSON file to change the software configuration.

2. Configure the hardware.

- Billing method
 - Like with ECS, both Subscription and Pay-As-You-Go modes are supported. If you select Subscription mode, you must also select the duration. You can select 1, 2, 3, 6, or 9 months, or 1, 2, or 3 years. This mode is applicable to short-term testing or flexible dynamic tasks, but is relatively expensive.
- Cluster network configuration
 - 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 [VPC](#).

- 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. Configure the basic information.

- 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.
- Running logs
 - Running logs: The function for saving running logs is turned on by default. In the default state, you can select the OSS directory as a location to save running logs, but you must have activated OSS before using this function. The cost depends on the number of uploaded files. We recommend that you open the OSS log saving function, which helps in debugging and error screening.
 - Log path: OSS path for saving logs.
 - Uniform Meta Database: This is provided by E-MapReduce to store all Hive metadata in the external database of the cluster. We recommend that you use this function when the cluster uses OSS as the main storage.
- Permission settings
 - 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.
- 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](#)
 - Logon 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.
- (Optional) Bootstrap actions: Before Hadoop is enabled in the cluster, you can run the customized script. For more information, see [Bootstrap actions](#).

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.

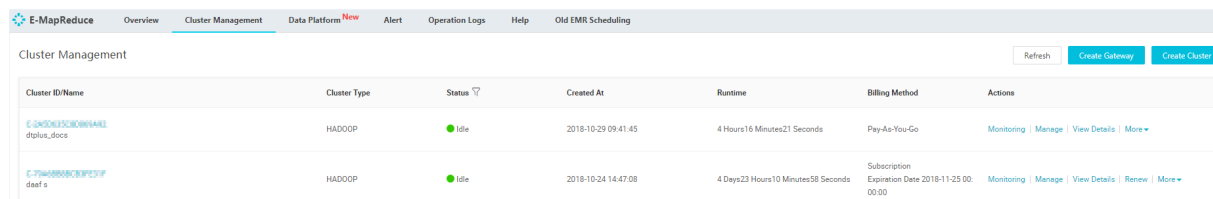
You do not need to perform any additional operations because the corresponding computing resources are not created. The cluster is automatically hidden after three days.

2.2 Cluster list and details

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:



Cluster ID/Name	Cluster Type	Status	Created At	Runtime	Billing Method	Actions
C-20181029094145-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
C-20181024144708-deaf_s	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 Renew 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:** Hadoop is the only cluster type available.
- **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: E-4T8BC7TTT12JFC071 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.
- **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 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-bp18guc2h0w50q3550qnd VPC/VSwitch: vpc-bp18guc2h0w50q3550qnd / vsw-bp18guc2h0w50q3550qnd

- **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.

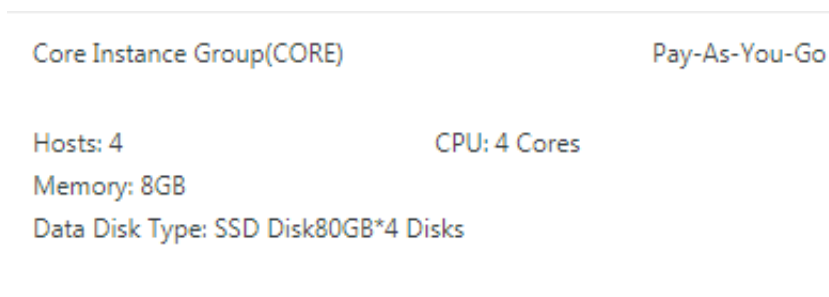


- **Hosts:** The number of current nodes. During the creation process, the number of current nodes is less than the number of nodes you applied for until the creation is complete.
- **CPU:** The number of cores in a node's CPU.
- **Memory:** Memory capacity of a node.
- **Data Disk Type:** Data disk 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-ko1-epp42h3j8lncvvcy0	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-bp1-4n9gnd4b3z6m8k4nqg	Normal		192.168.0.127	2018-10-25 10:36:59
i-bp1-7x3fagw9h4necv83n7	Normal		192.168.0.128	2018-10-25 10:37:00
i-bp1-6b0tqpmw834n9n8n32	Normal		192.168.0.129	2018-10-25 10:53:07
i-bp1-6b0tqpmw834n9n8n32	Normal		192.168.0.131	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 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-4t1m2c7t1t128f074 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.
- **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.

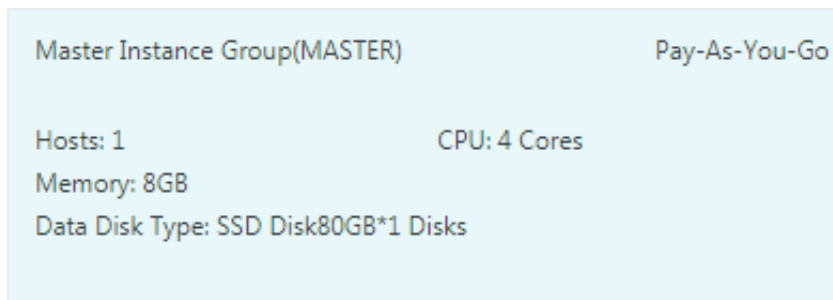
Network

Network
Region ID: cn-hangzhou-f
Network Type: vpc
Security Group ID: sg-bp18guc1v0t0g3q553q4ll
VPC/VSwitch: vpc-bp18guc1v0t0g3q553q4ll / vsw-bp18guc1v0t0g3q553q4ll


- **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-bp1-4ep4t9n0j0gwyv0	● Normal	47.98.199.51	192.168.0.45	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.


Core Instance Group(CORE)
Pay-As-You-Go

Hosts: 4 CPU: 4 Cores

Memory: 8GB

Data Disk Type: SSD Disk80GB*4 Disks

- **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-hang7-hangfushu3-2-80M4uagp	● Normal		192.168.0.127	2018-10-25 10:36:59
i-hang7-2733hangw9-10-80M4uagp	● Normal		192.168.0.128	2018-10-25 10:37:00
i-hang7-2733hangw9-10-80M4uagp	● Normal		192.168.0.129	2018-10-25 10:53:07
i-hang7-2733hangw9-10-80M4uagp	● Normal		192.168.0.130	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.

2.4 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

Select the cluster you want to expand from the list of clusters, click **More**, and select **Scale Up/Out**. You can also click **View Details** to the right of the cluster, and click **Scale Up/Out** in the upper right corner.

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-bp1c0d1hnoaicxddd57t	● 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 to the master host with SSH using the following command and obtain the public IP of the master cluster in the [Cluster Overview](#) 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 [Cluster Overview](#) panel.

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

4. Change your root password using the following command.

```
sudo passwd root
```

2.5 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.

2.6 Cluster renewal

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

Enter the renewal page

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

Renewal page

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

<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.


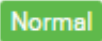


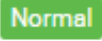
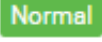
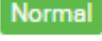
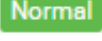
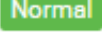
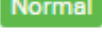
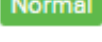
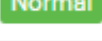
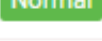
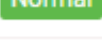
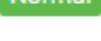
If an expired cluster is successfully renewed, its expiration date is updated to reflect the renewal. If an expired ECS instance is successfully renewed, its expiration date is usually updated around three to five minutes later.

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.

2.7 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 View 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/clusterScriptId`. If the cluster is configured with an OSS log directory, the running log is also uploaded to `osslogpath/clusterId/ip/cluster-scripts/clusterScriptId`.

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
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>
```

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.

[DO NOT TRANSLATE]

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 server in the same intranet as an E-MapReduce cluster. You can use gateways for load balancing and security isolation or to submit jobs to E-MapReduce clusters.

You can create a gateway in the following two ways:

- In [E-MapReduce](#) (recommended).
- Manually.

Create a gateway on the E-MapReduce console

E-MapReduce gateways only support Hadoop clusters. Before you create an E-MapReduce gateway, you first need to create a Hadoop cluster. To create an E-MapReduce gateway, complete the following steps.

1. Log on to the [E-MapReduce console](#) .
2. Click Create Gateway.

3. Enter the required information on the Create Gateway page.

- **Billing Method:**
 - **Subscription:** Charges for a specified period of time. This method is more cost-effective than Pay-As-You-Go, especially when you pay for three years at a time, as you get a larger discount.
 - **Pay-As-You-Go:** Charges by the hour. This method calculates fees based on the number of hours that you use the product.
- **Cluster:** Create a gateway that the gateway can submit jobs to. Gateways automatically configure the Hadoop environment that is consistent with the cluster.
- **Configuration:** The available ECS instance specifications in the zone.
- **System Disk Type:** The system disk type of the gateway node. There are two types of system disk: SSD cloud disk and efficient cloud disk. The displayed type varies according to the server model and region. By default, the system disk is released when the cluster is released.
- **System Disk Size:** The minimum size is 40 GB and the maximum is 500 GB. The default value is 300 GB.
- **Data Disk Type:** The data disk type of the gateway node. There are two types of data disk: SSD cloud disk and efficient cloud disk. The displayed type varies according to the server model and region. By default, the data disk is released when the cluster is released.
- **Data Disk Size:** The minimum size is 200 GB and the maximum is 4000 GB. The default value is 300 GB.
- **Quantity:** The number of data disks. The minimum is 1 and the maximum is 10.
- **Cluster Name:** The name of a gateway. It can contain between 1 and 64 characters. Only Chinese characters, letters, numbers, hyphens (-), and underscores (_) are allowed.
- **Password/Key Pair:**
 - **Password Mode:** Enter the password for logging on to the gateway in the text box.
 - **Key Pair Mode:** Select the key pair name for logging on to the gateway in the drop-down menu. If no key pair has been created yet, click Create Key Pair on the right to go to the ECS console. Do not disclose private key files in the .pem format that correspond to the key pair. After the gateway is created,

the public key of the key pair is automatically bound to the ECS where the gateway is located. When you log on to the gateway through SSH, enter the private key in the private key file.

4. Click Create to save the configurations.

If the creation is successful, the newly created gateway is displayed in the cluster list and the status in the Status column becomes Idle.

Create a gateway manually

- Network environment

Make sure that the gateway machine is in the security group of the corresponding E-MapReduce cluster. This allows the gateway nodes to have easy access to the E-MapReduce cluster. For more information about setting the security group of a machine, see [Create a security group](#).

- Software environment

- System environment: CentOS 7.2+ is recommended.
- Java environment: JDK 1.7 or later must be installed. OpenJDK 1.8.0 is recommended.

- Procedure

- E-MapReduce 2.7 or later, 3.2 or later

To create a gateway in these versions, we recommend that you use the E-MapReduce console.

If you want to set up a gateway manually, copy the following script to the gateway host and run it: `sh deploy . sh < master_ip > master_password_file .`

■ `deploy.sh` is the script name.

■ `master_ip` is the IP address of the master node in the cluster, which needs to be accessible.

■ `master_password_file` is the file for storing the password of the master node, which is written in the file.

```
#!/usr/bin/bash
If [ $ #! = 2 ]
then
    echo " Usage : $ 0  master_ip  master_pas  sword_file "
    exit 1 ;
fi
```

```

masterip=$ 1
masterpwdf ile=$ 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 $ masterpwdf ile scp -r -o ' StrictHost
KeyCheckin g no ' root @$ masterip :/ usr / lib / hadoop -
current / opt / apps /
echo " - copying hive - 2 . 0 . 1 "
sshpass -f $ masterpwdf ile scp -r root @$ masterip :/
usr / lib / hive - current / opt / apps /
echo " - copying spark - 2 . 1 . 1 "
sshpass -f $ masterpwdf ile 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 $ masterpwdf ile scp -r root @$ masterip :/
etc / ecm / hadoop - conf / etc / ecm / hadoop - conf
sshpass -f $ masterpwdf ile scp -r root @$ masterip :/
etc / ecm / hive - conf / etc / ecm / hive - conf
sshpass -f $ masterpwdf ile scp -r root @$ masterip :/
etc / ecm / spark - conf / etc / ecm / spark - conf
echo " Start to copy environmen t from $ masterip to
local gateway (/ etc / profile . d )"
sshpass -f $ masterpwdf ile scp root @$ masterip :/ etc /
profile . d / hdfs . sh / etc / profile . d /
sshpass -f $ masterpwdf ile scp root @$ masterip :/ etc /
profile . d / yarn . sh / etc / profile . d /
sshpass -f $ masterpwdf ile scp root @$ masterip :/ etc /
profile . d / hive . sh / etc / profile . d /
sshpass -f $ masterpwdf ile 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 / hdf s . sh
echo export JAVA_HOME =/ usr / lib / jvm / jre - 1 . 8 . 0
>>/ etc / profile . d / hdf s . sh
echo " Start to copy host info from $ masterip to
local gateway (/ etc / hosts )"
sshpass - f $ masterpwdf ile 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 2.7 or earlier, 3.2 or earlier

Copy the following script to the gateway host and run it: `sh deploy . sh < masteri_ip > master_pas sword_file .`

■ `deploy.sh` is the script name.

■ `masteri_ip` is the IP address of the master node in the cluster, which needs to be accessible.

