

Alibaba Cloud E-MapReduce

Quick Start

Issue: 20190220

Legal disclaimer

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








1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company, or individual in any form or by any means without the prior written consent of Alibaba Cloud.
3. The content of this document may be changed due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary, incidental, special, or punitive damages, including lost profits arising from the use

or trust in this document, even if Alibaba Cloud has been notified of the possibility of such a loss.

5. By law, all the content of the Alibaba Cloud website, including but not limited to works, products, images, archives, information, materials, website architecture, website graphic layout, and webpage design, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of the Alibaba Cloud website, product programs, or content shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates).
6. Please contact Alibaba Cloud directly if you discover any errors in this document.

Generic conventions

Table -1: Style conventions

Style	Description	Example
	This warning information indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	 Danger: Resetting will result in the loss of user configuration data.
	This warning information indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	 Warning: Restarting will cause business interruption. About 10 minutes are required to restore business.
	This indicates warning information, supplementary instructions, and other content that the user must understand.	 Notice: Take the necessary precautions to save exported data containing sensitive information.
	This indicates supplemental instructions, best practices, tips, and other content that is good to know for the user.	 Note: You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
Bold	It is used for buttons, menus, page names, and other UI elements.	Click OK.
Courier font	It is used for commands.	Run the <code>cd /d C:/windows</code> command to enter the Windows system folder.
<i>Italics</i>	It is used for parameters and variables.	<code>bae log list --instanceid Instance_ID</code>
[] or [a b]	It indicates that it is an optional value, and only one item can be selected.	<code>ipconfig [-all -t]</code>

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

Contents

Legal disclaimer.....	I
Generic conventions.....	I
1 Prerequisites.....	1
2 Create E-MapReduce.....	3
2.1 E-MapReduce quick start.....	3
2.2 Create a cluster.....	5
2.3 Create a job.....	9
2.4 Create an execution plan.....	13
3 Select configuration.....	16

1 Prerequisites

Make sure that the following prerequisites are met before you create an E-MapReduce cluster:

1. Create an Alibaba Cloud account

To create an E-MapReduce cluster, you must have an Alibaba Cloud account, which is used to uniquely identify you in the Alibaba Cloud ecosystem. You can use this account to create E-MapReduce clusters and activate other Alibaba Cloud services, including *Object Storage Service (OSS)* and *ApsaraDB for RDS (RDS)*.

For more information about creating an Alibaba Cloud account, see [Sign up with Alibaba Cloud](#).

2. Create an AccessKey (optional)

To use E-MapReduce, you must create a minimum of one AccessKey. Follow these steps to create an AccessKey:

- a. Log on to the [Alibaba Cloud website](#).
- b. Go to the Alibaba Cloud console.
- c. Hover over your profile picture and then click AccessKeys.



Notice:

If the following dialog box appears, click Continue to manage AccessKey .

Security Tips



AccessKey of your cloud account is the secret key to access Alibaba Cloud APIs. Since the AccessKey has full permissions of your cloud account, please make sure you keep it well. To avoid the AccessKey being used by others to cause [Sensitive information leakage](#) , do not release your AccessKey to any external channels (for example, Github). We strongly recommend you use the AccessKeys of RAM users in API calls, according to [Alibaba Cloud account security best practices](#) .

Continue to manage AccessKey

Get Started with Sub Users's AccessKey

- d. Click Create AccessKey.
- e. The AccessKey has been created.

3. Activate Alibaba Cloud OSS

E-MapReduce stores the job and execution logs on Alibaba Cloud OSS. Therefore, you must activate Alibaba Cloud OSS. For more information, see [Sign up for OSS](#).

Create OSS buckets in the same region as the EMR cluster that you need to create.

For more information, see [Create a bucket](#).

4. Create high-configuration instances (optional)

If you need to create instances with eight or more cores for a Pay-As-You-Go-based cluster, log on to your Alibaba Cloud account and submit a ticket for application.

For more information, see [Support and Services](#).

2 Create E-MapReduce

2.1 E-MapReduce quick start

This section describes clusters and jobs in E-MapReduce (EMR) and how to use them. For example, you can create a Spark job, run the job on the cluster to calculate Pi (π), and then view the result in the console.



