

Hortonworks Data Platform

Ambari Troubleshooting Guide

(Apr 13, 2015)

Hortonworks Data Platform : Ambari Troubleshooting Guide

Copyright © 2012-2014 Hortonworks, Inc. All rights reserved.

The Hortonworks Data Platform, powered by Apache Hadoop, is a massively scalable and 100% open source platform for storing, processing and analyzing large volumes of data. It is designed to deal with data from many sources and formats in a very quick, easy and cost-effective manner. The Hortonworks Data Platform consists of the essential set of Apache Hadoop projects including MapReduce, Hadoop Distributed File System (HDFS), HCatalog, Pig, Hive, HBase, Zookeeper and Ambari. Hortonworks is the major contributor of code and patches to many of these projects. These projects have been integrated and tested as part of the Hortonworks Data Platform release process and installation and configuration tools have also been included.

Unlike other providers of platforms built using Apache Hadoop, Hortonworks contributes 100% of our code back to the Apache Software Foundation. The Hortonworks Data Platform is Apache-licensed and completely open source. We sell only expert technical support, [training](#) and partner-enablement services. All of our technology is, and will remain free and open source. Please visit the [Hortonworks Data Platform](#) page for more information on Hortonworks technology. For more information on Hortonworks services, please visit either the [Support](#) or [Training](#) page. Feel free to [Contact Us](#) directly to discuss your specific needs.

Licensed 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.

Table of Contents

| | |
|---|---|
| 1. Troubleshooting Ambari Deployments | 1 |
| 1.1. Troubleshooting Ambari Issues | 1 |
| 1.2. Reviewing Ambari Log Files | 1 |
| 1.3. Resolving Ambari Installer Problems | 2 |
| 1.3.1. Problem: Browser crashed before Install Wizard completes | 2 |
| 1.3.2. Problem: Install Wizard reports that the cluster install has failed | 2 |
| 1.3.3. Problem: Ambari Agents May Fail to Register with Ambari Server. | 3 |
| 1.3.4. Problem: The "yum install ambari-server" Command Fails | 4 |
| 1.3.5. Problem: HDFS Smoke Test Fails | 4 |
| 1.3.6. Problem: yum Fails on Free Disk Space Check | 4 |
| 1.3.7. Problem: A service with a customized service user is not appearing properly in Ambari Web | 5 |
| 1.4. Resolving Cluster Deployment Problems | 5 |
| 1.4.1. Problem: Trouble Starting Ambari on System Reboot | 5 |
| 1.4.2. Problem: Metrics and Host information display incorrectly in Ambari Web | 5 |
| 1.4.3. Problem: On SUSE 11 Ambari Agent crashes within the first 24 hours | 5 |
| 1.4.4. Problem: Attempting to Start HBase REST server causes either REST server or Ambari Web to fail | 6 |
| 1.4.5. Problem: Multiple Ambari Agent processes are running, causing re- register | 6 |
| 1.4.6. Problem: Some graphs do not show a complete hour of data until the cluster has been running for an hour | 6 |
| 1.4.7. Problem: Ambari stops MySQL database during deployment, causing Ambari Server to crash. | 6 |
| 1.4.8. Problem: Cluster Install Fails with Groupmod Error | 7 |
| 1.4.9. Problem: Host registration fails during Agent bootstrap on SLES due to timeout. | 7 |
| 1.4.10. Problem: Host Check Fails if Transparent Huge Pages (THP) is not disabled. | 7 |
| 1.5. Resolving General Problems | 8 |
| 1.5.1. During Enable Kerberos, the Check Kerberos operation fails. | 8 |
| 1.5.2. Problem: Hive developers may encounter an exception error message during Hive Service Check | 8 |
| 1.5.3. Problem: API calls for PUT, POST, DELETE respond with a "400 - Bad Request" | 9 |

1. Troubleshooting Ambari Deployments

The first step in troubleshooting any problem in an Ambari-deploying Hadoop cluster is [Reviewing the Ambari Log Files](#).

Find a recommended solution to a troubleshooting problem in one of the following sections:

- [Resolving Ambari Installer Problems](#)
- [Resolving Cluster Deployment Problems](#)
- [Resolving General Problems](#)

1.1. Troubleshooting Ambari Issues

The first step in troubleshooting any problem in an Ambari-deploying Hadoop cluster is [Reviewing the Ambari Log Files](#).

