

# Elastic MapReduce

## FAQs

### Product Documentation



## Copyright Notice

©2013-2024 Tencent Cloud. All rights reserved.

Copyright in this document is exclusively owned by Tencent Cloud. You must not reproduce, modify, copy or distribute in any way, in whole or in part, the contents of this document without Tencent Cloud's the prior written consent.

## Trademark Notice

 Tencent Cloud

All trademarks associated with Tencent Cloud and its services are owned by Tencent Cloud Computing (Beijing) Company Limited and its affiliated companies. Trademarks of third parties referred to in this document are owned by their respective proprietors.

## Service Statement

This document is intended to provide users with general information about Tencent Cloud's products and services only and does not form part of Tencent Cloud's terms and conditions. Tencent Cloud's products or services are subject to change. Specific products and services and the standards applicable to them are exclusively provided for in Tencent Cloud's applicable terms and conditions.

---

# Contents

## FAQs

- Basic EMR Issues

- Billing

- Failed Operations on EMR Master Node Due to Low Configuration

- Cluster Network Settings

- Unhealthy guidance for Yarn nodemanager nodes

- HiveServer2 migrates to Router

- InternalError in the Console

- EMR on CVM

  - Billing

  - Cluster Management

  - Reasons Why a Triggered Auto Scaling Rule Is Not Executed

# FAQs

## Basic EMR Issues

Last updated : 2022-07-04 11:23:37

### How do I view task logs?

You can log in to any EMR server and run the following command to view task logs:

```
yarn logs -applicationId application_1507732460084_0057
```

Note :

You need to run this command with the `Hadoop` username. If it is a task of another user, you can add the `appOwner username` parameter.

To view the cause of a task exception, run the following command:

```
yarn logs -applicationId application_1507732460084_0057|grep -A20 Exception
```

### How do I adjust the computing resources of a cluster?

Cluster computing resources are determined by the following two configuration items in `yarn-site.xml` :

```
<property>
<name>yarn.nodemanager.resource.cpu-vcores</name>
<value>4</value>
</property>
<property>
<name>yarn.nodemanager.resource.memory-mb</name>
<value>14745</value>
</property>
```

By default, `cpu-vcores` is equal to number of CPU cores of the server, and `memory-mb` is equal to 91% of the memory size of the server. You can adjust them based on your actual needs, but if they are too large, there may be a risk of server failure.

### How do I fix an out of memory error?

If an out of memory error occurs when you are submitting a MapReduce task or running a SQL script through Hive, fix it by setting the following parameters:

```
set mapreduce.map.java.opts=-Xmx4096m;
set mapreduce.reduce.java.opts=-Xmx4096m;
```

The memory parameter can be adjusted based on your actual computation needs. It can also be written in the `~/.hiverc` file in Hive and will be executed automatically when submitted.

## How do I estimate the cluster size?

Suppose that you need to run a SQL query. If 64 vcores and 128 GB memory are needed for getting the query result in the specified time period, and the business requires 10 concurrencies, then the required resources would be 640 vcores and 1,280 GB memory. If the server specification you are using is 24 cores and 48 GB memory, then you need around  $1280 / 48 = 27$  servers.

## How do I set up a fetch query in Hive?

The default query in Hive is as follows:

```
select * from tablename where a='1' limit 10;
```

The default query does not start a computation task. You can start a distributed query by adding the `set hive.fetch.task.conversion=none` parameter.

## How do I choose the cluster storage media?

An EMR cluster supports the following storage media: HDD local disk, SSD local disk, HDD cloud disk, SSD cloud disk, and COS. You can choose the most appropriate one based on your actual needs:

- If your application scenario is large-scale data warehouse analysis which is not very latency-sensitive, we recommend you use COS as the underlying storage.
- If you are very familiar with HDFS and the cost of migration to COS would be too high, you can use HDD cloud disks.
- If your application is a massive columnar database (HBase) which requires efficient writing and querying, we recommend you use SSD local or cloud disks.

## Can I log in to an EMR cluster over the public network?

You can log in to the cluster after enabling the public IP as instructed in [Logging in to Clusters](#).