Notice:

Note: Confirm that all *prerequisites* have been met.

1. Create a cluster

- a. In the *Alibaba Cloud E-MapReduce console*, click the Cluster tab to go to the clusters list page. Click Create Cluster in the upper-right corner.
- b. Software Configuration
 - A. Select the latest EMR version. Example: EMR 3.13.0.
 - B. Select the default software configuration.
- c. Hardware Configuration
 - A. Select Pay-As-You-Go.
 - B. If no security group has been created, enter a name and then create one.
 - C. Select a master instance with 4 cores and 8 GB of memory.
 - D. Select two core instances with 4 cores and 8 GB of memory.
 - E. The remaining configuration all uses the default settings.
- d. Basic Configuration
 - A. Enter the name of the cluster.
 - B. Specify a path to store the job log. Make sure that the running log feature is enabled. *Create an OSS bucket* in the region where the cluster has been created.
 - C. Enter the password that is used to log on to the cluster.
- e. Click OK to create the cluster.

2. Create a job.

- a. Click the Data Platform tab to go to the project list page. Click New Project in the upper-right corner.

- b. In the New Project dialog box, enter the project name and description, and then click Create.
- c. Click Design Workflow to the right of the specified project and to go to the Edit Jobs page.
- d. On the left side of the Edit Jobs page, right-click the folder that you want to operate and select New Job.
- e. Enter the job name and description.
- f. Select Spark as the job type.
- g. Click OK.



Note:

You can also right-click a folder and then choose to create a subfolder, rename the folder, or delete the folder.

- h. Enter parameters as follows:

```
--class org.apache.spark.examples.SparkPi --master yarn-client  
--driver-memory 512m --num-executors 1 --executor-memory 1g  
--executor-cores 2 /usr/lib/spark-current/examples/jars/spark-  
examples_2.11-2.1.1.jar 10
```



Note:

The `/usr/lib/spark-current/examples/jars/spark-examples_2.11-2.1.1.jar` JAR file name is defined by the version of Spark in the cluster. For example, if the version of Spark is 2.1.1, then the JAR file is named as `spark-examples_2.11-2.1.1.jar`. If the version of Spark is 2.2.0, then the JAR file is named as `spark-examples_2.11-2.2.0.jar`.

- i. Click Run.
3. View the job log entries and confirm the results.

After you have executed a job, you can click the Log tab at the bottom of the page to view running log. Click View Details to go to the details page. On this page, you can view details, including the job submission log and YARN Container log.

2.2 Create a cluster

This section describes how to create and configure an E-MapReduce cluster.

Go to the cluster creation page

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. Complete RAM authorization. For more information, see [Role authorization](#).
3. Select a region where you want to create the cluster. The region cannot be changed after the cluster has been created.
4. Click Create Cluster in the upper-right corner.

Create a cluster

To create a cluster, follow these steps:

- Software configuration
- Hardware configuration
- Basic configuration

Step 1: Software configuration

Description

- EMR Version: By default, the latest version is selected.
- Cluster Type: Currently, E-MapReduce (EMR) provides the following cluster types:
 - Hadoop clusters. These clusters provide multiple ecosystem components, such as Hadoop, Hive, Spark, Spark Streaming, Flink, Storm, Presto, Impala, Oozie, and Pig. Hadoop, Hive, and Spark are semi-hosted services and are used for distributed large-scale data storage and computing. Spark Streaming, Flink, and Storm can provide stream computing. Presto and Impala are used to achieve interactive queries. For more information about the components, see Services List displayed on the cluster and service management page.
 - Kafka clusters. These clusters serve as a semi-hosted distributed message system with high throughput and high scalability. Kafka clusters provide a comprehensive service monitoring system that keeps the clusters running stably. Kafka clusters are more professional, reliable, and secure. You do not need to deploy or maintain these clusters. These clusters are commonly used in scenarios such as log collection and monitoring data aggregation. Offline

data processing and stream computing, and real-time data analysis are also supported.