Find a recommended solution to a troubleshooting problem in one of the following sections:

- [Resolving Ambari Installer Problems](#)
- [Resolving Cluster Deployment Problems](#)
- [Resolving General Problems](#)

1.2. Reviewing Ambari Log Files

Find files that log activity on an Ambari host in the following locations:

- Ambari Server logs

```
<your.Ambari.server.host>/var/log/ambari-server/ambari-server.log
```

- Ambari Agent logs

```
<your.Ambari.agent.host>/var/log/ambari-agent/ambari-agent.log
```

- Ambari Action logs

```
<your.Ambari.agent.host>/var/lib/ambari-agent/data/
```

This location contains logs for all tasks executed on an Ambari agent host. Each log name includes:

- command-N.json - the command file corresponding to a specific task.
- output-N.txt - the output from the command execution.

- errors-N.txt - error messages.

1.3. Resolving Ambari Installer Problems

Try the recommended solution for each of the following problems:

1.3.1. Problem: Browser crashed before Install Wizard completes

Your browser crashes or you accidentally close your browser before the Install Wizard completes.

1.3.1.1. Solution

The response to a browser closure depends on where you are in the process:

- The browser closes before you press the `Deploy` button.

Re-launch the same browser and continue the install process. Using a different browser forces you to re-start the entire process.

- The browser closes after you press `Deploy`, while or after the `Install`, `Start`, and `Test` screen opens.

Re-launch the same browser and continue the process, or log in again, using a different browser. When the `Install`, `Start`, and `Test` displays, proceed.

1.3.2. Problem: Install Wizard reports that the cluster install has failed

The `Install`, `Start`, and `Test` screen reports that the cluster install has failed.

1.3.2.1. Solution

The response to a report of install failure depends on the cause of the failure:

- The failure is due to intermittent network connection errors during software package installs.

Use the `Retry` button on the `Install`, `Start`, and `Test` screen.

- The failure is due to misconfiguration or other setup errors.

1. Use the left navigation bar to go back to the appropriate screen. For example, `Customize Services`.

2. Make your changes.

3. Continue in the normal way.

- The failure occurs during the start/test sequence.
 1. Click `Next` and `Complete`, then proceed to the `Monitoring Dashboard`.
 2. Use the `Services View` to make your changes.
 3. Re-start the service using `Service Actions`.
- The failure is due to something else.
 1. Open an SSH connection to the Ambari Server host.
 2. Clear the database. At the command line, type:

```
ambari-server reset
```
 3. Clear your browser cache.
 4. Re-run the Install Wizard.

1.3.3. Problem: Ambari Agents May Fail to Register with Ambari Server.

When deploying HDP using Ambari 1.4.x or later on RHEL CentOS 6.5, click the “Failed” link on the Confirm Hosts page in the Cluster Install wizard to display the Agent logs. The following log entry indicates the SSL connection between the Agent and Server failed during registration:

```
INFO 2014-04-02 04:25:22,669 NetUtil.py:55 - Failed to
connect to https://{ambari-server}:8440/cert/ca due to
[Errno 1] _ssl.c:492: error:100AE081:elliptic curve
routines:EC_GROUP_new_by_curve_name:unknown group
```

For more detailed information about this OpenSSL issue, see https://bugzilla.redhat.com/show_bug.cgi?id=1025598

1.3.3.1. Solution:

In certain recent Linux distributions, such as RHEL/Centos/Oracle Linux 6.x, the default value of `nproc` is lower than the value required to deploy the HBase service successfully. If you are deploying HBase, change the value of `nproc`:

1. Check the OpenSSL library version installed on your host(s):

```
rpm -qa | grep openssl openssl-1.0.1e-15.el6.x86_64
```

2. If the output reads `openssl-1.0.1e-15.x86_64 (1.0.1 build 15)`, you must upgrade the OpenSSL library. To upgrade the OpenSSL library, run the following command:

```
yum upgrade openssl
```

3. Verify you have the newer version of OpenSSL (1.0.1 build 16):

```
rpm -qa | grep openssl-1.0.1e-16.el6.x86_64
```

4. Restart Ambari Agent(s) and click Retry -> Failed in the wizard user interface.

1.3.4. Problem: The “yum install ambari-server” Command Fails

You are unable to get the initial install command to run.

1.3.4.1. Solution:

You may have incompatible versions of some software components in your environment. See Meet Minimum System Requirements in Installing HDP Using Ambari for more information, then make any necessary changes.