■ `master_password_file` is the file for storing the password of the master node, which is written in the file.

```

! / usr / bin / bash
if [ $# != 2 ]
then
    echo " Usage : $ 0 master_ip master_pas sword_file "
    exit 1 ;
fi
masterip =$ 1
masterpwdf ile =$ 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 $ masterpwdf ile scp - r - o ' StrictHost
KeyCheckin g no ' root @$ masterip :/ usr / lib / hadoop -
current / opt / apps /
echo " - copying hive - 2 . 0 . 1 "
sshpass - f $ masterpwdf ile scp - r root @$ masterip :/
usr / lib / hive - current / opt / apps /
echo " - copying spark - 2 . 1 . 1 "
sshpass - f $ masterpwdf ile 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-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 Auto Scaling

2.12.1 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. [DO NOT TRANSLATE]

2.12.2 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(Seconds): Valid range: 0 to 86400 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 to 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.

2.12.3 Scale By Rule

When you use an EMR Hadoop cluster, choose Scale By Rule as the rule type if you cannot predict the peak or valley periods for big data computing.

Configure the number of instances in a scaling group

- **Maximum Instances:** The maximum number of task nodes in a scaling group. Once the maximum number is exceeded, auto scaling will not be performed even if the

scaling rule is met. Currently, the maximum number of instances that you can set is 1,000.

- **Minimum Instances:** The minimum number of task nodes of a scaling group.

Configure a scaling rule

Scaling rules include scaling out rules and scaling in rules. After you disable the auto scaling feature for a cluster, all rules are deleted. You need to configure scaling rules again when you reactivate the auto scaling feature. After you switch the rule type, for example, switching from Scale by Rule to Scale By Time, the scaling rules based on the original rule type are canceled and are not to be triggered. Instances that have been added for scaling out are not to be released.

- **Rule Name:** A rule name is required to be unique for each cluster. This requirement applies to rules both for scaling out and scaling in.
- **Cluster Metric:** Load metrics of Hadoop YARN. For more information, see [Hadoop documentation](#).

E-MapReduce Auto Scaling Metrics	YARN Metrics	Description
YARN.AvailableVCores	availableVirtualCores	The number of available virtual cores
YARN.PendingVCores	pendingVirtualCores	The number of pending virtual cores, EMR collects supplementary metrics.
YARN.AllocatedVCores	allocatedVirtualCores	The number of allocated virtual cores
YARN.ReservedVCores	reservedVirtualCores	The number of reserved virtual cores
YARN.AvailableMemory	availableMB	The amount of memory available in MB
YARN.PendingMemory	pendingMB	The amount of memory pending in MB, EMR collects supplementary metrics.
YARN.AllocatedMemory	allocatedMB	The amount of memory allocated in MB

E-MapReduce Auto Scaling Metrics	YARN Metrics	Description
YARN.ReservedMemory	reservedMB	The amount of memory reserved in MB
YARN.AppsRunning	appsRunning	The number of applications running
YARN.AppsPending	appsPending	The number of applications pending
YARN.AppsKilled	appsKilled	The number of applications killed
YARN.AppsFailed	appsFailed	The number of applications failed
YARN.AppsCompleted	appsCompleted	The number of applications completed
YARN.AppsSubmitted	appsSubmitted	The number of applications submitted
YARN.AllocatedContainers	containersAllocated	The number of containers allocated
YARN.PendingContainers	containersPending	The number of containers pending
YARN.ReservedContainers	containersReserved	The number of containers reserved

- **Evaluation Period and Scaling Condition:** Once the aggregate of the Cluster Metric's values meets the scaling condition during an evaluation period, the number of occurrences increases by one.
- **Occurrences:** Auto scaling is performed when the number of occurrences reaches the set value of Occurrences.
- **Add/Remove Nodes:** The number of task nodes to add or to remove each time a cluster is scaled.
- **Cooldown (Seconds):** The minimum interval time that is allowed between auto scaling. Auto-scaling is not to be performed during a cooldown even if the scaling condition is met.

Configure the specifications of nodes

You can specify the hardware specification of the nodes that are used to scale in or scale out a cluster. Click **Enable Auto-scaling** to configure the specifications.

You cannot modify the configurations after clicking Save. For applying new configurations, click Disable Auto-scaling and click Enable Auto-scaling again.

- Instance types are displayed in the Available Instances list based on the specifications of the vCPU and memory you select for a node. Select the instance types that you want from the Available Instances list and move the instance types to the Selected Instances list on the right. By doing this, you have confirmed the instance types used to scale your cluster.
- You can select a maximum of three instance types. This helps to prevent auto scaling failures due to insufficient ECS resources.
- The minimum size of a data disk is 40 GB, and it does not matter whether you choose the Ultra Disk or the SSD Disk as the data disk type.

2.12.4 Preemptible instances in Auto Scaling

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](#).

2.12.5 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.

2.13 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

2.14 MetaService

MetaService allows you to access Alibaba Cloud resources in the E-MapReduce cluster without using an AccessKey (AK).

Default roles

When creating a cluster, you must authorize an application role (AliyunEmrEcsDefaultRole) to E-MapReduce. After you do so, you can perform operations on E-MapReduce to access Alibaba Cloud resources without

using an AK. By default, the following permission policies are granted to AliyunEmrEcsDefaultRole:

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

By default, operations based on MetaService can only access OSS data. If you want to use MetaService to access other Alibaba Cloud resources, such as LogService, you must grant permissions to AliyunEmrEcsDefaultRole. Perform the preceding operations on the [RAM console](#).



Notice:

MetaService only supports AK-free operations on OSS, LogService, and MNS data. Modify and delete the default role with caution. Otherwise, you may fail to create or perform operations on clusters.

Custom application roles

When creating a cluster, you can use a default role or create your own application role. In most cases, you only need to use or modify the default role. For more information about how to create and authorize a role to E-MapReduce, see [RAM](#).

Access to MetaService

MetaService is an HTTP service that can be accessed directly to obtain metadata. For example, by using the `curl http://localhost:10011/cluster-region` command, you can obtain the region where the current cluster is located.

MetaService supports the following types of information:

- Region: /cluster-region
- Role name: /cluster-role-name

- AccessKeyId: /role-access-key-id
- AccessKeySecret: /role-access-key-secret
- SecurityToken: /role-security-token
- Network type: /cluster-network-type

Use MetaService

You can use MetaService to access Alibaba Cloud resources without using an AK. This has the following advantages:

- Reduces the risk of an AK leak. The use of RAM also minimizes security risks, as only the required permissions are granted to the role.
- Improves user experience. For instance, when you access OSS resources interactively, you no longer need to write a long string of OSS paths.

The usage methods are as follows:


```
I . Using the Hadoop command line to display OSS
data
    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 . Using Hive to create a table
    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 " )
```

2.15 Metadata management

2.15.1 Table management

E-MapReduce 2.4.0 and later support the central management of metadata in the highly reliable Hive metastore. In earlier versions, clusters use the local MySQL database.

Introduction



When you create a cluster, you can enable the central metastore function so that the cluster uses an external metastore.



Note:

- The current metastore needs to be connected using the public IP address. The cluster must therefore have a public IP address. Do not change this IP address, otherwise the corresponding database whitelist becomes invalid.
- The table management function can only be used if the central metastore function was enabled when a cluster was created. Local metastores do not currently support table management. If you have a local metastore, use the Hue tool in the cluster to manage your tables.

The central metadata management function performs the following:

1. Provides long-term metadata storage.

If metadata is stored in the local MySQL database of the cluster, it is lost when the cluster is released. Clusters can be created and released at any time as required, especially when E-MapReduce supports flexible creation. To retain metadata, you must log on to the cluster and manually export the data. This issue is resolved with the central metadata management function.

2. Separates computing and storage.

E-MapReduce supports storing data in Alibaba Cloud OSS, which reduces usage costs, especially when dealing with high volumes of data. Meanwhile, E-MapReduce clusters are mainly used as computing resources and can be released at any time after use. Since data is stored in OSS, problems with metadata migration are non-existent.

3. Implements data sharing.

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



Notice:

Before central metadata management is supported, metadata is stored in the local MySQL database of each cluster and is lost when the cluster is released. With central 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.

Table management operations

Before E-MapReduce clusters can support metadata management, you have to first log on to the internal environment of a cluster to check, add, or delete tables. If more than one cluster exists, you have to log on to each, one by one. This is inconvenient. With the central metadata management function, E-MapReduce enables table management on the console. This includes checking the list of databases and tables, checking table details, creating and deleting databases and tables, and previewing data.

- Create a database or table list
- View table details
- Preview data
- Create a database
- Create a table

There are two ways of creating a table: manually and from a file.

- **Manual creation:** If no service data exists, you can manually input the table structure to create an empty table.
- **Creation from a file:** If service data already exists, you can use it as a table by parsing the table interface from the file. Make sure that the separators used for creating a table correspond to those used in the data file. This guarantees a proper table structure.

The separators can be common characters such as commas and spaces, or special characters TAB, ^A, ^B, and ^C.



Notice:

1. Databases and tables can only be created and deleted in E-MapReduce clusters.
2. 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.
3. The location of a database or table must be in a directory under the OSS bucket, rather than the OSS bucket.

Common issues

1. Wrong FS: `oss :// yourbucket / xxx / xxx / xxx .`

This error occurs when the table data on OSS is deleted but the table metadata is not. The table schema continues to exist, but the actual data does not 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_bucket /
your_folder '
```

You can complete this operation on the E-MapReduce interactive console.



Note:

`oss :// your_bucket / your_folder` must be an existing OSS path.

2. Wrong FS: `hdfs://yourhost:9000/xxx/xxx/xxx.`

This error occurs when the table data in the HDFS is deleted but the table schema is not. The error can be removed by following the preceding solution.

3. 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 is not cleaned 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 the databases and tables that are manually created in the HDFS of the cluster.

To resolve this problem, log on to the master node of the cluster using the command line, and find the address, user name, and password used to access the Hive metastore in `$ HIVE_CONF_ DIR / hive - site . xml`.

```
javac . jdo . option . Connection Username // Username for
accessing the database ;
javac . jdo . option . Connection Password // Password for
accessing the database ;
```



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

```
<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>
```

Log on to the Hive metastore on the master node of the cluster:

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

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

```
mysql> show databases;
+-----+
| Database |
+-----+
| information_schema |
| xxxxxxxxxx77ac43c3bd0efae77e0bf1947d45fb4c896fb99 |
+-----+

mysql> use xxxxxxxxxx77ac43c3bd0efae77e0bf1947d45fb4c896fb99;

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;
```

3 Modify configurations

3.1 Node upgrade

In real scenarios, the CPU or memory of a cluster node, especially master nodes, may be insufficient.



Note:

- Only E-MapReduce Subscription clusters can be upgraded.
- Ephemeral disk instances (such as instance types of D1 and I2) do not support upgrade configurations. You can only add nodes.
- Non-ephemeral disk instances only support upgrade configurations, and do not support downgrade configurations.

Procedure

1. In the Cluster Management panel, select a cluster, and click View Details to go to the Cluster Overview pane.
2. In the upper-right corner, click Renew & Upgrade > Upgrade.
3. Configure the nodes that you want to upgrade.
4. Click OK.
5. Pay for your order.
6. Return to the Cluster Management page, refresh the page to make sure that the node configuration has become the target specification. The following figure displays the upgraded node information.

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-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

7. Click The specification upgrades are complete. Restart the server for the upgrades to take effect to view to the following figure.

Restart Cluster

Cluster: C-HPHNR2TQCLFANGLA / T10PABWY


The server will be restarted. All services will be suspended during the restart. Make sure you restart your server during non-business hours to avoid service interruption.

☒ Rolling Start ☒ Only Restart Updated Nodes

OK

Cancel

8. Click OK.

 Notice:

- If you restart a cluster, you also restart ECS instances. ECS instance data is unavailable during restarting.
- Rolling Start
 - If selected, only ECS instances you select is restarted. Instances are restarted one after another. Each restart takes about five minutes.
 - If not selected, all ECS instances are restarted.
- Only Restart Updated Nodes
 - Configuration-updated nodes are nodes whose disks or configurations have been expanded or upgraded.
 - If selected, only configuration-updated nodes are restarted. For example, if you expanded a core group node, but did not expand a master group node's disk or upgrade its configurations, only the ECS instance in the core group is restarted, and the ECS instance in the master group will not be restarted.
 - If not selected, all nodes (all instances in the cluster) are restarted.

9. During the restart process, the prompt Restarting Servers in the following figure is displayed for the corresponding node group (such as a core group).

Host

(MASTER)

Subscription

Public Bandwidth: 8M

Hosts: 1

CPU: 4 Cores

Memory: 8GB

Data Disk Type: SSD

Disk 80GB*1 Disks

Core Instance Group

Restarting Servers

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

10. Log on to the EMR cluster to check upgrades. After the prompt in step 9 is not displayed, all upgraded configurations take effect.
11. If you just upgraded the CPU and did not upgrade the memory, ignore this step. If you upgraded the memory, you need to modify cluster configurations so that YARN can use new resources.
- On the Clusters and Services pane, click YARN.
 - Click the Configuration tab, locate `yarn . nodemanage r . resource . memory - mb` and `yarn . scheduler . maximum - allocation - mb` two configuration items, and change the value of the two configuration items to `machine memory * 0.8`. The unit is MB. For example, in this new configurations, the memory is 32 GB, configure the following value 26214 ($32 \times 1024 \times 0.8$).
- ```
yarn . nodemanage r . resource . memory - mb = 26214
yarn . scheduler . maximum - allocation - mb = 26214
```
- In the upper right corner, click Save.
  - In the upper right corner, click Actions > CONFIGURE All Components.
  - In the upper right corner, click View Operation Logs and wait for the status of the CONFIGURE YARN operation type to Successful.
  - In the upper right corner, click Actions > RESTART All Components.
  - Click View Operation Logs, wait for the status of the RESTART YARN operation type to Successful, and new resources can be used by YARN.

## 3.2 Disk expansion

In the event that you do not have enough data storage space, cluster disks of E-MapReduce 3.11.0 or later can be expanded on the cluster overview page.

Enter the expansion interface

On the Cluster Management page, find the target cluster. Click View Details, and in the upper-right corner, click Disk Capacity Expansion.