- **Druid clusters.** These clusters provide semi-hosted and real-time interactive analysis services. Druid clusters also support querying large amounts of data in milliseconds and writing data through multiple methods. Druid clusters offer flexible and stable real-time queries together with other services such as EMR Hadoop, EMR Spark, OSS, and RDS.
- **Data Science clusters.** These clusters are commonly applied in big data and AI scenarios. Data Science clusters provide Hive and Spark offline big data ETL, and TensorFlow model training. You can choose the CPU plus GPU heterogeneous computing framework and the deep learning algorithm supported by NVIDIA GPUs to run computing tasks efficiently.
- **Required Services:** Select the default configuration. You can add, enable, or disable services on the management page later.
- **High Security Mode:** Indicates whether to enable the Kerberos authentication feature for the cluster. This mode is disabled by default because typically this mode is not required by clusters created for general users.
- **Enable Custom Setting:** You can specify a JSON file to modify software configuration. For more information about the procedure, see [Software configuration](#).

Step 2: Hardware configuration

Description

- **Billing Configuration**
 - **Billing Method:** You can select Pay-As-You-Go when testing the cluster. After all tests have been passed, you can create and use a Subscription-based cluster.
- **Network configuration**
 - **Zone:** Typically, the default zone is used.
 - **Network Type:** VPC is selected by default. If you have not created a VPC, go to the [VPC console](#) to create one.
 - **Zone:** Geographical areas that are located in the same region. These zones are interconnected through VPCs.
 - **VPC:** Select a VPC that has been created in the specified region. If no VPC is available, click Create VPC/VSwitch to create one in the current zone.

- VSwitch: Select a VSwitch for the specified VPC in the current zone. If no VSwitch is available, go to the VPC console and create one in the current zone.
- Security Group Name: By default, no security group is available if this is the first time that you have created a cluster. You need to enter a name and then create a security group. If you have already created security groups, select a security group.
- Cluster configuration
 - High Availability: After this feature has been enabled, two master nodes will be provided to guarantee the high availability of Resource Manager and Name Node. HBase clusters support high availability by default. An HBase cluster has to use a core node as one of the two master nodes. When the high availability feature is enabled, the HBase cluster only needs one master node to support high availability, which is more secure and reliable. If you need to create clusters that support high availability, enable high availability during testing.
 - Master node: Deploys processes, such as Resource Manager and Name node.
 - Master Instance Type: Select instance specifications based on your needs. For more information, see [Instance type families](#).
 - System Disk Type: Select Ultra Disk or SSD Disk based on your needs.
 - System Disk Size: You can resize the disk based on your needs. The recommended minimum disk size is 120 GB.
 - Data Disk Type: Select Ultra Disk or SSD Disk based on your needs.
 - Data Disk Size: You can resize the disk. The recommended minimum disk size is 80 GB.
 - Master Instances: The default number of master instances is one.
 - Core instance: Stores all cluster data. You can scale up the instance based on your needs.
 - Core Instance Type: Select instance specifications based on your needs. For more information, see [Instance type families](#).
 - System Disk Type: Select Ultra Disk or SSD Disk based on your needs.
 - System Disk Size: You can resize the disk based on your needs. The recommended minimum disk size is 80 GB.
 - Data Disk Type: Select Ultra Disk or SSD Disk based on your needs.

- **Data Disk Size:** You can resize the disk based on your needs. The recommended minimum disk size is 80 GB.
- **Core Instances:** The default number of core instances is two. You can adjust the number of core instances based on your needs.
- **Task Instance Group:** No data is stored on task instance groups. Task instance groups are used to adjust the computing capability of clusters. This feature is disabled by default. You can enable it based on your needs.

Step 3: Basic configuration

Description

- **Basic Information**
 - **Cluster Name:** Enter the name of the cluster. It must be 1-64 characters in length and can contain Chinese characters, uppercase English letters, lowercase English letters, numbers, hyphens (-), and underscores (_).
- **Running Logs**
 - **Running Logs:** You can enable or disable this feature. This feature is enabled by default. After you enable this feature, you must specify an OSS path to save the logs. All running logs are stored in the specified OSS directory. To enable this feature, you must first activate OSS. The uploaded data is billed based on the space usage. We recommend that you enable this feature, which helps in debugging and troubleshooting.
 - **Log Path:** Specify an OSS path to store the log.
 - **Uniform Meta Database:** Hive uses a unified metadatabase, which is independent of the cluster. The meta information will not be deleted after the cluster is released. We recommend that you disable this feature if you are an E-MapReduce beginner.
- **Permission Settings:** Typically, the default setting is used.
- **Logon Settings**
 - **Remote Logon:** Indicates whether to open port 22 of the security group. This feature is enabled by default.
 - **Password:** Set the password used to log on to the master instance. The logon password must be 8 to 30 characters in length and can contain uppercase letters, lowercase letters, numbers, and special characters including exclamation

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

- **Bootstrap Operation (optional):** You can set the cluster to run your custom script before it starts Hadoop. For more information, see Bootstrap action.

