

Alibaba Cloud E-MapReduce

Errors and Troubleshooting

Issue: 20190905

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 <i>Instance_ID</i></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 Appendix.....	1
1.1 Cluster management page.....	1
1.2 Configure cluster ports.....	2
1.3 Notes on E-MapReduce versions.....	4
1.4 Error code list.....	4
1.5 Status list.....	6

1 Appendix

1.1 Cluster management page

You can purchase a one vCPU 2 GB ECS instance that runs the Ubuntu system and deploy the instance in a VPC network. You can use this instance as a management client to access the management pages.

The following table lists the endpoints of services in the cluster.

Software	Service	Endpoint
Hadoop		
	yarn resourcemanager	k masternode1_private_ip:8088,masternode2_private_ip:8088
	jobhistory	masternode1_private_ip:19888
	timeline server	masternode1_private_ip:8188
	hdfs	masternode1_private_ip:50070,masternode2_private_ip:50070
Spark		
	spark ui	masternode1_private_ip:4040
	history	masternode1_private_ip:18080
Tez		
	tez-ui	masternode1_private_ip:8090/tez-ui2
Hue		
	hue	masternode1_private_ip:8888
Zeppelin		

Software	Service	Endpoint
	zeppelin	masternode1_private_ip: 8080
Hbase		
	hbase	masternode1_private_ip: 16010
Presto		
	presto	masternode1_private_ip: 9090
Oozie		
	oozie	masternode1_private_ip: 11000
Ganglia		
	ganglia	masternode1_private_ip: 8085/ganglia

1.2 Configure cluster ports

Hadoop HDFS

Service	Port number	Parameter	Description
NameNode	9000	fs.default.name or fs.defaultFS	fs.default.name is expired but still available
NameNode	50070	dfs.http.address or dfs.namenode.http- address	dfs.http.address is expired but still available.

Hadoop YARN (MRv2)

Service	Port number	Parameter
JobHistory Server	10020	mapreduce.jobhistory. address
JobHistory Server	19888	mapreduce.jobhistory. webapp.address
ResourceManager	8025	yarn.resourcemanager. resource-tracker.address

Service	Port number	Parameter
ResourceManager	8032	yarn.resourcemanager.address
ResourceManager	8030	yarn.resourcemanager.scheduler.address
ResourceManager	8088	yarn.resourcemanager.webapp.address

Hadoop MapReduce (MRv1)

Service	Port number	Parameter
JobTracker	8021	mapreduce.jobtracker.address

Hadoop HBase

Service	Port number	Parameter
HMaster	16000	hbase.master.port
HMaster	16010	hbase.master.info.port
HRegionServer	16020	hbase.regionserver.port
HRegionServer	16030	hbase.regionserver.info.port
ThriftServer	9099	-

Hadoop Spark

Service	Port number
SparkHistory	18080

Storm

Service	Port number	Parameter
Storm UI	9999	ui_port

Druid

Service	Port number	Parameter
overlord	18090	overlord.runtime -> druid.port

Service	Port number	Parameter
coordinator	18081	coordinator.runtime -> druid.port
middleManager	18091	middleManager.runtime -> druid.port
historical	18083	historical.runtime -> druid.port

1.3 Notes on E-MapReduce versions

- EMR is regularly updated.
- Versions of software installed within each EMR version are fixed. Currently, EMR does not support choosing different versions of software. We recommend that you do not manually change the versions of software. For example, Hadoop 3.6.0 and Spark 1.4.1 are installed in EMR V1.0.
- If you have selected a version and created a cluster, the version that the cluster uses will not be automatically updated. If you select V1.0, the Hadoop remains at V2.6.0 and Spark remains at V1.4.1. When EMR is updated to V1.1, Hadoop is updated to V2.7.0 and Spark is updated to V1.5.0. These updates do no effect the clusters that you have created. Only new clusters use new mirror images.
- When updating the cluster version, for example, from V1.0 to V1.1, test the jobs in the news software environment, to see if they run successfully and to avoid exceptions caused by incompatibility and a change of software environment.