Select the instance group that needs to be expanded.



### Notice:

Make sure your account balance is sufficient. The fee of expanding a data disk is deducted automatically. When your account balance is insufficient the, the process of expanding a data disk will be interrupted.

Expansion interface

The following figure shows the disk expansion interface displayed.

- Configuration: Data disk configurations of the current instance group.
- Expand to: Data disk size after the disk is expanded.



### Note:

- E-MapReduce currently supports the expansion disks, but not their reduction.
- The expansion of disks is only supported for data disks.

Restart a cluster

1. After completing the disk expansion, the corresponding core instance group displays The disk capacity expansion is complete. Restart the servers for the changes to take effect, as shown in the following figure.

|                               |               |                                                                                                                   |                             |           |             |
|-------------------------------|---------------|-------------------------------------------------------------------------------------------------------------------|-----------------------------|-----------|-------------|
| Host                          |               | Core Instance Group  The disk capacity expansion is complete. Restart the servers for the changes to take effect. |                             |           |             |
| Master Instance Group(MASTER) | Pay-As-You-Go | ECS ID                                                                                                            | Component Deployment Status | Public IP | Intranet IP |
|                               |               | Created At                                                                                                        |                             |           |             |

2. Click The disk capacity expansion is complete. Restart the servers for the changes to take effect, as shown in the following figure.

Restart Cluster ✕

---

Cluster: C-976F407DC1F1542A / TESTABBY

The server will be restarted. All services will be suspended during the restart. Make sure you restart your server during non-business hours to avoid service interruption.

☒ Rolling Start ☒ Only Restart Updated Nodes



#### Note:

If you restart a cluster, you also restart ECS instances. ECS instance data is unavailable during restarting.

- Rolling Start
    - If selected, only ECS instances you select is restarted. Instances are restarted one after another. Each restart takes about five minutes.
    - If not selected, all ECS instances are restarted.
  - Only Restart Updated Nodes
    - Configuration-updated nodes are nodes whose disks or configurations have been expanded or upgraded.
    - If selected, only configuration-updated nodes are restarted. For example, if you expanded a core group node, but did not expand a master group node's disk or upgrade its configurations, only the ECS instance in the core group is restarted, and the ECS instance in the master group will not be restarted.
    - If not selected, all nodes (all instances in the cluster) are restarted.
3. During the restart process, the prompt in the following figure is displayed for the corresponding node group (such as a core group).

|                                             |                                         |                             |           |             |            |
|---------------------------------------------|-----------------------------------------|-----------------------------|-----------|-------------|------------|
| Host <span>⌵</span>                         | Core Instance Group  Restarting Servers |                             |           |             |            |
| Master Instance Group(MASTER) Pay-As-You-Go | ECS ID                                  | Component Deployment Status | Public IP | Intranet IP | Created At |

4. After the prompt in step 3 is no longer displayed, all data disks are expanded and available to use.

## Node data disk expansion

1. Log on to the E-MapReduce console, find the target cluster, and click View Details to go to the cluster overview page.
2. View the ECS ID of the core node in the cluster you want to expand, such as `i-bp1bsithym-5hh9h93xxx`. All the node disks in the cluster are expanded uniformly by default to make sure the disk spaces of all nodes in the cluster are consistent.
3. Copy the ECS ID and go to the ECS console. Select Instance Details in the pane on the left, then select the instance ID and enter the ECS ID in the search box. Note that the same region must be selected.
4. When the corresponding ECS node is found, click Manage, and click Disks in the navigation pane on the left.
5. Expand the data disks. You must repeat the following steps to scale out each disk, because you cannot expand multiple disks in batches.
6. Expand all the disks in the ECS console and then restart the nodes.

See [Linux \\_ Resize a data disk](#) for more information on expanding data disks.



### Note:

If the unmount operation fails, YARN and HDFS services must be disabled in the cluster. Disk1 operations may encounter unmount failures because of log writing by `ilogtail`. `ilogtail` must be temporarily killed using the command `sudo pgrep ilogtail | sudo xargs kill`. `ilogtail` can then be recovered by restarting the node.

7. You can then see that all disks are expanded using the `df -h` command on the node.

## Node system disk expansion

Cluster disks of E-MapReduce 3.11.0 or earlier can be scaled out in the ECS console.

1. In the E-MapReduce console, enter the details page of the cluster you want to expand.
2. View the ECS ID of the master node in the cluster you want to expand, such as `i-bp1bsithym-5hh9h93xxx`. All the node disks in the cluster are expanded uniformly by default to make sure the disk spaces of all nodes in the cluster are consistent.

3. Copy the ECS ID and go to the ECS console. Select Instance Details in the pane on the left, then select the instance ID and enter the ECS ID in the search box. Note that the same region must be selected.
4. When the corresponding ECS node is found, click Manage, and click Disks in the navigation pane on the left.
5. Find the system disk (there is only one).
6. See [Increase system disk size](#) for more information on expanding system disks.

**Notice:**

- For non-HA clusters, while you are expanding the system disk, the cluster becomes unavailable.
- ECS disk operations may have changed the `/etc/hosts` file in the node. After the disk has been expanded, the file must be repaired. SSH log-on-free is also damaged, but the service remains unaffected and can be repaired manually.

### 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

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

Add Host Group

Host Group Name:

Task Instance Type:  ecs.sn2.large

System Disk Type: ☒ SSD Disk ☐ Ultra Disk

System Disk Size:  GB \* Disks IOPS 5400

Data Disk Type: ☒ SSD Disk ☐ Ultra Disk

Data Disk Size:  GB \* 4 Disks IOPS 4200

Task Instances:  Instances

Billing Method:

**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

Software such as Hadoop, Hive, and Pig contain a large number of configurations that can be modified using the software configuration function. For example, you may want to increase the number of service threads in the `dfs.namenode.handler.count` HDFS server from 10 to 50, or decrease the default size of HDFS file block `dfs.blocksize` from 128 MB to 64 MB because the system only contains small files. The software configuration function can only be performed once when a cluster is started.

### 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. All of the software and the corresponding versions can be seen in the Software Configuration step during cluster creation. If you want to change the configuration of the cluster, select a corresponding .json configuration file in the Enable Custom Setting box. You can then override or add to the default cluster parameters. An example of a .json file is 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 suffix of the actual incoming FileName parameter needs to be removed.
- Capitalize all letters of the ServiceName parameter name.

ConfigKey is the name of a configuration, whereas ConfigValue is a specific value.

Configuration files of all services are 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 setting configuration files, click Next to complete software configurations.**

## 5 Workflow development

---

### 5.1 Manage a workflow project

After creating an E-MapReduce cluster, you can create workflow projects so that multiple jobs can be run simultaneously or sequentially.

#### Create a project

1. At the top of the page, click the Data Platform tab to enter the Projects page.

Under the master account, you can view all of its projects and RAM user accounts. RAM users can only view projects if they have development permissions. The granting of project development permissions must be configured in the master account. For more information about authorization, see [User management](#) below.

2. In the upper-right corner, click New Project. The New Project dialog box is displayed.
3. Enter the project name and description and click Create.



#### Note:

You can only create a project with the master account. New Projects are only visible to the administrator of the master account.

#### User management

After creating a new project, you can grant operational permissions for the project to RAM user accounts.

1. In the Project List page, click View Details in the Actions column.
2. Click the User Management tab.
3. Click Add User to add RAM users to the project under the master account.

The added RAM users become members of the project and are able to view and develop the jobs and workflows under the project. If you remove a RAM user from a project, click Delete in the Actions column.



#### Note:

You can only add project members with the master account. The User Management tab is only visible to the administrator of the master account.



## Associate clusters

After creating a new project, you need to associate it with a cluster so that the workflow in the project can run on it.

1. In the Projects page, click View Details in the Actions column.
2. Click the Cluster Settings tab.
3. Click Add Cluster. From the drop-down menu, you can select a Subscription or Pay-As-You-Go cluster. (Clusters created by temporary jobs are not listed here.)
4. Click OK.

To disassociate the cluster, click Delete in the Operation column.



### Note:

You can only associate cluster with the master account. The Cluster Settings tab is only visible to the administrator of the master account.

To set both the queue and user to submit jobs to the cluster, click Modify Configuration in the Operation column. The configuration items are as follows:

- **Default Submit Job User:** Sets the default Hadoop user who submits the job to the selected cluster in the project. The default value is hadoop. There can only be one default user.
- **Default Submit Job Queue:** Sets the default queue that the jobs are submitted to in the project. If you leave this blank, the job will be submitted to the default queue.
- **Submit Job User Whitelist:** Sets Hadoop users who can submit jobs to the selected cluster in the project. If there is more than one user, they can be separated by a comma (,).
- **Submit Job Queue Whitelist:** Sets the queue of the selected cluster that jobs in the project can run in. If there is more than one queue, they can be separated by a comma (,).
- **Client whitelist:** Configures the client that can submit jobs. You can select either the E-MapReduce master node or the E-MapReduce gateway. Gateways that you have built are not currently listed here.

## 5.2 Job operations

In a project, you can create jobs such as Shell, Hive, Spark, SparkSQL, MapReduce, Sqoop, Pig and Spark Streaming jobs.

### Create a job

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. Click the Data Platform tab to enter the Projects page.
3. Click Design Workflow next to the target project in the Actions column.
4. On the left side of the job editing page, right-click the folder you want to operate on and select New Job.
5. In the New Job dialog box, enter a name and description for the job and select a job type.

Once the job type is selected, it cannot be modified.

6. Click OK.



Note:

You can also create subfolders, rename folders, and delete folders by right-clicking on them.

### Develop a job

For more information about the various types of jobs, see [Jobs](#).



Note:

When you insert an OSS path, if you select the OSSREF file prefix, the OSS file is downloaded to the cluster and added to the classpath.

- Basic settings

Click the Job Settings in the upper-right corner of the page to enter the Job Running Configuration page.

- **Number of Retries:** Sets the number of retries for if a job fails during a workflow. This option does not take effect if you run the job directly on the Job Editing page.
- **Failure Policy:** Sets whether to run the next job or suspend the current workflow in the event that a job fails during a workflow.
- **Resource File:** If you add resources such as JAR packages or UDFs that the current job depends on, you need to upload the resources to OSS first. After doing so, you can reference the resources directly in the job code.
- **Parameter Configuration:** Specifies the value of the variable referenced in the job code. You can reference variables in the code with the format `${variable name}`. Click the plus icon (+) on the right to add the key and value. The key is the variable name, whereas the value is the value of the variable. You can also customize the time variable according to the schedule time. The rules are as follows:

- yyyy indicates the year (4-digit format).
- MM indicates the month.
- dd indicates the day.
- HH24 indicates the hour. If the 12-hour clock is used, this is displayed as hh.
- mm indicates the minute.
- ss indicates the second.

The time variable can be any combination of time containing yyyy. You can also use the plus symbol (+) to advance time and the minus symbol (-) to delay time. For example, if `${yyyy-MM-dd}` indicates the current date, then:

- One year from now is: `${yyyy+1y}` or `${yyyy-MM-dd hh:mm:ss+1y}`.
- Three months from now is: `${yyyyMM+3m}` or `${hh:mm:ss yyyy-MM-dd+3m}`.
- Five days ago is: `${yyyyMMdd-5d}` or `${hh:mm:ss yyyy-MM-dd-5d}`.

- Advanced settings

To configure advanced settings, click the Advanced tab on the Job Settings page.

- **Mode:** Job running modes, including YARN and LOCAL. In YARN mode, the job is submitted on YARN by the Launcher. In LOCAL mode, jobs run directly on the assigned host.
- **Environment Variables:** Add environment variables to run jobs, or export environment variables directly in the job script.
- **Scheduling Parameters:** Set job configurations, such as the YARN queue, CPU, memory, and Hadoop users. If you do not set this parameter, the job adopts the default value of the Hadoop cluster.

### Run a job

Once a job has been developed and configured, you can click Run in the top right corner to run the job.

### View a log

After the job has been set to run, you can view its running log in the View Records tab at the bottom. Click Workflow to enter the detailed log page. Here, you can see information such as the job's submitting log and the YARN Container log.

## 5.3 Ad hoc queries

Only three types of ad hoc query are supported: HiveSQL, SparkSQL, and Shell. When you execute an ad hoc query statement, the log and query results are displayed at the bottom of the log and query page.

### Create a job

When you execute a job on the Edit Jobs page and click Details, you are directed to the Details page that shows the operation logs and run logs of this job. Ad hoc queries and jobs are used in different places. Ad hoc queries are usually used by data analysts. You also need to use SQL as a tool to implement an ad hoc query.

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. Click the Data Platform tab to enter the Projects page.
3. Click Design Workflow next to the target project to enter the Edit Jobs page.
4. In the navigation pane on the left, click the Query tab to enter the Query page.

5. In the navigation pane on the left, right-click a folder as required and select New Job.
6. In the New Job dialog box, enter the job name and description, and select a job type.

The job type cannot be modified once the job has been created.

7. Click OK.



**Note:**

By right-clicking a folder, you can rename it, delete it, and create new subfolders.

### Develop a job

For more information, see [HiveSQL](#), [SparkSQL](#), and [Shell](#) job types.



**Note:**

When you insert an OSS UNI and select OSSREF as the prefix, E-MapReduce downloads OSS files to your cluster and adds them to the classpath.

#### · Basic job settings

In the top-right corner, click Configure Jobs. The Job Settings dialog box is displayed.