## Does an EMR cluster support LDAP authentication?

LDAP authentication is subject to the product version. On v2.3.0 and later, it is supported and enabled by default and cannot be disabled. It is not supported on earlier versions.

### **Can I add more Common nodes?**

No.

### **Does EMR support ClickHouse version upgrade?**

At present, there are three releases in an EMR ClickHouse cluster, and ClickHouse has been upgraded in each release. For more information, see [Version Overview](#).

### **Will the big data components selected in the EMR console be automatically installed?**

Components can be directly installed in the console.

### **Does EMR HDFS enable WebHDFS by default?**

Yes.

### **Can I modify the project of an EMR cluster?**

No. If this affects your business, [submit a ticket](#) for assistance.

### **Can I run backend JAR packages on MapReduce nodes?**

When creating a cluster, you can use bootstrap actions to set a custom bootstrap script to achieve this (after node initialization, before cluster start, or after cluster start). For more information, see [Bootstrap Action](#).

### **How do I select cluster specifications based on business use cases?**

EMR provides six types of clusters for you to choose from based on your business needs. For more information, see [Business Evaluation](#).

# Billing

Last updated : 2022-08-12 14:59:19

## **When I purchase a pay-as-you-go cluster, why am I prompted that my account balance is insufficient even though it is actually sufficient?**

As some pay-as-you-go clusters are costly and fees will be frozen, you may be prompted that your account balance is insufficient at the time of purchase.

## **Can I change a non-HA EMR cluster to an HA one through configuration upgrade or scale-out?**

No. You need to purchase a new one.

## **Will EMR servers be displayed in CVM bills after fees are deducted by EMR?**

EMR servers in the format of `emr-vm-xxxxxxx` will be displayed in both EMR and CVM bills, but fees will be displayed as 0 in CVM bills.

# Failed Operations on EMR Master Node Due to Low Configuration

Last updated : 2021-09-23 09:26:33

## How do I fix failed operations on the EMR master node due to low configuration?

### Symptoms

As the master node's configuration is too low, Hive or Spark jobs submitted to it report errors or are directly killed.

### Cause analysis

The memory of the master node is insufficient, causing other applications to be killed due to OOM.

### Solution

1. Too many businesses are deployed on the EMR master node, which usually becomes the bottleneck of the entire cluster. However, the master node cannot be scaled out; instead, it can only be upgraded as described below:

- First, find the node where the standby NameNode resides in the cluster.
  - Run the following command on the standby NameNode to enter the safe mode.

```
hdfs dfsadmin -fs 10.0.0.9(standby node IP):4007 -safemode enter Enter the safe mode
```

- Run the following command on the standby NameNode to save the metadata.

```
hdfs dfsadmin -fs 10.0.0.9(standby node IP):4007 -saveNamespace Save the metadata
```

- Run the following command on the standby NameNode to exit the safe mode.

```
hdfs dfsadmin -fs 10.0.0.9(standby node IP):4007 -safemode leave Exit the safe mode
```

- Then, in the EMR Console (or the CVM Console for a legacy cluster), upgrade the active node.
- Upgrade the standby node and make the configuration of the master's active node the same as that of the standby node.

If your cluster is not a high-availability one, then it will become unavailable for a while during the upgrade.

2. In Spark, jobs are committed in client mode by default, and the driver runs on the master node. You can change the mode to master mode and then commit jobs.



- For the Hive component, enable the router node, migrate HiveServer2 to it, and then disable the Hive component on the master node. For detailed directions, please see [Migrating HiveServer2 to Router](#).
- Disable components that are not commonly used on the master node or migrate Hue to the router node.

Directions for migrating Hue to the router node:

- Enter the EMR Console, Add a router node on the Cloud Hardware Management page, and select the Hue component.
- After the scale-out, disable the original Hue component on the master node, retain that on the router node, bind a public EIP to the router node, and open the source policy and ports in the security group.