Configuration List

Confirm the configured items and fees in the configuration list.

Confirm the creation

After you configure the settings and make sure that all settings are valid, the Create button is highlighted. Confirm the information, and click Create to create a cluster.



Notice:

- The cluster is created immediately if the billing method is Pay-As-You-Go. You are directed to the cluster list page. You can find an initializing cluster in the Clusters list. It may take several minutes to complete the cluster creation. After the cluster has been created, it changes its status to Idle.
- An order is generated if the billing method is Subscription. The cluster will be created after you complete the payment.

Creation failed

If the creation fails, CREATE_FAILED is displayed on the clusters list. Move the pointer over the red exclamation mark (!) to view the cause.

You do not to handle clusters that are failed to be created because no computing resources have been created for these clusters. These clusters will be automatically hidden after three days.

2.3 Create a job

This section describes how to create a job for scheduling in an early E-MapReduce version.

To run a computing task, you need to follow these steps to define a job:

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. Select a region where you want to create the job.
3. Click the Old EMR Scheduling tab to go to the jobs list page.

- Click **Create job** in the upper-right corner to go to the job creation page, as shown in the following figure:

Create job

×

* Name :

Length: 1 to 64 characters. Only Chinese characters, English letters, numbers '-', and '_' are allowed

* Type :

☒ Spark

☐ Hadoop

☐ Hive

☐ Pig

☐ Sqoop

☐ Spark SQL

☐ Shell

* Parameter :

+ Select OSS path oss console Upload

* Actual execution :

spark-submit

Fail retry

☐ Yes

☒ No

* Failure policy :

☐ Pause current execution plan

☒ Continue execution of next job

OK

Cancel

- Enter the job name.
- Select a job type.
- Enter parameters of the job. The application parameters must include the following information: the JAR package run by the job, data input and output addresses of the job, and certain command line parameters. You must enter all parameters in the command line in this field. If you need to use an OSS path, click

Select OSS Path to select the OSS resource path. For more information about the parameters of all job types, see the *Job* section in *User Guide*.

8. Actual execution. The job command that has been executed on ECS will be displayed in this field. You can copy the displayed command and directly run it in the command line environment on an E-MapReduce cluster.
9. Fail Retry. You can set the number of retries and their intervals. This feature is disabled by default.
10. Select a failure policy. Pause the current execution plan: Indicates that the current execution plan will be paused if this job fails and will wait for you to process.
Continue execution of the next job: Indicates that errors will be ignored and the next job will be executed after this job fails.
11. Click OK to complete the creation.

Example

This is a Spark job, which sets parameters such as the input and output paths in the application parameters.



Notice:

This example is for reference only.

Create job

✕

* Name :

SparkPi

Length: 1 to 64 characters. Only Chinese characters, English letters, numbers '-', and '_' are allowed

* Type :

☒ Spark
 ☐ Hadoop
 ☐ Hive
 ☐ Pig
 ☐ Sqoop
 ☐ Spark SQL
 ☐ Shell

* Parameter :

```
spark-submit --class org.apache.spark.examples.SparkPi --master
yarn-client --driver-memory 512m --num-executors 1 --executor-
memory 1g --executor-cores 2 /opt/apps/spark-1.6.1-bin-
hadoop2.7/lib/spark-examples-1.6.1-hadoop2.7.2.jar 100
```

+ Select OSS path

* Actual execution :

```
spark-submit spark-submit --class
org.apache.spark.examples.SparkPi --master yarn-client --driver-
memory 512m --num-executors 1 --executor-memory 1g --
executor-cores 2 /opt/apps/spark-1.6.1-bin-hadoop2.7/lib/spark-
examples-1.6.1-hadoop2.7.2.jar 100
```

* Failure policy :

☐ Pause current execution plan
 ☒ Continue execution of next job

OSS and ossref

The `oss://` prefix indicates that the data path points to an OSS path, which specifies the operation path when reading/writing the data. This is similar to `hdfs://`.