- **Resource File:** If you want to add resources such as JAR packages or UDFs that a job execution depends on, you must upload these files to OSS. When you select a resource, you can use this resource in a job.
- **Parameter Configuration:** Specifies the values of variables used in a job. You can use variables in your code. The format is `${variable name}`. Click the plus icon (+) on the right to add key-value pairs. The key is the name of the variable,

whereas value is the value of the variable. You can also customize the time variable according to the schedule time. The rules are as follows:

- yyyy indicates the year (4-digit format).
- MM indicates the month.
- dd indicates the day.
- HH 24 indicates the hour. If the 12-hour clock is used, this is displayed as hh.
- mm indicates the minute.
- ss indicates the second.
- The time variable can be any combination of time containing yyyy. You can also use the plus symbol (+) to advance time and the minus symbol (-) to delay time. For example, if `${yyyy-MM-dd}` indicates the current date, then:
  - One year from now is: `${yyyy+1y}` or `${yyyy-MM-dd hh:mm:ss+1y}`.
  - Three months from now is: `${yyyyMM+3m}` or `${hh:mm:ss yyyy-MM-dd+3m}`.
  - Five days ago is: `${yyyyMMdd-5d}` or `${hh:mm:ss yyyy-MM-dd-5d}`.

#### · Advanced job settings

To configure advanced settings, click the Advanced tab on the Job Settings page.

- **Mode:** Job running modes, including YARN and LOCAL. In YARN mode, the job is submitted on YARN by the Launcher. In LOCAL mode, jobs run directly on the assigned host.
- **Scheduling Parameters:** Set job configurations, such as the YARN queue, CPU, memory, and Hadoop users. If you do not set this parameter, the job adopts the default value of the Hadoop cluster.

-

#### Execute a job

Once a job has been developed and configured, you can click Run in the top right corner to run the job.

#### View logs

After you execute a job, you can view its running logs on the Log tab at the bottom of the query page.

## 5.4 Manage a workflow

E-MapReduce workflows support the parallel execution of big data jobs based on DAG. You can also suspend, stop, rerun workflows, and view their running statuses in the webUI.

### Create a workflow

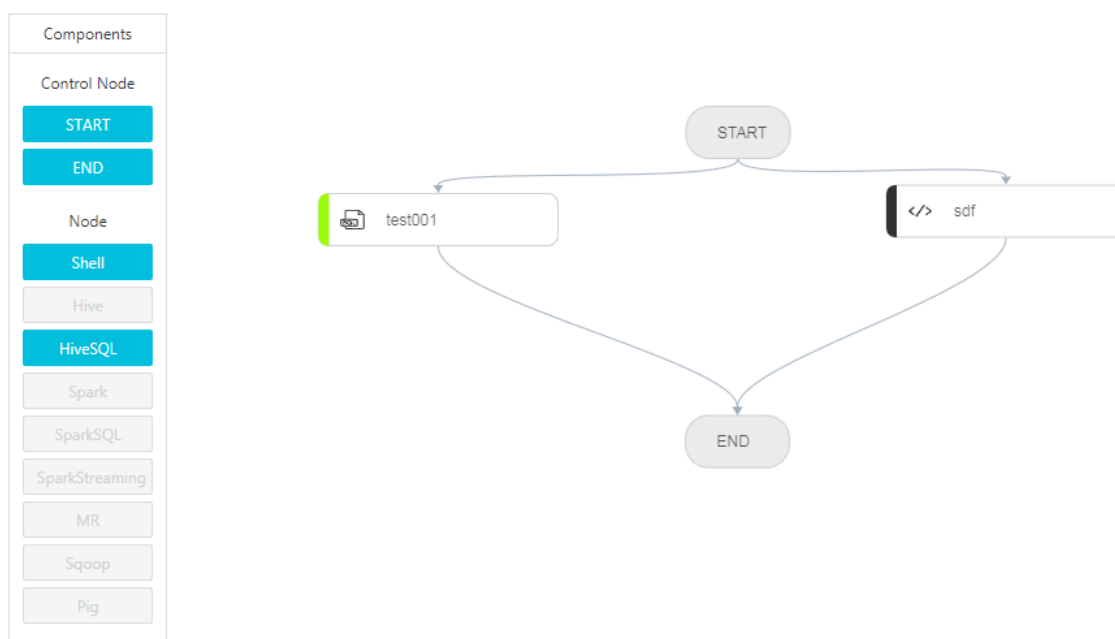
1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. At the top of the page, click the Data Platform tab.
3. Click Design Workflow next to the target project in the Actions column. Then select the Design Workflow tab.
4. On the left side, right-click the folder you want to operate on and select New Workflow.
5. In the New Workflow dialog box, enter the workflow name and description, and select the E-MapReduce cluster where you want to run the workflow.

You can select a Subscription or Pay-As-You-Go E-MapReduce cluster that has been created and associated with the project. Alternatively, you can create a new temporary cluster using the cluster template.

6. Click OK.

### Edit workflow

You can drag different types of jobs to the workflow editing canvas and specify the order of job instances by curve. After the job has been dragged, drag the END component from the control node area to the canvas. This indicates that the entire workflow is complete.



## Configure a workflow

On the right of the Workflow Design page, click **Configure** to configure the workflow scheduling.

- **Run In:** The E-MapReduce cluster where the workflow is to run can be modified.
- **Scheduling Policy:** After workflow scheduling has been enabled, period schedule are mandatory by default, and dependency schedule can be added.
  - **Time Scheduler:** Sets the start and end times for the workflow scheduling. The system then runs the workflow according to the schedule you set.
  - **Dependency:** Select the dependency workflow of the current workflow from the selected project. After the dependency workflow has been completed, the current workflow is scheduled to run. Currently, only one workflow can be selected.

## Run a workflow

Once a workflow has been developed and configured, you can click **Run** in the top right corner to run the workflow.

## View and operate workflow instances

After the workflow is running, click the **View Records** tab on the left to view the running status of the workflow instance. Click **View Details** next to the workflow



instance to view the running status of the job instance. You can also suspend, resume, stop, and rerun workflow instances.

Home Page > Data Platform > Project (FP-19DF27ECF61E5AF1) > View Records > Workflow Instance (FI-EC6641ECF29D9E2F) > View Details

| Cluster Information             |                               | Chart |  |
|---------------------------------|-------------------------------|-------|--|
| ID: FI-EC6641ECF29D9E2F         | Name: test                    |       |  |
| Workflow ID: F-798F6400F5485142 | Run In: C-555D28C90826D188    |       |  |
| Status: <b>FAILED</b>           | Run Time: 11 Seconds          |       |  |
| Start Time: 2018-11-12 17:08:42 | End Time: 2018-11-12 17:08:53 |       |  |

Name  Enter

| Job Instance ID ↕    | Name ↕  | Run In             | Type ↕   | Host                       | Start Time ↕        | Completed At ↕      | Run Time   | Execution State ▾ | Actions                      |
|----------------------|---------|--------------------|----------|----------------------------|---------------------|---------------------|------------|-------------------|------------------------------|
| FIN-7CDAF5769E68D2FB | test001 | C-555D28C90826D188 | HIVE_SQL | emr-header-1.cluster-84010 | 2018-11-12 17:08:43 | 2018-11-12 17:08:53 | 10 Seconds | <b>FAILED</b>     | <a href="#">View Details</a> |
| FIN-05CD626FF3B17CC9 | odf     | C-555D28C90826D188 | SHELL    | emr-header-1.cluster-84010 | 2018-11-12 17:08:44 | 2018-11-12 17:08:53 | 9 Seconds  | OK                | <a href="#">View Details</a> |

- **Suspend workflow instance:** The job instance continues to run, but subsequent instances do not. By clicking **Resume Workflow**, the system continues to run the subsequent jobs.
- **Stop workflow instance:** All running job instances stop immediately.
- **Rerun workflow instance:** The system runs the workflow from the start component
- 

## 5.5 Jobs

### 5.5.1 Configure a Hadoop MapReduce job

#### Procedure

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. At the top of the navigation bar, click **Data Platform**.
3. In the **Actions** column, click **Design Workflow** next to the specified project.
4. On the left of the **Job Editing** page, right-click the folder you want to operate and select **New Job**.
5. In the **New Job** dialog box, enter the job name and description.
6. Select a Hadoop job type to create a Hadoop MapReduce job. This type of job is submitted in the background using the following process.

```
hadoop jar xxx.jar [MainClass] - Dxxx
```

7. Click **OK**.



**Note:**

You can also create subfolders, rename folders, and delete folders by right-clicking on them.

8. Enter the parameters in the Content field that are required to submit this job. Enter the parameters after the Hadoop jar, followed by other command line parameters.

For instance, if you want to submit a Hadoop sleep job that doesnot read or write any data, this will only succeed if you submitMapper/Reducer tasks to the cluster and wait for each task to sleep for a while. In Hadoop, this job is packaged in the Hadoop release version's hadoop-mapreduce-client-jobclient-2.6.0-tests.jar. If this job is submitted from the command line, the command should read as follows.

```
hadoop jar / path / to / hadoop - mapreduce - client - jobclient
- 2 . 6 . 0 - tests . jar sleep - m 3 - r 3 - mt 100 -
rt 100
```

To configure this job in E-MapReduce, enter the following content in the Content field.

```
/ path / to / hadoop - mapreduce - client - jobclient - 2 . 6 . 0 -
tests . jar sleep - m 3 - r 3 - mt 100 - rt 100
```



**Note:**

The jar package path used here is an absolute path on the E-MapReduce host. However, the user may put these jar packages anywhere, and as clusters are created and released, the packages become unavailable. Therefore, upload the jar package by performing the following steps:

- a. Users send their own jar packages to the bucket of OSS for storage. When you configure the parameters for Hadoop, click Select OSS path to select and execute the jar package you want from the OSS directory. The system will then auto-complete the OSS address for jar packages. Be sure to switch the prefix of the jar to ossref by clicking Switch resource type. This ensures that the jar package is downloaded correctly by MapReduce.
- b. Click OK. The OSS path for this package will be auto-completed in the Content field. When a job is submitted, the system will find the corresponding jar packages automatically based on this path.
- c. Behind the jar package path for this OSS, other command line parameters for running jobs will be filled in further.

9. Click Save.

In the example above, the sleep job has no data input/output. If you want the job to read data and process input results, such as wordcounts, the data input and output paths need to be specified. You can read/write data on the HDFS of the E-MapReduce cluster as well as on OSS. To read/write data on OSS, write the data path as the OSS path when specifying the input and output paths. For instance:

```
jar ossref :// emr / checklist / jars / chengtao / hadoop / hadoop
- mapreduce - examples - 2 . 6 . 0 . jar randomtext writer - D
mapreduce . randomtext writer . totalbytes = 320000 oss :// emr /
checklist / data / chengtao / hadoop / Wordcount / Input
```

## 5.5.2 Configure a Hive job

When you apply for clusters in E-MapReduce, you are provided with a Hive environment by default. Using Hive, you can create and operate tables and data.

### Procedure

1. Prepare the Hive script in advance. For example:

```
USE DEFAULT ;
DROP TABLE uservisits ;
CREATE EXTERNAL TABLE IF NOT EXISTS uservisits
(sourceIP STRING , destURL STRING , visitDate STRING ,
adRevenue DOUBLE , user
Agent STRING , countryCode STRING , languageCode STRING
, searchWord STRING , duration INT) ROW FORMAT DELIMITED
FIELDS TERMINATED BY ',' STORED AS SEQUENCEFILE LOCATION '/
HiBench / Aggregation / Input / uservisits ' ;
DROP TABLE uservisits _aggre ;
CREATE EXTERNAL TABLE IF NOT EXISTS uservisits _aggre
(sourceIP STRING , sumAdRevenue DOUBLE) STORED AS
SEQUENCEFILE LOCATION '/ HiBench / Aggregation / Output / uservisits _aggre
' ;
INSERT OVERWRITE TABLE uservisits _aggre SELECT
sourceIP , SUM (adRevenue) FROM uservisits GROUP BY
sourceIP ;
```

2. Save this script into a script file, such as `uservisits _aggre_hdfs .`

`hive` , and upload it to an OSS directory (for example, `oss :// path / to / uservisits _aggre_hdfs . hive` ).

3. Log on to the [Alibaba Cloud E-MapReduce console](#).
4. At the top of the navigation bar, click Data Platform.
5. In the Actions column, click Design Workflow next to the specified project.
6. On the left of the Job Editing page, right-click the folder you want to operate and select New Job.
7. In the New Job dialog box, enter the job name and description.

8. Select the Hive job type to create a Hive job. This type of job is submitted in the background using the following method.

```
hive [user provided parameters]
```

9. Click OK.



**Note:**

You can also create subfolders, rename folders, and delete folders by right-clicking on them.

10. Enter the parameters in the Content field after the Hive commands. For example, if you want to use a Hive script uploaded to OSS, enter the following.

```
- f ossref :// path / to / uservisits _aggre_hdf s . hive
```

You can also click Select OSS path to view and select from OSS. The system will automatically complete the path of the Hive script on OSS. Switch the Hive script prefix to ossref by clicking Switch resource type. This ensures that the file is correctly downloaded by E-MapReduce.

11. Click Save to complete the Hive job configuration.

### 5.5.3 Configure a Pig job

When you apply for clusters in E-MapReduce, a Pig environment is provided by default. Using Pig, you can create and operate tables and data.