1.3.5. Problem: HDFS Smoke Test Fails

If your DataNodes are incorrectly configured, the smoke tests fail and you get this error message in the DataNode logs:

```
DisallowedDataNodeException  
org.apache.hadoop.hdfs.server.protocol.  
DisallowedDatanodeException
```

1.3.5.1. Solution:

1. Make sure that reverse DNS look-up is properly configured for all nodes in your cluster.
2. Make sure you have the correct FQDNs when specifying the hosts for your cluster. Do not use IP addresses - they are not supported.
3. Restart the installation process.

1.3.6. Problem: yum Fails on Free Disk Space Check

If you boot your Hadoop DataNodes with/as a ramdisk, you must disable the free space check for yum before doing the install. If you do not disable the free space check, yum will fail with the following error:

```
Fail: Execution of '/usr/bin/yum -d 0 -e 0 -y install unzip'  
returned 1. Error Downloading Packages: unzip-6.0-1.el6.x86_64:  
Insufficient space in download directory /var/cache/yum/x86_64/6/  
base/packages * free 0 * needed 149 k
```

1.3.6.1. Solution:

To disable free space check, update the DataNode image with a directive in `/etc/yum.conf`:

```
diskspacecheck=0
```

1.3.7. Problem: A service with a customized service user is not appearing properly in Ambari Web

You are unable to monitor or manage a service in Ambari Web when you have created a customized service user name with a hyphen, for example, `hdfs-user`.

1.3.7.1. Solution

Hyphenated service user names are not supported. You must re-run the Ambari Install Wizard and create a different name.

1.4. Resolving Cluster Deployment Problems

Try the recommended solution for each of the following problems:

1.4.1. Problem: Trouble Starting Ambari on System Reboot

If you reboot your cluster, you must restart the Ambari Server and all the Ambari Agents manually.

1.4.1.1. Solution:

Log in to each machine in your cluster separately:

1. On the Ambari Server host machine:

```
ambari-server start
```

2. On each host in your cluster:

```
ambari-agent start
```

1.4.2. Problem: Metrics and Host information display incorrectly in Ambari Web

Charts appear incorrectly or not at all despite Host health status is displayed incorrectly.

1.4.2.1. Solution:

All the hosts in your cluster and the machine from which you browse to Ambari Web must be in sync with each other. The easiest way to assure this is to enable NTP.

1.4.3. Problem: On SUSE 11 Ambari Agent crashes within the first 24 hours

SUSE 11 ships with Python version 2.6.0-8.12.2 which contains a known defect that causes this crash.

1.4.3.1. Solution:

Upgrade to Python version 2.6.8-0.15.1.

1.4.4. Problem: Attempting to Start HBase REST server causes either REST server or Ambari Web to fail

As an option you can start the HBase REST server manually after the install process is complete. It can be started on any host that has the HBase Master or the Region Server installed. If you install the REST server on the same host as the Ambari server, the http ports will conflict.

1.4.4.1. Solution

In starting the REST server, use the `-p` option to set a custom port. Use the following command to start the REST server. `/usr/lib/hbase/bin/hbase-daemon.sh start rest -p <custom_port_number>`

1.4.5. Problem: Multiple Ambari Agent processes are running, causing re-register

On a cluster host `ps aux | grep ambari-agent` shows more than one agent process running. This causes Ambari Server to get incorrect ids from the host and forces Agent to restart and re-register.

1.4.5.1. Solution

On the affected host, kill the processes and restart.

1. Kill the Agent processes and remove the Agent PID files found here: `/var/run/ambari-agent/ambari-agent.pid`.
2. Restart the Agent process:

```
ambari-agent start
```

1.4.6. Problem: Some graphs do not show a complete hour of data until the cluster has been running for an hour

When you start a cluster for the first time, some graphs, such as `Services View > HDFS` and `Services View > MapReduce`, do not plot a complete hour of data. Instead, they show data only for the length of time the service has been running. Other graphs display the run of a complete hour.

1.4.6.1. Solution

Let the cluster run. After an hour all graphs will show a complete hour of data.

1.4.7. Problem: Ambari stops MySQL database during deployment, causing Ambari Server to crash.

The Hive Service uses MySQL Server by default. If you choose MySQL server as the database on the Ambari Server host as the managed server for Hive, Ambari stops this database during deployment and crashes.