## Preset values of memory size for master node components in EMR cluster and recommendations

### 1. List of heap memories of common components

Component	Process	Configuration File	Configuration Item	Default Heap Memory (in MB)
HDFS	Namenode	hadoop-env.sh	NNHeapsize	4,096
YARN	Resourcemanager	yarn-env.sh	Heapsize	2,000
Hive	Hiveserver2	hive-env.sh	HS2Heapsize	4,096
HBase	Hmaster	hbase-env.sh	Heapsize	1,024
Presto	Coordinator	jvm.config	Maximum JVM	3,072
Spark	spark-driver	spark-defaults.conf	spark.driver.memory	1,024
Oozie	Oozie	-	-	1,024
Storm	Nimbus	-	-	1,024

### 2. Suggested preset values for components

Component	Suggested Heap Memory Size
HDFS (NameNode)	Minimum heap memory = 250 x number of files + 290 x number of directories + 368 x number of blocks
YARN (ResourceManager)	It can be increased as needed
Hive (HiveServer2)	It can be increased as needed
HBase (HMaster)	The master node only receives DDL requests and performs load balancing. The default size of 1 GB is generally sufficient

---

Presto (Coordinator)	Use the default value
Spark (spark-driver)	It can be increased as needed
Oozie (oozie)	Use the default value
Storm (Nimbus)	Use the default value

3. Suggested idle memory size for servers: 10–20% of the total memory size.

4. You can deploy EMR components in independent mode or hybrid mode as needed.

- Independent deployment: it is suitable for HDFS clusters for storage, HBase clusters for analysis of massive amounts of data, and Spark clusters for job computation.
- Hybrid deployment: multiple components can be deployed in a cluster in this mode, which is suitable for testing clusters or scenarios where the business volume is not high or resource preemption is negligible.

# Cluster Network Settings

Last updated : 2020-04-21 09:45:56

**What should I do if the error "No subnet in the AZ selected for the cluster" is displayed when I try to set the cluster network on the purchase page?**

**Cluster Public Network**  Enable public network of cluster's Master node

Enable the public network of the cluster's Master nodes, which is mainly used for SSH login and viewing the [component webui](#).

Master node Master. 1. The public network of the node will be enabled and billed by traffic. The maximum bandwidth is 5 Mbps. After the cluster is created, you can adjust the network in the console.

**Cluster network**    Total 0 subnet IPs, 0 available

If the existing networks do not meet your needs, you can [create a VPC](#) or [create a subnet](#) on the console.

(Under the current network selection, only Default-VPC VPC devices can access this cluster.)

A VPC can be deployed across AZs, and all VPCs in the current region will be loaded by default. A subnet belongs to an AZ, and all existing subnets in the AZ where the cluster resides will be loaded by default. Please check whether there are available subnets in the selected AZ.