#### Procedure

1. Prepare the Pig script in advance. For example:

```
``` shell
/*
 * Licensed to the Apache Software Foundation (ASF)
 * under one
 * or more contributor license agreements. See the
 * NOTICE file
 * distributed with this work for additional
 * information
 * regarding copyright ownership. The ASF licenses
 * this file
 * to you under the Apache License, Version 2.0 (
 * the
 * "License"); you may not use this file except in
 * compliance
 * with the License. You may obtain a copy of
 * the License at
 *
 * http://www.apache.org/licenses/LICENSE-2.0
 *
 * ```
```

```

* Unless required by applicable law or agreed to
in writing, software
* distributed under the License is distributed on
an "AS IS" BASIS,
* WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND,
either express or implied.
* See the License for the specific language
governing permissions and
* limitations under the License.
*/
-- Query Phrase Popularity (Hadoop cluster)
-- This script processes a search query log file
from the Excite search engine and finds search
phrases that occur with particular high frequency
during certain times of the day.
-- Register the tutorial JAR file so that the
included UDFs can be called in the script.
REGISTER oss :// emr / checklist / jars / chengtao / pig /
tutorial . jar ;
-- Use the PigStorage function to load the excite
log file into the "raw" bag as an array of
records .
-- Input : ( user , time , query )
raw = LOAD ' oss :// emr / checklist / data / chengtao / pig /
excite . log . bz2 ' USING PigStorage ('\ t ') AS ( user ,
time , query );
-- Call the NonURLDetector UDF to remove records
if the query field is empty or a URL .
clean1 = FILTER raw BY org . apache . pig . tutorial .
NonURLDetector ( query );
-- Call the ToLower UDF to change the query field
to lowercase .
clean2 = FOREACH clean1 GENERATE user , time , org .
apache . pig . tutorial . ToLower ( query ) as query ;
-- Because the log file only contains queries for
a single day , we are only interested in the hour
.
-- The excite query log timestamp format is
YYMMDDHHMM SS .
-- Call the ExtractHour UDF to extract the hour (
HH ) from the time field .
houred = FOREACH clean2 GENERATE user , org . apache .
pig . tutorial . ExtractHour ( time ) as hour , query ;
-- Call the NGramGenerator UDF to compose the n -
grams of the query .
ngramed1 = FOREACH houred GENERATE user , hour , flatten
( org . apache . pig . tutorial . NGramGenerator ( query )) as
ngram ;
-- Use the DISTINCT command to get the unique n -
grams for all records .
ngramed2 = DISTINCT ngramed1 ;
-- Use the GROUP command to group records by n -
gram and hour .
hour_frequency1 = GROUP ngramed2 BY ( ngram , hour );
-- Use the COUNT function to get the count (
occurrences ) of each n - gram .
hour_frequency2 = FOREACH hour_frequency1 GENERATE
flatten ( $ 0 ), COUNT ( $ 1 ) as count ;
-- Use the GROUP command to group records by n -
gram only .
-- Each group now corresponds to a distinct n -
gram and has the count for each hour .
uniq_frequency1 = GROUP hour_frequency2 BY group ::
ngram ;

```

```
-- For each group , identify the hour in which
this n - gram is used with a particularly high
frequency .
-- Call the ScoreGenerator UDF to calculate a "
popularity " score for the n - gram .
uniq_freqency2 = FOREACH uniq_freqency1 GENERATE
flatten ($ 0 ), flatten ( org . apache . pig . tutorial .
ScoreGenerator ($ 1 ));
-- Use the FOREACH - GENERATE command to assign names
to the fields .
uniq_freqency3 = FOREACH uniq_freqency2 GENERATE $
1 as hour , $ 0 as ngram , $ 2 as score , $ 3 as
count , $ 4 as mean ;
-- Use the FILTER command to move all records
with a score less than or equal to 2 . 0 .
filtered_uniq_freqency = FILTER uniq_freqency3 BY
score > 2 . 0 ;
-- Use the ORDER command to sort the remaining
records by hour and score .
ordered_uniq_freqency = ORDER filtered_uniq_freqency
BY hour , score ;
-- Use the PigStorage function to store the results
.
-- Output : ( hour , n - gram , score , count , average_co
unts_among_all_hours )
STORE ordered_uniq_freqency INTO ' oss :// emr /
checklist / data / chengtao / pig / script1 - hadoop - results '
USING PigStorage ();
```

```

2. Save this script into a script file, such as `script1 - hadoop - oss . pig` , and upload it to an OSS directory (for example, `oss :// path / to / script1 - hadoop - oss . pig` ).
3. Log on to the [Alibaba Cloud E-MapReduce console](#).
4. At the top of the navigation bar, click Data Platform.
5. In the Actions column, click Design Workflow next to the specified project.
6. On the left of the Job Editing page, right-click the folder you want to operate and select New Job.
7. In the New Job dialog box, enter the job name and description.
8. Select the Pig job type to create a Pig job. This type of job is submitted in the background using the following method.

```
pig [user provided parameters]
```

9. Click OK.



**Note:**

You can also create subfolders, rename folders, and delete folders by right-clicking on them.

10. Enter the parameters in the Content field after the Pig commands. For example, if you want to use a Pig script uploaded to OSS, enter the following.

```
- x mapreduce ossref :// emr / checklist / jars / chengtao / pig / script1 - hadoop - oss . pig
```

You can click Select OSS path to view and select from OSS. The system will automatically complete the path of Pig script on OSS. Switch the Pig script prefix to ossref by clicking Switch resource type. This ensures that the file is correctly downloaded by E-MapReduce.

11. Click Save to complete the Pig job configuration.

## 5.5.4 Configure a Spark job

In this tutorial, you will learn how to configure a Spark job.

### Procedure

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. At the top of the navigation bar, click Data Platform.
3. In the Actions column, click Design Workflow next to the specified project.
4. On the left of the Job Editing page, right-click the folder you want to operate and select New Job.
5. In the New Job dialog box, enter the job name and description.
6. Click OK.



#### Note:

You can also create subfolders, rename folders, and delete folders by right-clicking on them.

7. Select the Spark job type to create a Spark job. This type of job is submitted in the background using the following method.

```
spark - submit [options] -- class [MainClass] xxx . jar
args
```

8. Enter the parameters in the Content field that are required to submit this job. Only the parameters after `spark - submit` can be entered. The following example shows how to enter the parameters for creating a Spark job and a PySpark job.

- Create a Spark job

Create a Spark WordCount job:

- Job name: WordCount
- Type: Select Spark
- Parameters:

■ Enter the following command:

```
spark - submit -- master yarn - client -- driver -
memory 7G -- executor - memory 5G -- executor -
cores 1 -- num - executors 32 -- class com . aliyun
. emr . checklist . benchmark . SparkWordC ount emr
- checklist_ 2 . 10 - 0 . 1 . 0 . jar oss :// emr /
checklist / data / wc oss :// emr / checklist / data / wc
- counts 32
```

■ Enter the following in the E-MapReduce job Content field:

```
-- master yarn - client -- driver - memory 7G --
executor - memory 5G -- executor - cores 1 -- num -
executors 32 -- class com . aliyun . emr . checklist .
benchmark . SparkWordC ount ossref :// emr / checklist /
jars / emr - checklist_ 2 . 10 - 0 . 1 . 0 . jar oss ://
emr / checklist / data / wc oss :// emr / checklist / data
/ wc - counts 32
```



**Notice:**

Job jar packages are saved in OSS. In the example above, the way to reference the Jar package is `ossref :// emr / checklist / jars / emr - checklist_ 2 . 10 - 0 . 1 . 0 . jar`. Click Select OSS path to view and select one from OSS. The system will automatically complete the



absolute path of the Spark script on OSS. Switch the default OSS protocol to the ossref protocol.

- Create a PySpark job

In addition to Scala and Java job types, E-MapReduce also supports Python job types in Spark. Create a Spark K-means job for the Python script:

- Job name: Python-Kmeans
- Type: Spark
- Parameters:

```
-- master yarn - client -- driver - memory 7g -- num -
executors 10 -- executor - memory 5g -- executor - cores
1 -- jars ossref :// emr / checklist / jars / emr - core
- 0 . 1 . 0 . jar ossref :// emr / checklist / python /
wordcount . py oss :// emr / checklist / data / kddb 5
32
```

- References of Python script resources are supported, and the ossref protocol is used.
- For PySpark, the online Python installation kit is not supported.

9. Click Save to complete the Spark job configuration.

### 5.5.5 Configure a Spark SQL

In this tutorial, you will learn how to configure a Spark SQL job.



Note:

By default, the mode of Spark SQL used for submitting a job is YARN.

#### Procedure

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. At the top of the navigation bar, click Data Platform.
3. In the Actions column, click Design Workflow next to the specified project.
4. On the left of the Job Editing page, right-click the folder you want to operate and select New Job.
5. In the New Job dialog box, enter the job name and description.
6. Click OK.



Note:

You can also create subfolders, rename folders, and delete folders by right-clicking on them.

7. Select the Spark SQL job type to create a Spark SQL job. This type of job is submitted in the background using the following method.

```
spark - sql [options] [cli option]
```

8. Enter the parameters in the Content field after the Spark SQL commands.

- -e option

-e options can be written to the running SQL by inputting them into the Content field of the job. For example:

```
- e " show databases ;"
```

- -f option

-f options can be used to specify a Spark SQL script file. Uploading well-prepared Spark SQL script files to OSS can provide greater flexibility. We recommend that you use this operation mode. For example:

```
- f ossref :// your - bucket / your - spark - sql - script . sql
```

9. Click Save to complete Spark SQL job configuration.

## 5.5.6 Configure a Shell job

In this tutorial, you will learn how to configure a Shell job.



### Notice:

By default, Shell scripts are currently run by Hadoop. If you need to use the root user, the `sudo` command can be used. Use Shell script jobs with caution.

### Procedure

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. At the top of the navigation bar, click Data Platform.
3. In the Projects area, select a target project ID to go to the Project Management tab page.
4. In the left-side navigation bar, click Edit Jobs next to the specified project.
5. On the left of the Edit Jobs tab page, right-click the folder you want to operate and select New Job.

6. In the New Job dialog box, enter the job name and description.
7. Select the Shell job type to create a Bash Shell job.
8. Click OK.

**Note:**

You can also create subfolders, rename folders, and delete folders by right-clicking on them.

9. Enter the parameters in the Content field after the Shell commands.

- -c option

-c options can be used to set Shell scripts to run by inputting them into the Content field of the job. For example:

```
- c " echo 1 ; sleep 2 ; echo 2 ; sleep 4 ; echo 3
 ; sleep 8 ; echo 4 ; sleep 16 ; echo 5 ; sleep 32
 ; echo 6 ; sleep 64 ; echo 8 ; sleep 128 ; echo
 finished "
```

- -f option

-f options can be used to run Shell script files. By uploading a Shell script file to OSS, Shell scripts on OSS can be defined in the job parameters, making it more flexible than the -c option. For example:

```
- f ossref :// mxbucket / sample / sample - shell - job . sh
```

10. Click Save to complete Shell job configurations.

## 5.5.7 Configure a Sqoop job

In this tutorial, you will learn how to configure a Sqoop job.

**Note:**

Only E-MapReduce products with version V1.3.0 or later support the Sqoop job type. Running a Sqoop job on lower versions will fail and errlog will report "Not supported" errors. For more information on parameters, see [Sqoop](#).

### Procedure

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. At the top of the navigation bar, click Data Platform.
3. In the Actions column, click Design Workflow next to the specified project.

4. On the left of the Job Editing page, right-click the folder you want to operate and select New Job.
5. In the New Job dialog box, enter the job name and description.
6. Select the Sqoop job type to create a Sqoop job. This type of job is submitted in the background using the following method.