1.4.7.1. Solution

If you plan to use the default MySQL Server setup for Hive and use MySQL Server for Ambari - make sure that the two MySQL Server instances are different.

If you plan to use the same MySQL Server for Hive and Ambari - make sure to choose the existing database option for Hive.

1.4.8. Problem: Cluster Install Fails with Groupmod Error

The cluster fails to install with an error related to running `groupmod`. This can occur in environments where groups are managed in LDAP, and not on local Linux machines. You may see an error message similar to the following one:

```
Fail: Execution of 'groupmod hadoop' returned 10. groupmod: group 'hadoop' does not exist in /etc/group
```

1.4.8.1. Solution

When installing the cluster using the Cluster Installer Wizard, at the `Customize Services` step, select the `Misc` tab and choose the `Skip group modifications during install` option.

1.4.9. Problem: Host registration fails during Agent bootstrap on SLES due to timeout.

When using SLES and performing host registration using SSH, the Agent bootstrap may fail due to timeout when running the `setupAgent.py` script. The host on which the timeout occurs will show the following process hanging:

```
c6401.ambari.apache.org:/etc/ # ps -ef | grep zypper root 18318
18317 5 03:15 pts/1 00:00:00 zypper -q search -s --match-exact
ambari-agent
```

1.4.9.1. Solution

1. If you have a repository registered that is prompting to accept keys, via user interaction, you may see the hang and timeout. In this case, run `zypper refresh` and confirm all repository keys are accepted for the `zypper` command to work without user interaction.
2. Another alternative is to perform manual Agent setup and not use SSH for host registration. This option does not require that Ambari call `zypper` without user interaction.

1.4.10. Problem: Host Check Fails if Transparent Huge Pages (THP) is not disabled.

When installing Ambari on RHEL/CentOS 6 using the Cluster Installer Wizard at the `Host Checks` step, one or more host checks may fail if you have not disabled Transparent Huge Pages on all hosts.

Host Checks will warn you when a failure occurs.

1.4.10.1. Solution

Disable THP. On all hosts,

1. Add the following command to your `/etc/rc.local` file:

```
if test -f /sys/kernel/mm/transparent_hugepage/enabled; then
echo never > /sys/kernel/mm/redhat_transparent_hugepage/enabled
fi if test -f /sys/kernel/mm/transparent_hugepage/defrag; then
echo never > /sys/kernel/mm/redhat_transparent_hugepage/defrag
fi
```

2. To confirm, reboot the host then run the following command:

```
$ cat /sys/kernel/mm/transparent_hugepage/enabled always madvise
[never]
```

1.5. Resolving General Problems

1.5.1. During Enable Kerberos, the Check Kerberos operation fails.

When enabling Kerberos using the wizard, the Check Kerberos operation fails. In `/var/log/ambari-server/ambari-server.log`, you see a message: `02:45:44,490 WARN [qtp567239306-238] MITKerberosOperationHandler:384 - Failed to execute kadmin.`

1.5.1.1. Solution 1:

Check that NTP is running and confirm your hosts and the KDC times are in sync. A time skew as little as 5 minutes can cause Kerberos authentication to fail.

1.5.1.2. Solution 2: (on RHEL/CentOS/Oracle Linux)

Check that the Kerberos Admin principal being used has the necessary KDC ACL rights as set in `/var/kerberos/krb5kdc/kadm5.acl`.

1.5.2. Problem: Hive developers may encounter an exception error message during Hive Service Check

MySQL is the default database used by the Hive metastore. Depending on several factors, such as the version and configuration of MySQL, a Hive developer may see an exception message similar to the following one:

```
An exception was thrown while adding/validating classes) :
Specified key was too long; max key length is 767 bytes
```

1.5.2.1. Solution

Administrators can resolve this issue by altering the Hive metastore database to use the Latin1 character set, as shown in the following example: `mysql> ALTER DATABASE <metastore.database.name> character set latin1;`

1.5.3. Problem: API calls for PUT, POST, DELETE respond with a "400 - Bad Request"

When attempting to perform a REST API call, you receive a 400 error response. REST API calls require the "X-Requested-By" header.

1.5.3.1. Solution

Starting with Ambari 1.4.2, you must include the "X-Requested-By" header with the REST API calls.

For example, if using curl, include the `-H "X-Requested-By: ambari"` option.

```
curl -u admin:admin -H "X-Requested-By: ambari" -X DELETE http://  
<ambari-host>:8080/api/v1/hosts/host1
```