# Unhealthy guidance for Yarn nodemanager nodes

Last updated : 2020-10-13 15:51:31

## How do I fix an unhealthy node in YARN NodeManager?

### Symptoms

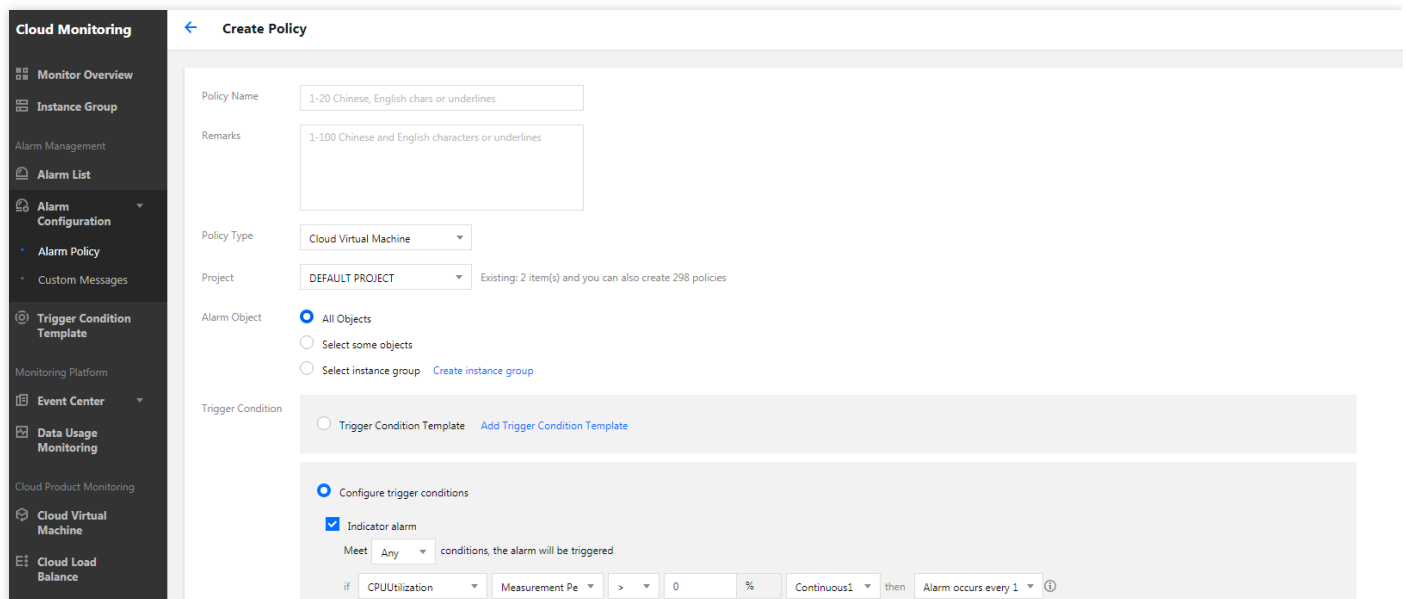
When disk utilization of a core node exceeds 90%, NodeManager will set it as "Unhealthy".

### Solution

1. We recommend that you use Cloud Monitor and set the alarm threshold for ERM CVM disk utilization to 80-85%, so that the node disk utilization won't exceed 90%.

Configure the EMR disk utilization threshold in Cloud Monitor at the following address:

```
https://console.tencentcloud.com/monitor/policyTemplate
```



2. If the disk capacity is insufficient, you can add core nodes and perform load balancing to reduce the load of HDFS storage on the core nodes.
3. Clean up the disk regularly, especially the following parts:
  - Storage space of the core nodes.
  - Entire storage space of HDFS.

# HiveServer2 migrates to Router

Last updated : 2020-12-04 17:50:36

## How do I migrate HiveServer2 to a router node?

1. Log in to the [EMR console](#), and click the specific cluster ID in the **ID/Name** column on the **Cluster List** page to go to the cluster details page. Click **Cluster Resource > Resource Management** to go to the resource management page. Click **Scale Out** to go to the cluster scale-out page.

**Resource Management** Help Documentation

The current cluster resources were purchased based on the official EMR billing rules and the renewal statuses of CVM resources are managed by EMR in a unified manner. For monthly-subscription CVM resources, the billing statuses shown on the current page shall prevail. It is not allowed to adjust the fee policy on the CVM Console.

All Nodes
Master
Core
Common
Task
Router
Metadb
Recyle
Renew

Scale Out
Scale In
More 

Search IP

<input type="checkbox"/>	Node Type	Resou...	Resource ...	IP	Configura...	Deployment Process	Expiry Time	Operation
<input type="checkbox"/>	Master	HOST			EMR StandardSA 2 CPU: 4-core; memory: 8GB Efficient cloud disk : 100G x 1	knox,JobHistoryServer,NameNode,Zookeeper,ResourceManager	Pay-as-you-go	<a href="#">Change Configuration</a>
<input type="checkbox"/>	Core	HOST			EMR StandardSA 2 CPU: 4-core; memory: 8GB Efficient cloud disk : 100G x 1	DataNode,NodeManager	Pay-as-you-go	<a href="#">Change Configuration</a>

On the cluster scale-out page, select **Router** in **Node Type** and **Hive-2.3.5** in **Scale-out Service**. Configure other

settings as required.

### Cluster Scale-out ✕

Node Type Core Task Router

Billing Mode Pay-as-you-go


Services HDFS-2.8.5 YARN-2.8.5

Deployment Process  Show deployment process

Current Specification **EMR StandardSA2 / 4-core 8 GB / Efficient cloud disk 100G \* 1**  
The default node specification for scale-out. You can adjust the specification in [Node Specification](#).