1.4 Error code list

Common error codes

Error code	Error message
4001	The error message returned when the request parameter is invalid, for example , the parameter is missing or the format of the parameter is invalid.
4005	The error message returned when you are not authorized to access resources of other users.

Error code	Error message
4006	The error message returned when the cluster is in the Abnormal status and the job cannot be submitted. Check whether the cluster associated with the execution plan has been released.
4007	The error message returned when the name of the security group is empty.
4009	The error message returned when your account has an overdue payment or is suspended. Check the status of your account.
4011	The error message returned when the cluster is in the Abnormal status and cannot be scheduled. Check whether the cluster associated with the execution plan has been released.
5012	The error message returned when the number of security groups you can create has exceeded the upper limit. Go to the security group page and delete security groups that are not in use.
5038	The error message returned when the job is in a running or pending execution plan and cannot be modified. You can modify the job only after the associated execution plan has been successfully executed. You can clone the job, then modify and use the cloned job.
5039	The error message returned when you fail to lock the cluster role. You must have certain permissions to use EMR. For role authorization , click here to create cluster roles.
5050	The error message returned when accessing the database. Try again later.
6002	The error message returned when cluster updating failure occurs.

Error code	Error message
8002	The error message returned when you are not authorized to perform the specified operation. Click RAM to apply for authorization.
8003	The error message returned when you are not authorized to perform the PassRole operation. Click RAM to apply for authorization.
9006	The error message returned when the ID of the cluster does not exist. You need to verify the ID.
9007	The error message returned when the password used to log on to the master node is invalid. The password must be 8-30 characters in length and can contain uppercase letters, lowercase letters, and numbers.

ECS-related errors

Error message	Description
The specified InstanceType is not authorized for use.	You have not applied for the specified types of instances that are used to create clusters. You can apply for high-configuration instances on the ECS buy page.
No zone or cluster resources are available .	No ECS resources are available in this zone.

1.5 Status list

Cluster status list



Note:

You can view the cluster status in the cluster list or on the cluster details page.

Status	Status code	Description
Creating	CREATING	The cluster is being created. The creation task includes two stages : creating physical ECS machines and activating Spark clusters. It takes a moment for the clusters to start running.
Failed	CREATE_FAILED	An exception occurred during creation. The ECS instance that you have created automatically rolls back. You can click the question mark (?) to the right of the status on the cluster list page to view exception details.
Running	RUNNING	The computing cluster is running.
Idle	IDLE	The cluster is not running any execution plan.
Releasing	RELEASING	Click Release in the status list to set the cluster to this status. This status indicates that the cluster is in the releasing process . It may take a moment to complete this process.
Release Failed	RELEASE_FAILED	An exception occurred when releasing the cluster . You can click the question mark (?) to the right of the status on the cluster list page to view the exception details. When the cluster is in this status, click Release to release the cluster again .

Status	Status code	Description
Released	RELEASED	The computing cluster and the ECS instance that hosts the computing cluster have been released.
Abnormal	ABNORMAL	Unrecoverable errors occurred on one or more nodes in the computing cluster. Click Release to release the cluster.

Job status list



Note:

View job status in the job status list

Status	Description
Ready	The creation information is complete, correct, and successfully saved. The job is ready to be added to the submission queue. It may take a moment for the job to change its status to Submitting.
Submitting	The job is in the submission queue of the computing cluster. It has not been submitted to the computing cluster.
Failed	An exception occurred when submitting the job to the computing cluster. You need to clone and submit the job if you want to submit this job again.
Running	The job is running in the cluster. Wait a moment and click the corresponding log button in the job list to view output log entries in real time.
Succeeded	The job has been successfully executed in the cluster. Click the corresponding log button to view the log entry.
Failed	An exception occurred when executing the job. Click the corresponding log button to view the log entry.