The `ossref://` prefix also indicates that the data path points to an OSS path. However, it is used to download the corresponding code to a local disk, and then replace the path in the command line with this local path. It is easier for you to run native code. You do not need to log on to the computer to upload the code and the dependent resource packages.

In this example, the `ossref://xxxxxx/xxx.jar` parameter represents the JAR package of job resources. This JAR package is stored on OSS. When this path is executed in the code, the JAR package will automatically download to the cluster and be executed. The two `oss://xxxx` and the two values following the JAR package are processed by the main class in the JAR package as parameters.



Notice:

The ossref cannot be used to download large amounts of data. Otherwise, an error may occur.

2.4 Create an execution plan

This section describes how to create an execution plan in an earlier E-MapReduce version.

After you have created a job, you need to create an execution plan if you want to run a predefined job on a cluster. An execution plan may contain multiple jobs. You can define the execution sequence of these jobs. For example, if a scenario handles data in the following order: prepare data, process data, and clean up data, then you can define the following three jobs: prepare-data, process-data, and cleanup-data, and then you can create an execution plan that includes these three jobs.

To create an execution plan, follow these steps:

1. Log on to the [Alibaba Cloud E-MapReduce console](#).
2. Select a region.
3. Click Old EMR Scheduling.
4. Click the left-side Execution Plan tab to go to the execution plan page.
5. Click Create an execution plan in the upper-right corner to go to the execution plan creation page.
6. The following two options are available on the Select the cluster page: Create as needed and Existing Clusters. Create as needed indicates that you have not created any clusters. You will run the execution plan in a temporary cluster, which will be automatically released after you have executed the plan. Existing Cluster indicates that you have at least one running cluster, and you have to submit the execution plan to the existing cluster for execution.
 - a. If you select Create as needed, you can follow the same steps as creating a cluster to create an on-demand cluster. Select configuration of the on-demand cluster, and click OK.
 - b. If you select Existing Cluster, you will be directed to the cluster selection page. You can select clusters that you want to associate with the execution plan.

×

Create an execution plan

1 : Select the
2 : Select the
3 : Configure the
4 : Configure the

* Associated cluster :

* Cluster info :

HBase

Status : Idle
 Type : Hadoop
 Running time : 5hour(s)36minute(s)43second(s)
 created time : 2018/10/24 14:57:13

Prev

Next

Cancel

Currently, you can only submit clusters that are in Running or Idle status.

- Click Next to go to the Configure the job page. The page displays the predefined jobs on the left-side list, and the jobs to be executed by the newly created execution plan on the right-side list. To define the execution plan, select jobs from the left-side list based on the execution sequence and add them to the right-side list. You can click the question mark (?) to view parameters of the jobs. After you have defined the execution plan, click Next.

×

Create an execution plan

1 : Select the cluster mode
2 : Select the cluster
3 : Configure the job
4 : Configure the scheduling

Job list

ID

Search

ID	Name	Type
J-224FB0D8 84AE6DEF	test_wordcount (?)	Hadoop
J-6C48C6EB 9D5AF93A	Pi (?)	Spark
J-B21F6DBC 51A68F7F	wcr (?)	Hadoop

>

<

Configured job

ID	Name	Type
----	------	------

^

v

Prev

Next

Cancel

8. Enter an execution plan name.

9. Select a scheduling policy.

- Manually executed indicates that the execution plan will be executed only when you click the execution plan.
- Periodic scheduling execution plan indicates that you must set the scheduling cycle and the scheduling start time.

Create an execution plan ✕

1 : Select the

2 : Select the

3 : Configure the

4 : Configure the

* Name :

* Scheduling policy :

Periodic scheduling execution plan ▼

* Set the scheduling cycle :

day(s) ▼

* Set the scheduling cycle :

per

1 ▲ ▼

 day(s)

* first execute time :

2018-10-25

09 ▲ ▼

 :

47 ▲ ▼

First run time 2018-10-25 9:47

Subsequent intervals1 day(s) run1Times

When the execution plan is completed, you can make the relevant settings for the execution plan in the execution plan management page.