```
sqoop [args]
```

7. Click OK.



Note:

You can also create subfolders, rename folders, and delete folders by right-clicking on them.

8. Enter the parameters in the Content field after the Sqoop commands.
9. Click Save to complete Sqoop job configuration.

## 5.5.8 Job operations

You can create, clone, modify, and delete jobs.

### Job creation

A new job can be created at any time. Currently, a job can only be used in the region where it is created.

### Job cloning

Configurations that already exist for a job can be cloned. A cloned job can also only be used in the region where it is created.

### Job modification

Before you can modify a job that needs to be added to an execution plan, you must first ensure that the execution plan is not running and that its periodic scheduling is not in progress.

Before you can modify a job that needs to be added to several execution plans, you must first ensure that none of the execution plans are running and that none of their periodic scheduling is in progress. Modifying a job may result in changes to all of the execution plans that use this job.

If you need to debug, we recommend that you perform cloning instead. After you debug, the original jobs in the execution plan are replaced.

## Job deletion

As with modification, a job can only be deleted when the execution plan where the job is located is not running and its periodic scheduling is not in progress.

## 5.5.9 Time and date variables

When you are creating a job, variable wildcards are supported in the job parameters for both time and date.

### Variable wildcard format

The format of the variable wildcards supported by E-MapReduce is either `${ dateexpr - 1d }` or `${ dateexpr - 1h }`. For example, assuming the current date and time is `2016 / 04 / 27 12 : 08 : 01 :`

- If `${ yyyyMMdd HH : mm : ss - 1d }` is displayed, the parameter wildcard is replaced with `20160426 12 : 08 : 01` when executed, which is the current date minus one day, and time accurate to the second.
- If `${ yyyyMMdd - 1d }` is displayed, the parameter wildcard is replaced with `20160426` when executed, which is the current date minus one day.
- If `${ yyyyMMdd }` is displayed, the parameter wildcard is replaced with `20160427`, which is the current date.

`dateexpr` represents the standard format of expressing time. Time is therefore formatted according to this expression and is followed by the amount of time that you want to add or deduct, which can be written as `N`. For example, `${ yyyyMMdd - 5d }`, `${ yyyyMMdd + 5d }`, `${ yyyyMMdd + 5h }`, or `${ yyyyMMdd - 5h }`.



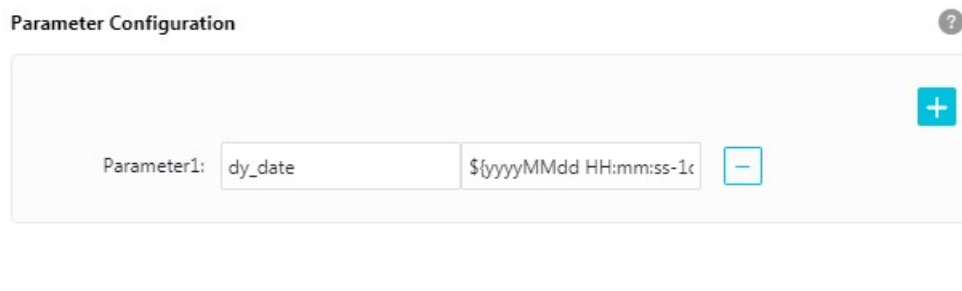
#### Note:

E-MapReduce currently supports the addition and deduction of hours and days only.

### Example

1. Click Job Settings on the top right of the Edit Jobs page.

2. Click the add icon to add new parameters on the Parameter Configuration part, and fill in the parameter according to the Variable wildcard format that mentioned above.



3. You can now use the reference of the parameter key in the job editing.

## 5.6 Old EMR Scheduling (Soon will be unavailable)

### 5.6.1 Notebooks

#### 5.6.1.1 Introduction

Notebooks allow you to compile and run Spark, Spark SQL, and Hive SQL tasks directly on the E-MapReduce console. You can then view the running results in the notebook. Notebooks are ideal for processing debugging tasks that require a shorter runtime and whose results need to be viewed directly. For tasks that have a longer runtime and require regular execution, the job and execution plan function must be used. This section describes how to create and run a notebook demo task.

#### Create a demo task

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. At the top of the navigation bar, click Old EMR Scheduling.
3. In the navigation bar on the left, click Notebook.

#### 4. Click New notebook demo.

notebook list

EMR-Hive-Demo

EMR-SparkSQL-Demo

EMR-Spark-Demo

EMR-SparkSQL-Demo

EMR-Spark-Demo

EMR-Hive-Demo

5. A confirmation box is displayed, indicating the required cluster environment. Click **OK** to create a demo task. Three examples of interactive tasks are created.

notebook list

EMR-Hive-Demo

EMR-SparkSQL-Demo

EMR-Spark-Demo

#### Run a Spark demo task

1. Click EMR-Spark-Demo to display the example of a Spark notebook. Before running the notebook, you need to associate the task to a created cluster. Select a created cluster in the list of available clusters. Note that the associated cluster must

be E-MapReduce 2.3 or later and have no less than three nodes, each with at least 4 cores and 8 GB of memory.



2. After a cluster is associated, click Run. When the associated cluster executes the Spark or Spark SQL notebook for the first time, it takes about one minute to




build the Spark context and running environment. It does not need to be built in subsequent executions. The running result is displayed under the Run button.

[Save Paragraph](#) [Hide results](#) [delete](#)

```
> %spark
import scala.math.random
import org.apache.spark._

val slices = 20
val n = math.min(100000L * slices, Int.MaxValue).toInt
val count = sc.parallelize(1 until n, slices).map { i =>
 val x = random * 2 - 1
 val y = random * 2 - 1
 if (x*x + y*y < 1) 1 else 0}.reduce(_ + _)
println("Pi is roughly " + 4.0 * count / n)
```

 run


Run results :

```
import scala.math.random
import org.apache.spark._
slices: Int = 20
n: Int = 2000000
count: Int = 1570075
Pi is roughly 3.14015
```

status : **FINISHED** , run 0second(s) , finish Time : Nov 9, 2018 11:25:52 AM

## Run a SparkSQL demo task

1. Click EMR-Spark-Demo to display the SparkSQL notebook example. Before running the notebook, you need to associate it to a created cluster. In the upper-right corner, select a created cluster from the list of available clusters.



2. The SparkSQL demo contains several demo sections that can be run individually or together by clicking Run All. After running, you can see the returned data results of each section.



Note:

If the section for creating a table is run multiple times, an error is reported indicating that the table already exists.

Save Paragraph
Hide results
delete

```

> %sql
-- get data
select * from uservisit_sparksql limit 10

```

run

Run results :

| ip            | uri                                                                      | birth     | score         | ua                      | country | state  | name                | level | dt         |
|---------------|--------------------------------------------------------------------------|-----------|---------------|-------------------------|---------|--------|---------------------|-------|------------|
| 170.131.2.2   | 13rdgckzlcblruc.html                                                     | 1984-8-7  | 336.869186722 | NuSearch Spider         | HUN     | HUN-NL | remnants            | 3     | 2016-01-01 |
| 162.114.4.2   | 6xpjrzjeytxdjsmwtmyeugkesratmpvamliekrijlgmvyyslqwgw.html                | 1978-1-9  | 331.791153595 | Superdownloads Spiderma | AUT     | AUT-ZR | MHD                 | 8     | 2016-01-01 |
| 177.110.45.18 | 11zvmoamsyaameokoeylbkivgquksibqbalnmailbiyfxitbhfdroyxesixbjndkyqz.html | 1986-9-25 | 411.968497603 | Mozilla/4.0             | FLK     | FLK-GB | apj@as.arizona.edu. | 7     | 2016-01-01 |
| 157.111.12.37 | 44mvdnl.html                                                             | 2002-7-3  | 486.660926201 | PHP/4.0.                | FIN     | FIN-CZ | diffuse             | 3     | 2016-01-01 |

status : **FINISHED** , run 0second(s) , finish Time : Nov 9, 2018 11:18:38 AM

### Run a Hive demo task

1. Click EMR-Hive-Demo to display the Hive notebook example. Before running the notebook, you need to associate it to a created cluster. In the upper-right corner, select a created cluster from the list of available clusters.
2. The Hive demo task contains several demo sections that can be run individually or together by clicking Run All. After running, you can see the returned data results of each section.



#### Note:

- When the associated cluster executes the Hive notebook for the first time, it takes a few seconds to build the Hive client running environment. It does not need to be built in subsequent executions.

- If the section for creating a table is run multiple times, an error is reported indicating that the table already exists.

Save Paragraph
Hide results
delete

```

> %hive
 -- get data
 select * from uservisit_hive limit 10

```

//

run

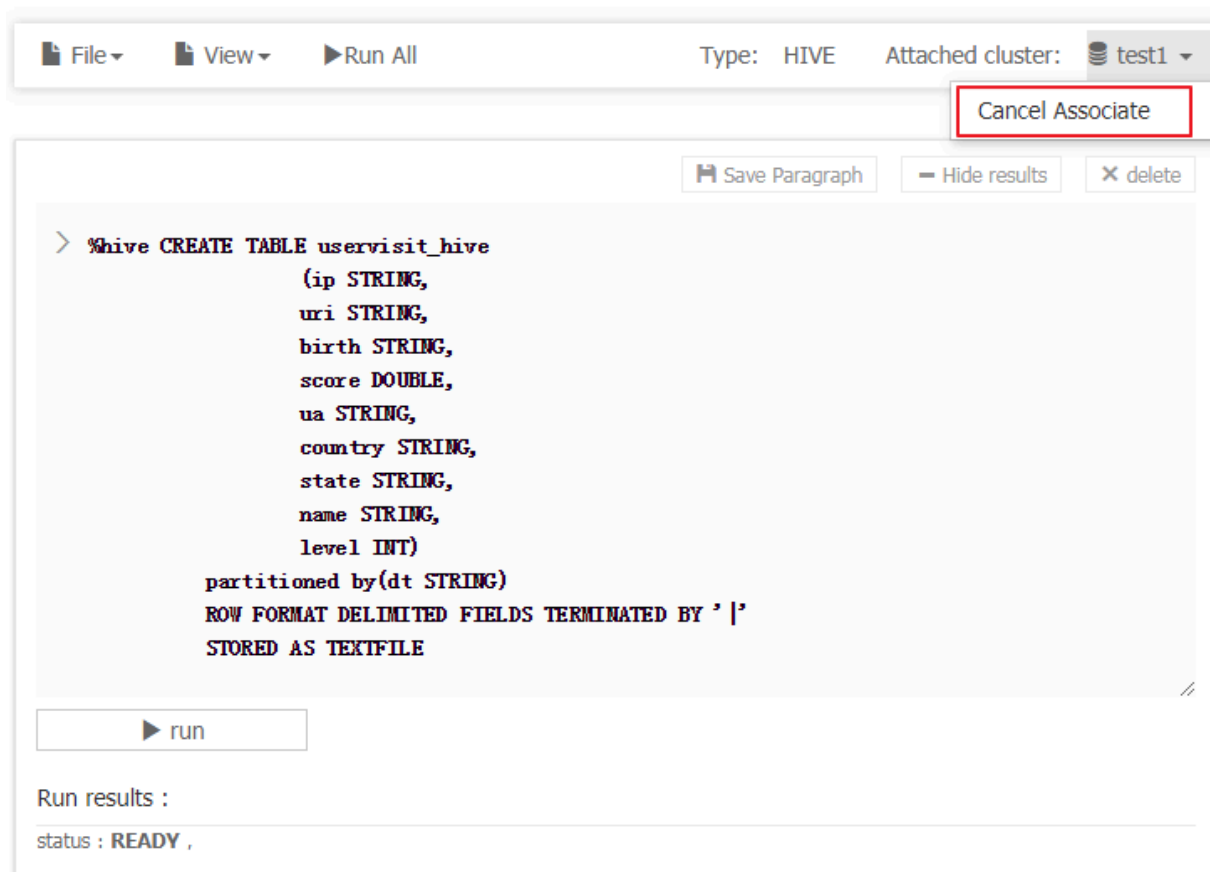
Run results :

| uservisit_hive.ip | uservisit_hive.uri                                                         | uservisit_hive.birth | uservisit_hive.score | uservisit_hive.ua        | uservisit_hive.country | uservisit_hive.state | uservisit_hive.name | uservisit_hive.level | uservisit_hive.dt |
|-------------------|----------------------------------------------------------------------------|----------------------|----------------------|--------------------------|------------------------|----------------------|---------------------|----------------------|-------------------|
| 170.131.22.2      | 13rdgckzlcblruc.html                                                       | 1984-8-7             | 336.869186722        | NuSearch Spider          | HUN                    | HUN-NL               | remnants            | 3                    | 2016-01-01        |
| 162.114.42        | 6xpjrzjeytxdjsmwtmyeugkesratmpvamliekrijlgmvyysrlqwgw.html                 | 1978-1-9             | 331.791153595        | Superdown loads Spiderma | AUT                    | AUT-ZR               | MHD                 | 8                    | 2016-01-01        |
| 177.110.45.18     | 11zvmoamsyaameokoeylbkivgquksibqbalnpmailbiyxitbhfdroyxesixbjndkyqzll.html | 1986-9-25            | 411.968497603        | Mozilla/4.0              | FLK                    | FLK-GB               | apj@as.arizona.edu. | 7                    | 2016-01-01        |

status : **FINISHED** , run 0second(s) , finish Time : Nov 9, 2018 11:21:08 AM

## Cancel the association with clusters

After a notebook is run in a cluster, the cluster creates a process for caching some context running environments in order to ensure a quick response upon re-execution. If you do not need to execute other notebooks, and you want to release the cluster resources occupied by caching, you can disassociate all interactive tasks that have been run from the associated clusters. In this way, you can release the memory resources occupied on the original associated clusters.



### 5.6.1.2 Operations

This section details how to perform a number of notebook operations, including how to create a new notebook task on the E-MapReduce console.

Create a new notebook task




Note:

The cluster on which an interactive task is run must be E-MapReduce 2.3 or later and have no less than three nodes, each with at least 4 cores and 8 GB of memory.

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. At the top of the navigation bar, click Old EMR Scheduling.
3. In the navigation bar on the left, click Notebook.

4. Click New notebook or File > -New notebook.



5. Enter a name and select the default type. Associating a cluster is optional. Click OK to create a notebook.

The screenshot shows the 'New Notebook' dialog box. At the top, the title is 'Notebook' with a close button (X) on the right. The dialog contains three main sections: 1. 'Name': A text input field containing 'test'. Below it, a note states: 'Length: 1 to 64 characters. Only Chinese characters, English letters, numbers '-', and '\_' are allowed'. 2. 'defaultType': Three radio buttons are shown: 'Spark' (selected), 'Spark SQL', and 'Hive'. Below them, a note says: 'when you run notebook without specify any type , notebook will use this as default type'. 3. 'Attached cluster': A dropdown menu that is currently empty. At the bottom right, there are two buttons: 'OK' (blue) and 'Cancel' (gray).

Three types of notebook task are supported. Spark can be used to write Scala code, Spark SQL can be used to write SQL statements supported by Spark, and Hive can be used to write SQL statements supported by Hive.

6. An associated cluster must be E-MapReduce 2.3 or later and have no less than three nodes, each with at least 4 cores and 8 GB of memory. You can also associate the cluster before running the task.

Up to 20 interactive tasks can be created in one account.

#### Enter and save a section


A paragraph is the smallest unit for running a notebook. Multiple paragraphs can be entered into a notebook. Each paragraph starts with either `% spark`, `% sql`, or `% hive`, indicating whether it is a Scala code paragraph, Spark SQL paragraph, or Hive SQL paragraph. The type prefix is separated by a blank space or by line feed and actual content. If the type prefix is not specified, the default type of the interactive task is used as the run type of this paragraph.

The following example shows how to create a temporary Spark table:

Paste the following code into the section and a red \* symbol is displayed, indicating that this notebook has been changed. Click Save Paragraph or run to save the modifications to the paragraph. Click + under the paragraph to create a new paragraph. Up to 30 paragraphs can be created in one notebook.