Scale-out Quantity - 1 +

Tag + New Tag  
Up to 5 tags can be bound

Cost 

Confirm Cancel

2. Log in to the router node and modify the `hive-site.xml` configuration file.

```
1. vim /usr/local/service/hive/conf/hive-site.xml
2. <property>
3.   <name>hive.metastore.uris</name>
4.   <value>thrift://<router-ip>:7004</value>
5. </property>
```

3. Disable the Hive service on the master node.

On the **Cluster Service** page, select **Operation > Role Management** for the Hive component, pause all the Hive

processes on the master node, and restart the Hive processes on the router node.

Cluster Services / ZOOKEEPER Help Documentation

Service Status **Role Management** Configuration Management Configuration Record

Restart Service Enter Maintenance Exit Maintenance   Nodes pending restart

Start Pause

<input type="checkbox"/>	Roles	Role Status	Configura...	Node Type	Maintenance	Node IP	Last Restarted
<input type="checkbox"/>	Zookeeper	Running	zookeeper-defa...	Master	Normal mode		

Total 1 item Lines per page 20 1 /1 page

#### 4. Conduct a test.

On the router node, check whether HiveServer2 can be properly connected to and existing tables can be queried; if yes, the migration is successful.

1. `cd /usr/local/service/hive`
2. `./bin/beeLine -u "jdbc:hive2://<router-ip>:7001" -n hadoop -p hadoop`
3. `show tables;`

#### 5. Modify the Hue configuration file to route requests to the Hive component on the router node.

```
vim /usr/local/service/knox/conf/topologies/emr.xml Modify `HIVE` and `HIVEUI`.
<service>
<role>HIVE</role>
<url>http://Router-ip:7003</url>
<param>
<name>replayBufferSize</name>
<value>8</value>
</param>
</service>
<service>
<role>HIVEUI</role>
<url>http://Router-ip:7003</url>
</service>
```

Run the following commands to restart Knox:

```
su hadoop
/usr/local/service/knox/bin/gateway.sh stop ; /usr/local/service/knox/bin/gatew
ay.sh start
```





# InternalError in the Console

Last updated : 2020-05-15 17:26:14

## How do I fix the "InternalError" reported by the console?

1. "InternalError" is reported when a non-root account attempts to purchase an EMR cluster.

**Cause:** the current account lacks permission.

**Solution:** confirm whether the current account has completed identity verification and been granted the payment permission.

2. "InternalError" is reported when a non-root account clicks **Hardware Management** in the console.

**Cause:** the current account lacks permission.

**Solution:** open the following link:

```
https://console.tencentcloud.com/cam/role/grant?roleName=EMR_QCSRole&policyName=QcloudAccessForEMRRole&principal=eyJzZXJ2aWNlIjoizW1yLmNsb3VkLnRlbnNlbnQuY29tIn0=&serviceType=EMR . Then, grant the EMR permission by using the root account.
```

# EMR on CVM

## Billing

Last updated : 2023-07-14 10:57:58

### **Are there any additional charges for adding components after a cluster is created?**

EMR charges you fees based on the nodes used in the cluster deployment, with no additional charge for adding components.

### **Why the EMR on CVM console shows node auto-renewal, but the associated CVM instance in the CVM console does not?**

EMR cluster nodes are composed of CVM instances, and their billing type and auto-renewal are managed in the EMR console.

### **Does renewing an EMR on CVM node renew the associated cloud disk (CBS)?**

A system disk has a renewal period same as the nodes and does not need to be renewed separately. The renewal of a cloud data disk depends on how the node renewal is set. If nodes are renewed manually or automatically in the EMR console, the associated cloud data disks do not need to be renewed separately; if they are renewed through the Billing Center, the cloud data disks need to be renewed separately.

# Cluster Management

Last updated : 2023-07-14 10:57:58

## Can I upgrade a non-HA cluster to an HA one?

A non-HA cluster and an HA cluster differ greatly in their deployment architecture, so they are not interchangeable. You can purchase a new cluster with an appropriate deployment architecture as needed. Please note that a non-HA cluster can be used only for test but not for production.