Prev

OK

Cancel

10. Click OK to complete the creation of the execution plan.

3 Select configuration

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

Big data scenarios

EMR has been applied in the following enterprise big data scenarios:

- **Batch operations.** This scenario requires high disk I/O throughput and high network throughput, but has low real-time capacity requirements. If you need to process large amounts of data but do not require real-time processing, you can use MapReduce, Pig, and Spark. This scenario does not require a high memory capacity. Therefore, you need to focus on the CPU, memory, and network requirements when you perform data shuffling.
- **Ad hoc queries.** Data scientists and data analysts use ad hoc queries to retrieve data. This scenario requires real-time queries, high disk I/O throughput, and network throughput. In this scenario you can use Impala and Presto. This scenario also has high memory requirements. You need to consider the amount of data and concurrent queries.
- **Stream computing, high network throughput, and compute-intensive scenarios.** For these scenarios, you can use Flink, Spark Streaming, and Storm.
- **Message queues.** This scenario requires high disk I/O throughput and network throughput, consumes large amounts of memory, and the storage does not depend on HDFS. Therefore, you can choose to use Kafka clusters in EMR. EMR clusters are divided into Kafka clusters and Hadoop clusters to avoid impacting Hadoop.
- **Cold backup.** This scenario does not require high disk I/O throughput or computing throughput and is low cost. We recommend that you use EMR d1 instances for cold backups. The storage cost of d1 instance is 0.03 USD/month/GB.

EMR instances

An EMR cluster consists of *three types of instances*: master instances, core instances, and task instances.

You can select ultra disks, SSD disks, and local disks for EMR storage. Performance of different disks is: SSD disks > local disks > ultra disks.

EMR underlying storage supports OSS (standard OSS only) and HDFS. OSS has a higher data availability than HDFS. The data availability of OSS is 99.99999999%, while the data availability of HDFS depends on the reliability of cloud disk or local disk storage.

Storage prices are as follows:

- Instance with local disks: USD 0.003/GB/month.
- OSS standard storage: 0.02 USD/GB/month.
- Ultra disk storage: 0.05 USD/GB/month.
- SSD disk storage: 0.143 USD/GB/month.

Select EMR configuration

- Select master instance configuration
 - Master instances are used to deploy the master processes of Hadoop, such as NameNode and ResourceManager.
 - We recommend that you enable high availability on production clusters. High availability is available for EMR components such as HDFS, YARN, Hive, and HBase. We recommend that you enable high availability in the "Cluster configuration" step when you create the cluster. If high availability is not enabled when you create an EMR cluster, it cannot be enabled later.
 - Master instances are used to store HDFS metadata and component log files. They are compute-intensive with low disk I/O requirements. HDFS metadata is stored in memory. The minimum recommended memory size is 16 GB based on the number of files.
- Select core instance configuration
 - Core instances are used to store data, run computing tasks and processes such as DataNode and NodeManager.
 - If the amount of data stored in HDFS (3 backups) exceeds 60 TB, we recommend that you use instances with local disks (ECS d1instances and ECS d1ne instances). The local disk capacity is calculated as follows: $(\text{the number of CPU cores}/2) * 5.5 \text{ TB} * \text{the number of instances}$. For example, if you purchase four d1 instances with 8 cores, the local disk capacity is: $8/2 * 5.5 * 4=88 \text{ TB}$. HDFS requires three backups. Therefore, you need to buy at least three instances that use local disks. We recommend that you purchase at least four instances for data reliability and disk recovery.

- If the amount of data stored in HDFS is less than 60 TB, you can use ultra disks or SSD disks.
- Select task instance configuration
 - Task instances are used when the CPUs and memory of core instances have insufficient computing capability. Task instances do not store data or run DataNode. You can estimate the number of instances based on CPU and memory requirements.

EMR lifecycle

EMR supports auto-scaling, which allows you to quickly *scale up a cluster*. You can flexibly adjust the configuration of cluster nodes, or *upgrade or downgrade ECS instance configuration*.

Select a zone

We recommend that you deploy EMR and your business system in *the same zone and region*.