```
% spark
import org.apache.commons.io.IOUtils
import java.net.URL
import java.nio.charset.Charset
// load bank data
val bankText = sc.parallelize (
 IOUtils.toString (
 new URL ("http://emr-sample-projects.oss-cn-hangzhou.aliyuncs.com/bank.csv"),
 Charset.forName ("utf8")).split ("\n"))
case class Bank (age: Integer, job: String, marital: String, education: String, balance: Integer)
val bank = bankText.map (s => s.split(";")).filter (s => s(0) != "\" age \"").map (
 s => Bank (s(0).toInt,
 s(1).replaceAll ("\"", ""),
 s(2).replaceAll ("\"", ""),
 s(3).replaceAll ("\"", ""),
 s(5).replaceAll ("\"", ").toInt
)
).toDF ()
```

```
bank . registerTempTable (" bank ")
```



### Run a paragraph

Before running a notebook, you must first associate it to a created cluster. If a created notebook is not associated with a cluster, Not Attached is displayed in the upper-right corner of the page. Click it to select a cluster from the list of available clusters. Note that the associated cluster must be E-MapReduce 2.3 or later and have no less than three nodes, each with at least 4 cores and 8 GB of memory.



File View Run All

Type: SPARK Attached cluster: Not Attached

\* Save Paragraph Hide results delete

```
> %spark import org.apache.commons.io.IOUtils
import java.net.URL
import java.nio.charset.Charset
// load bank data
val bankText = sc.parallelize(
 IOUtils.toString(
 new URL("http://emr-sample-projects.oss-cn-hangzhou.aliyuncs.com/bank.csv"),
 Charset.forName("utf8")).split("\n")
)
case class Bank(age: Integer, job: String, marital: String, education: String, balance: Integer)
val bank = bankText.map(s => s.split(";")).filter(s => s(0) != "\age\").map(
s => Bank(s(0).toInt,
s(1).replaceAll("\\", ""),
s(2).replaceAll("\\", ""),
s(3).replaceAll("\\", ""),
s(5).replaceAll("\\", "").toInt)).toDF()
bank.registerTempTable("bank")
```

run

Run results :  
status : READY ,

Click Run to save the current paragraph and run the content. If this is the last paragraph, a new paragraph is created automatically.

PENDING indicates that the paragraph has not run yet, RUNNING indicates that the paragraph is running, FINISHED indicates that the running has finished, and ERROR indicates that an error has occurred. The running result is displayed beneath the Run button. During running, you can click Cancel beneath the Run button to cancel running. ABORT is displayed after running has been canceled.

Save Paragraph

Hide results

delete

```
> %spark
import org.apache.commons.io.IOUtils
import java.net.URL
import java.nio.charset.Charset // load bank data
val bankText = sc.parallelize(
 IOUtils.toString(
 new URL("http://emr-sample-projects.oss-cn-hangzhou.aliyuncs.com/bank.csv"),
 Charset.forName("utf8")).split("\n")
)
case class Bank(age: Integer, job: String, marital: String, education: String, balance: Integer)
val bank = bankText.map(s => s.split(";")).filter(s => s(0) != "\age\").map(
 s => Bank(s(0).toInt,
 s(1).replaceAll("\\", ""),
 s(2).replaceAll("\\", ""),
 s(3).replaceAll("\\", ""),
 s(5).replaceAll("\\", ").toInt
)
).toDF()
bank.registerTempTable("bank")
```

run

Run results :

```
import org.apache.commons.io.IOUtils
import java.net.URL
import java.nio.charset.Charset
bankText: org.apache.spark.rdd.RDD[String] = ParallelCollectionRDD[29] at parallelize at <console>:36
defined class Bank
bank: org.apache.spark.sql.DataFrame = [age: int, job: string, marital: string, education: string, balance: int]
```


status : **FINISHED** , run 0second(s) , finish Time : Nov 9, 2018 11:43:32 AM

The paragraph can be run multiple times, but only the result of the last running is retained. You cannot modify the content of a paragraph while it is running. It can only be modified after the running has finished.

### Run all

For a notebook, you can click Run All on the menu bar to run all paragraphs. The paragraphs are then submitted sequentially for running. Different types have independent execution queues. If a notebook contains multiple paragraph types, the order for executing them on the cluster is decided based on type after they have been submitted sequentially. Spark and Spark SQL support one-by-one execution. Hive supports concurrent execution, with the maximum number of concurrently executed interactive paragraphs on the same cluster is 10. Note that all concurrently executed

paragraphs are restricted by cluster resources. If the cluster size is small and many paragraphs need to be executed concurrently, the paragraphs still need to queue in YARN.



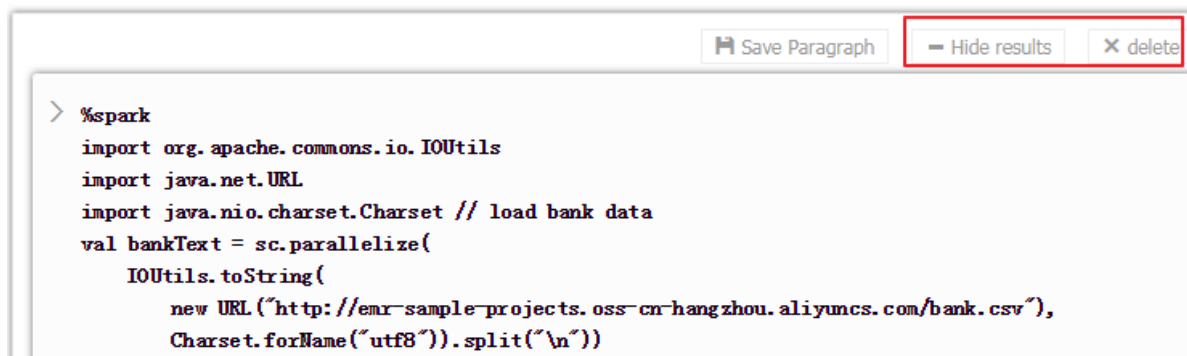
### Cancel the association with clusters

After a notebook is run in a cluster, the cluster creates a process for caching some context running environments to ensure a quick response upon re-execution. If you do not need to run other notebooks, and you want to release the cluster resources occupied by caching, you can disassociate all notebooks that have been run from the associated clusters. In this way, you can release the memory resources occupied on the original associated clusters.



## Other operations

- Paragraph operations



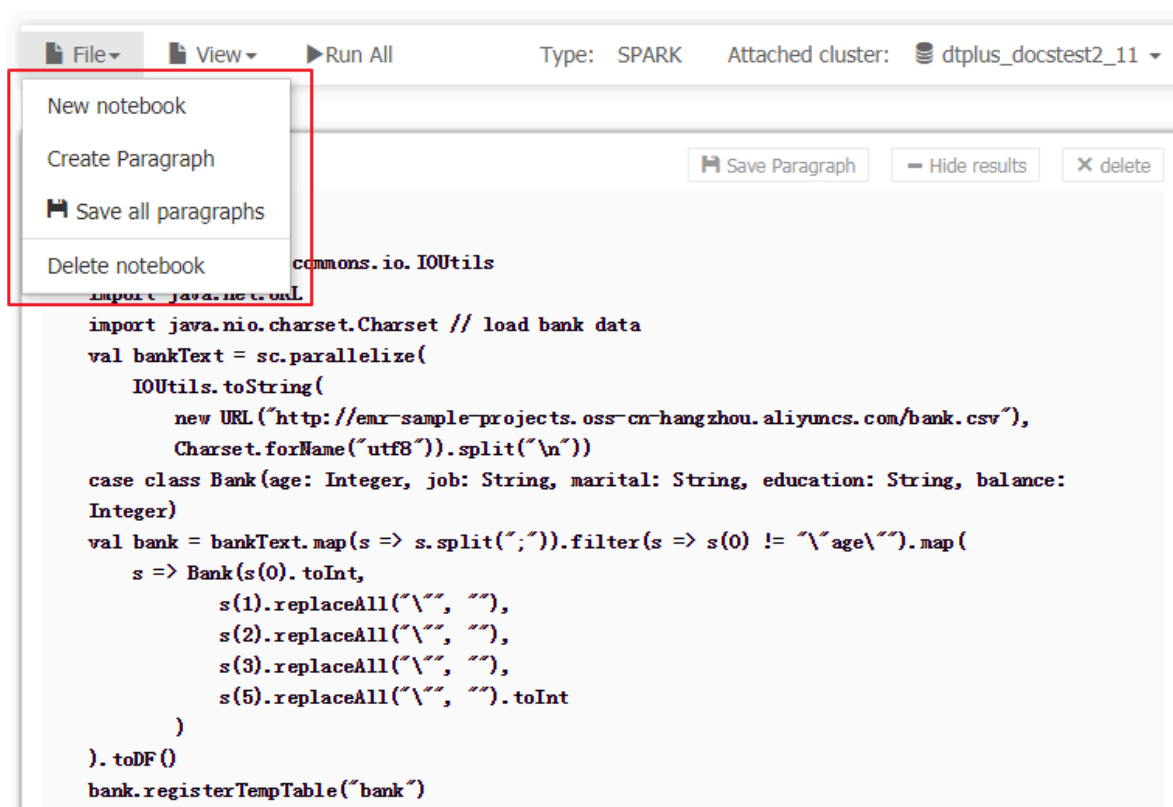
- Hide and display the results

Hide the paragraph results and only display the entered content of the paragraph.

- Delete a paragraph

Delete the current paragraph. Paragraphs that are running can also be deleted.

- File menu



- New notebook

Create a notebook and switch to the created notebook interface.

- Create Paragraph

Add a new paragraph to the end of a notebook. A notebook can have up to 30 paragraphs.

- Save all paragraphs

Save all modified paragraphs.

- Delete notebook

Delete the current notebook. If a cluster has been associated, it will be disassociated.

- View

Only display codes or display codes and results.

## 5.6.1.3 Examples

### 5.6.1.3.1 Query bank employee information

#### 1. Create a temporary table

```
% spark
import org.apache.commons.io.IOUtils
import java.net.URL
import java.nio.charset.Charset
// Zeppelin creates and injects sc (SparkContext) and
sqlContext (HiveContext or SqlContext)
// So you don't need create them manually
// load bank data
val bankText = sc.parallelize (
 IOUtils.toString (
 new URL ("http://emr-sample-projects.oss-cn-hangzhou.aliyuncs.com/bank.csv"),
 Charset.forName ("utf8")).split ("\n"))
case class Bank (age: Integer, job: String, marital: String, education: String, balance: Integer)
val bank = bankText.map (s => s.split(";")).filter (s => s(0) != "\" age \"").map (
 s => Bank (s(0).toInt,
 s(1).replaceAll ("\"", ""),
 s(2).replaceAll ("\"", ""),
 s(3).replaceAll ("\"", ""),
 s(5).replaceAll ("\"", "").toInt)
).toDF ()
bank.registerTempTable ("bank")
```

#### 2. Query the table structure

```
% sql
desc bank
```

#### 3. Query the number of employees in each age group under 30

```
% sql select age, count(1) value from bank where age < 30 group by age order by age
```

#### 4. Query the information of employees younger than or equal to 20

```
% sql select * from bank where age <= 20
```

### 5.6.1.3.2 Video playback data

#### Preparations

In this example, you need to download data from OSS and upload it to your OSS bucket. This data includes:

- [User table sample data](#)
- [Video table sample data](#)

- [Video playback table sample data](#)

Upload this sample data respectively to the specified UserInfo, Videoinfo, and Playvideo on your OSS bucket. For example, upload the data to the Demo or UserInfo directory under Bucket Example.

In the following table, replace the SQL [bucketname] with your bucket name, replace [region] with your OSS region name, and replace [bucketpath] with your specified OSS path prefix, such as Demo.

#### 1. Create a user table

```
% hive
CREATE EXTERNAL TABLE user_info (id int , sex int , age
int , marital_status int) ROW FORMAT DELIMITED FIELDS
TERMINATED BY ',' LOCATION ' oss ://[bucketname]. oss - cn -[
region]- internal . aliyuncs . com /[bucketpath]/ userinfo '
```

#### 2. Create a video table

```
% hive
CREATE EXTERNAL TABLE video_info (id int , title string
, type string) ROW FORMAT DELIMITED FIELDS TERMINATED
BY ',' LOCATION ' oss ://[bucketname]. oss - cn -[region]-
internal . aliyuncs . com /[bucketpath]/ videoinfo '
```

#### 3. Create a video playback table

```
% hive
CREATE EXTERNAL TABLE play_video (user_id int , video_id
int , play_time bigint) ROW FORMAT DELIMITED FIELDS
TERMINATED BY ',' LOCATION ' oss ://[bucketname]. oss - cn -[
region]- internal . aliyuncs . com /[bucketpath]/ playvideo '
```

#### 4. Count the user tables

```
% sql select count (*) from user_info
```

#### 5. Count the video tables

```
% sql select count (*) from video_info
```

#### 6. Count the video playback tables

```
% sql select count (*) from play_video
```

#### 7. Count the video playbacks for each video type

```
% sql select video . type , count (video . type) as count
from play_video play join video_info video on (play .
```

```
video_id = video . id) group by video . type order by
count desc
```

## 8. Display the video information for the top 10 video playbacks

```
% sql select video . id , video . title , video . type ,
video_coun t . count from (select video_id , count (video_id
) as count from play_video group by video_id order
by count desc limit 10) video_coun t join video_info
video on (video_coun t . video_id = video . id) order by
count desc
```

## 9. Display the age of the viewers watching the video with the most video playbacks

```
% sql select age , count (*) as count from (select
distinct (user_id) from play_video where video_id = 49
) play join user_info userinfo on (play . user_id =
userinfo . id) group by userinfo . age
```

## 10. Display the gender, age, and marital status of the viewers watching the video with the most video playbacks