## What components and versions are available?

EMR releases new versions on a regular basis. For components and versions available, see [Product Releases and Component Versions](#).

## Can I upgrade components after creating a cluster?

This is not supported.

## Can I add components after creating a cluster?

Yes. You can add components that are supported by your current cluster version but are not deployed.

## Can I change the project after creating a cluster?

No. The project cannot be changed after a cluster is created. We recommend you use tagging to realize authentication, cost allocation, and other relevant features.

## Can I set different model specs for nodes of the same type?

Nodes of the same type are deployed with the same services. We recommend you set an identical model spec for nodes of the same type to facilitate Ops and management. If your desired model spec is sold out or the existing spec does not meet your needs, you can change the specs of added nodes.

## Can I change existing nodes to lower configurations?

In the EMR on CVM scenario, various services are deployed on nodes of various types and share the compute, memory, and disk resources of the same node. Changing a node to lower configurations may cause the resource usage in the actual service configuration file exceeding the physical resource volume of the node, which will make the cluster unavailable. Therefore, existing nodes cannot be changed to lower configurations. If existing configurations of core, task, and router nodes are too high, you can add nodes with lower configurations first, and then remove those with higher configurations to meet your needs.

## Can I remove core nodes?

Core nodes store data and cannot be removed by default. If you have such a need, [submit a ticket](#).

# Reasons Why a Triggered Auto Scaling Rule Is Not Executed

Last updated : 2023-05-30 11:24:00

1. The number of elastic nodes reaches the lower limit. To continue scale-in, you can adjust the minimum number of nodes allowed.

Cause: The scale-in rule is triggered, but the current number of elastic nodes is smaller than the minimum number of nodes allowed.

Solution: To continue scale-in, reset the minimum number of nodes allowed.
2. The number of elastic nodes reaches the upper limit. To continue scale-out, you can adjust the maximum number of nodes allowed.

Cause: The scale-out rule is triggered, but the current number of elastic nodes is equal to the maximum number of nodes allowed.

Solution: To continue scale-out, reset the maximum number of nodes allowed.
3. Scale-out is impossible because no scaling specification is configured. You can add a scaling specification and try again.

Cause: The auto scaling rule is triggered, but no node specification is added in the **Auto-scaling > Scaling Specification Management** section in the console, as shown in the figure above.

Solution: Click **Add Specification** in the top-right corner and select the required node specification.
4. Scaling cannot be performed due to insufficient resources. You can switch to a specification with sufficient resources or [submit a ticket](#) to contact us.

Cause: The scale-out rule is triggered but the selected specification has insufficient resources in the current availability zone.

Solution: Switch to a node specification with sufficient resources.
5. The current retry period upon expiration is too short. We recommend extending the retry period.

Cause: The scaling rule is triggered, but there are other auto scaling processes in the cluster, which prevent the rule from being executed within current retry period upon expiration.

Solution: Edit the rule to extend the retry period upon expiration to ensure that the rule can be executed.
6. Scale-out cannot be performed due to insufficient account balance.

Cause: The scale-out rule is triggered, but the account balance is insufficient at the time when an order is placed.

Solution: Go to the [Billing Center](#) to top up your account.

7. No eligible elastic resources are available for scale-in.

Cause: The scale-in rule is triggered, but no elastic nodes are available currently or timed termination is set for all nodes.

Solution: Perform manual scale-in for nodes that are set to be terminated at a certain time.

8. The cluster is not in a scalable state.

Cause: The scale-out rule is triggered, but the cluster is not running due to component installation, previously triggered scale-out, or other factors.

Solution: Perform manual scale-out or edit the rule to extend the retry period upon expiration to ensure that the rule can be executed.

9. The cluster is in a cooldown period and scale-out cannot be triggered for the moment. We recommend adjusting the cooldown period for scaling rules.

Cause: The scale-out rule is triggered, but the cluster is in a cooldown period due to other scaling processes.

Solution: Shorten the cooldown period for other scaling rules or extend the retry period upon expiration for the current scale-out rule.