```
% sql select if (sex = 0 , ' Female ' , ' Male ') as title ,
count (*) as count , ' Gender ' as type from (select
distinct (user_id) from play_video where video_id = 49
) play join user_info userinfo on (play . user_id =
userinfo . id) group by userinfo . sex
union all
select case when userinfo . age < 15 then ' Less than
15 ' when age < 25 then ' 15 - 25 ' when age < 35 then
' 25 - 35 ' else ' More than 35 ' end , count (*) as
count , ' Age Group ' as type from (select distinct (
user_id) from play_video where video_id = 49) play join
user_info userinfo on (play . user_id = userinfo . id)
group by case when userinfo . age < 15 then ' Less than
15 ' when age < 25 then ' 15 - 25 ' when age < 35 then
' 25 - 35 ' else ' More than 35 ' end
union all
select if (marital_st atus = 0 , ' Unmarried ' , ' Married ') as
title , count (*) as count , ' Marital Status ' as type
from (select distinct (user_id) from play_video where
video_id = 49) play join user_info userinfo on (play .
user_id = userinfo . id) group by marital_st atus
```

## 5.6.2 Execution plans

### 5.6.2.1 Create an execution plan

An execution plan is a set of jobs that can be executed either at one time or periodically by means of scheduling. It can be executed on an existing E-MapReduce



cluster and can also create a temporary cluster to execute the jobs dynamically. Its biggest advantage is that it only uses the resources it needs during execution.

## Procedure

To create an execution plan, follow these steps:

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. Select a region.
3. In the upper-right corner, click Old MER Scheduling to go to the Jobs page.
4. In the navigation panel on the left, click Execution plan.
5. In the upper-right corner, click Create an execution plan.
6. In the Create an execution plan page, select between Create as needed and Existing clusters.
  - a. Create as needed: Create a new cluster to run jobs.
    - Execution plan for one-time scheduling: Clusters with corresponding configurations are created when the execution starts and are then released upon completion of the operation. For more information about creation parameters, see [Create a cluster](#).
    - Execution plan for periodic scheduling: A new cluster is created based on the scheduling settings you define and is then released upon completion of the operation.
  - b. Existing clusters: Use an existing cluster that complies with the following requirement:
    - Execution plans can only be added to clusters that are Running or Idle.Select Existing clusters and then enter the Select Cluster page. Here, you can select a cluster to associate with the execution plan.
7. Click Next to enter the job configuration page. All user jobs are listed in the table on the left. You can select jobs for execution from this table. By clicking the right-facing button, the checked jobs are added to the job queue. Jobs in the queue are then submitted to the cluster for execution in order. The same job can be added and executed several times. If you have not created any jobs, see [Jobs](#).

8. Click Next to enter the scheduling mode configuration page. The configuration items are as follows:
  - a. Name: Must be between 1-64 characters and may only consist of Chinese characters, English letters, numbers, hyphens (-), and underscores (\_).
  - b. Scheduling policy
    - Manual execution: The execution plan is not executed automatically after it is created. Instead, it must be executed manually. Once the execution is in progress, it cannot be executed again.
    - Periodic scheduling: If you select this function, it is enabled immediately after the execution plan is created. The execution then begins from the configured scheduling time. Periodic scheduling can be disabled in the list page. If a scheduling execution starts, but its last execution is not completed, the scheduling is ignored.
  - c. Set the scheduling cycle: There are two scheduling periods: days and hours. The day cycle is one day by default and cannot be changed. However, you can set a specific time interval for hours. The range must be 1-23.
  - d. First execution time: The effective start-time of the scheduling. From this point onwards, periodic scheduling is conducted according to the intervals specified.
9. Click OK to complete the creation of the execution plan.

## Other information

- Example of periodic scheduling

\* Set the scheduling cycle :

\* Set the scheduling cycle : per    day(s)

\* first execute time :     :

First run time 2018-11-1 17:35

Subsequent intervals 1 day(s) run 1 Times

These configurations indicate that the scheduling started on 11/01/2018 at 17:35 with an interval of one day. This means that the second scheduling was conducted on 11/02/2018, at 17:35.

- Sequence of jobs

Jobs in the execution plan are executed from first to last according to the sequence that you defined in the job list.

- Sequence of multiple execution plans

When multiple execution plans are submitted to the same cluster, each one submits jobs from its own job sequence. This means that jobs run parallel with each other.

- Example of early job debugging

During the debugging of a job, it may take some time to create and start a cluster on demand. We recommend that you create a cluster manually first, select Associate the cluster in the execution plan to run jobs, and then set the scheduling mode to Execute immediately. During debugging, you can view the results by clicking Run now on the execution plan list page. Once the debugging is finished, modify the execution plan, modify the way you associate an existing cluster to create a new cluster on demand, and then modify the scheduling mode to periodic scheduling as required. Jobs are then executed automatically on demand.

## 5.6.2.2 Manage an execution plan

You can view, manage, and modify your execution plans as follows.

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. Select a region.
3. In the upper-right corner, click Old MER Scheduling to go to the Jobs page.
4. In the navigation panel on the left, click Execution plan.
5. Click Manage next to a plan to go to the execution plan detail page. Here, you can perform the following operations:

- View details of the execution plan

You can view the basic information of the execution plan, such as its name, associated clusters, job configurations, scheduling mode and status, and alarm information.

- Modify the execution plan



**Notice:**

Jobs can only be modified if they are not currently running or being scheduled. For an execution plan to be executed immediately, it can only be modified when it is not currently running. If the execution plan is scheduled periodically, wait

for the completion of its current operation and verify whether it is in periodical scheduling. If it is, click Stop scheduling before modifying it.

Each separate module can be modified independently. Click the pen icon to modify information.

- Configure alarm notifications

There are three types of alarm notifications:

- Booting timeout: If the periodical scheduling has not been conducted correctly at the specified time and is not executed within 10 minutes of timeout, an alarm is sent.
- Failed execution: If any job in the execution plan fails, an alarm is sent.
- Successful execution: If all jobs in the execution plan are executed successfully, a notification is sent.

- Run and view results

If the execution plan can be run, in Basic Information, there will be a Run now button to the right of Scheduling status. If you click this button, a schedule will be executed.

At the bottom of the page, there are running records displaying the execution plan instances executed each time, making it easy to view the corresponding job list and logs.

### 5.6.2.3 Execution plan list

An execution plan list displays basic information about all of your execution plans, as shown in the following figure.

| ID/Name                       | Last run cluster | Last run                                                                                                                 | Scheduling status | Operation                                                                 |
|-------------------------------|------------------|--------------------------------------------------------------------------------------------------------------------------|-------------------|---------------------------------------------------------------------------|
| WF-63300C427CA72165<br>test1  | dtplus_docs      | Start Time : 2018/10/29 10:40:15<br>Running time : 4second(s)<br>Running Status : Complete                               |                   | <a href="#">Manage</a>   <a href="#">Run now</a>   <a href="#">More</a> ▼ |
| WF-C719251FFA366B3C<br>周期调度测试 | 高配测试             | Start Time : 2018/10/29 10:37:24<br>Running time : 8second(s)<br>Running Status : Complete                               | Scheduling paused | <a href="#">Manage</a>   <a href="#">Run now</a>   <a href="#">More</a> ▼ |
| WF-A388AA222DA06359<br>test1  | 高配测试             | Start Time : 2018/10/31 17:32:00<br>Running time :<br>1day(s)16hour(s)36minute(s)34second(s)<br>Running Status : Running | Scheduling        | <a href="#">Manage</a>   <a href="#">Run now</a>   <a href="#">More</a> ▼ |

- ID/Name: The ID and name of the execution plan.
- Last run cluster: The last cluster to execute this execution plan. This can either be a cluster created on demand or an existing associated cluster. If a cluster is created automatically on demand, (Automatically created) is displayed beneath

it, indicating that the cluster was created on demand by E-MapReduce and will be released automatically after running.

- **Last run:** The running status of the last execution plan.
  - **Start time:** The time at which the last execution plan started.
  - **Running time:** The duration for which the last plan ran.
  - **Running status:** The running status of the last execution plan.
- **Scheduling status:** This indicates whether scheduling is in progress or has been stopped. Only periodic jobs have a scheduling status.
- **Operation**
  - **Manage:** View and modify execution plans.
  - **Run now:** A job can only be run manually when it is neither running nor being scheduled. Click Run now to run the execution plan immediately.
  - **More**
    - **Start/Stop scheduling:** If the scheduling is stopped, Enable scheduling is displayed, which you can click to start the scheduling. If Stop scheduling is displayed during scheduling, you can click it to stop the scheduling. This button is only available for periodic execution plans.
    - **Running log:** Click to enter the job log viewing page.
    - **Delete:** Deletes an execution plan. A running execution plan or one in the process of scheduling cannot be deleted.

#### 5.6.2.4 View job results and logs

In this tutorial, you will learn how to view job results and logs.

##### View execution records

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. Select a region.
3. In the upper-right corner, click Old MER Scheduling to go to the Jobs page.
4. In the navigation panel on the left, click Execution plan.

## 5. To the right of the execution plan, click More > Running log.

Running log [Back to execution plan list](#) [Refresh](#)

ID/NameWF-A388AA222DA06359/test1

| Execution order ID | Running Status | Start Time          | Running time                           | Execute cluster      | Operation                                                     |
|--------------------|----------------|---------------------|----------------------------------------|----------------------|---------------------------------------------------------------|
| 7                  | Running        | 2018/10/31 17:32:00 | 1day(s)17hour(s)18minute(s)17second(s) | <a href="#">高配测试</a> | <a href="#">Stop all jobs</a>   <a href="#">View job list</a> |
| 6                  | Complete       | 2018/10/30 17:32:00 | 7second(s)                             | <a href="#">高配测试</a> | <a href="#">View job list</a>                                 |
| 5                  | Complete       | 2018/10/29 17:32:00 | 11second(s)                            | <a href="#">高配测试</a> | <a href="#">View job list</a>                                 |
| 4                  | Complete       | 2018/10/28 17:32:00 | 5second(s)                             | <a href="#">高配测试</a> | <a href="#">View job list</a>                                 |
| 3                  | Complete       | 2018/10/27 17:32:00 | 9second(s)                             | <a href="#">高配测试</a> | <a href="#">View job list</a>                                 |
| 2                  | Complete       | 2018/10/26 17:31:58 | 5second(s)                             | <a href="#">高配测试</a> | <a href="#">View job list</a>                                 |
| 1                  | Complete       | 2018/10/26 17:23:29 | 10second(s)                            | <a href="#">高配测试</a> | <a href="#">View job list</a>                                 |

- **Execution order ID:** The sequence of execution for the execution record, which indicates its position in the execution queue. For example, 1 stands for the first position.
- **Running status:** The running status of each execution record.
- **Start time** The time at which the execution plan starts.
- **Running time:** The total running time until the page is viewed.
- **Execute cluster:** The cluster run by the execution plan can either be created on demand or it can be an existing associated cluster. Click to view the cluster details page.
- **Operation**

**View job list:** Click to enter the job list page.

### View job records

On the Job list page, you can view the job list in the execution records of a single execution plan as well as the details of each job, as shown in the following figure.

| Job execution order ID | Name | Status | Type  | Start Time          | Running time | Operation                                                                                                |
|------------------------|------|--------|-------|---------------------|--------------|----------------------------------------------------------------------------------------------------------|
| WNE-92693C5408D39E03   | test | Failed | Spark | 2018/10/26 17:23:39 | 1second(s)   | <a href="#">Stop job</a>   <a href="#">stdout</a>   <a href="#">stderr</a>   <a href="#">workers log</a> |

- **Job execution order ID:** After a job is executed, a corresponding ID is created, which is different from the job ID. The job execution ID is the unique identifier for viewing logs on OSS.
- **Name:** The name of the job.
- **Status:** The running status of the job.
- **Type:** The type of job.

- **Start time:** The time at which the job starts. This is converted into local time.
- **Running time:** The total running time of the job, in seconds.
- **Operation**
  - **Stop job:** You can stop a job if it is in the process of submission or running. If a job is in submission, stopping it will cancel execution. If the job is running, it will be killed.
  - **stdout:** Records all output content from the standard output (Channel 1) of the master process. If log saving is not enabled for the cluster where jobs are run, this function cannot be executed.
  - **stderr:** Records all output content from the diagnostic output (Channel 2) of the master process. If log saving is not enabled for the cluster where jobs are run, this function cannot be executed.
  - **Workers log:** Views the logs of all job worker nodes. If log saving is not enabled for the cluster where jobs are run, this function cannot be executed.

#### View job worker logs

- **Cloud server instance IP:** The ECS instance ID of a running job and the corresponding intranet IP address.
- **Container ID:** The container ID that YARN runs.
- **Type:** Different log types. stdout and stderr come from different outputs.
- **Operation**

View the log: Click different types to view the corresponding logs.

### 5.6.2.5 Parallel execution of multiple execution plans

To maximize the use of a cluster's available computing resources, multiple execution plans can be associated to the same cluster and executed in parallel.

The main points are summarized as follows:

- Jobs in the same execution plan are executed in sequence. By default, preceding jobs are executed before new jobs can be submitted and executed.
- If you have enough cluster resources, you can create multiple different execution plans and associate them to the same cluster to run and execute jobs in parallel. Clusters support a maximum of 20 execution plans by default.
- The management and control system currently supports the submission to YARN of multiple execution plans associated to the same cluster. However, if the cluster



itself has insufficient resources, it may take some time for jobs in the YARN queue to wait for scheduling.

For more information on how to create execution plans and associate them to a cluster, see [Create an execution plan](#).