

Hortonworks Data Platform

Installing Hadoop Using Apache Ambari

(Aug 1, 2014)

Hortonworks Data Platform : Installing Hadoop Using Apache Ambari

Copyright © 2012, 2014 Hortonworks, Inc. Some 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.



Except where otherwise noted, this document is licensed under
Creative Commons Attribution ShareAlike 3.0 License.
<http://creativecommons.org/licenses/by-sa/3.0/legalcode>

Table of Contents

I. Setting Up Ambari	i
1. Getting Ready	2
1.1. Determine Version Compatibility	2
1.2. Understand the Basics	3
1.3. Meet Minimum System Requirements	4
1.3.1. Hardware Recommendations	4
1.3.2. Operating Systems Requirements	4
1.3.3. Browser Requirements	5
1.3.4. Software Requirements	5
1.3.5. JDK Requirements	6
1.3.6. Database Requirements	6
1.3.7. File System Partitioning Recommendations	6
1.4. Collect Information	6
1.5. Prepare the Environment	7
1.5.1. Check Existing Installs	7
1.5.2. Set Up Password-less SSH	8
1.5.3. Set up Users and Groups	9
1.5.4. Enable NTP on the Cluster and on the Browser Host	9
1.5.5. Check DNS	9
1.5.6. Configuring iptables	10
1.5.7. Disable SELinux and PackageKit and check the umask Value	11
1.6. Optional: Configure Ambari for Local Repositories	11
2. Running the Ambari Installer	13
2.1. Set Up the Bits	13
2.1.1. RHEL/CentOS/Oracle Linux 5.x	13
2.1.2. RHEL/CentOS/Oracle Linux 6.x	14
2.1.3. SLES 11	14
2.2. Set Up the Server	14
2.2.1. Setup Options	15
2.3. Optional: Set Up LDAP or Active Directory Authentication	16
2.4. Optional: Set Up Security for Ambari	19
2.4.1. Set Up HTTPS for Ambari Server	19
2.4.2. Set Up HTTPS for Ganglia	21
2.4.3. Set Up HTTPS for Nagios	22
2.4.4. Optional: Encrypt Database and LDAP Passwords	23
2.5. Optional: Set Up Two-Way SSL Between Ambari Server and Ambari Agents	25
2.6. Optional: Change the Ambari Server Port	26
2.7. Optional: Configure Ambari Server for Internet Proxy	26
2.8. Start the Ambari Server	26
II. Hadoop 2.x - Deploying, Configuring, and Upgrading Ambari	xxviii
3. Hadoop 2.x - Installing, Configuring, and Deploying the Cluster	29
3.1. Log into Apache Ambari	29
3.2. Welcome	29
3.3. Select Stack	29
3.4. Install Options	30
3.5. Confirm Hosts	31
3.6. Choose Services	31

3.7. Assign Masters	31
3.8. Assign Slaves and Clients	32
3.9. Customize Services	32
3.9.1. Service Users and Groups	33
3.9.2. Properties That Depend on Service Usernames/Groups	34
3.10. Review	34
3.11. Install, Start and Test	35
3.12. Summary	35
4. Hadoop 2.x - Troubleshooting Ambari Deployments	36
4.1. Review Ambari Log Files	36
4.2. Quick Checks	36
4.3. Specific Issues	37
4.3.1. Problem: Browser crashed before Install Wizard completed	37
4.3.2. Problem: Install Wizard reports that the cluster install has failed	37
4.3.3. Problem: "Unable to create new native thread" exceptions in HDFS DataNode logs or those of any system daemon	38
4.3.4. Problem: The "yum install ambari-server" Command Fails	39
4.3.5. Problem: HDFS Smoke Test Fails	39
4.3.6. Problem: The HCatalog Daemon Metastore Smoke Test Fails	39
4.3.7. Problem: MySQL and Nagios fail to install on RightScale CentOS 5 images on EC2	40
4.3.8. Problem: Trouble starting Ambari on system reboot	40
4.3.9. Problem: Metrics and Host information display incorrectly in Ambari Web	41
4.3.10. Problem: On SUSE 11 Ambari Agent crashes within the first 24 hours	41
4.3.11. Problem: Attempting to Start HBase REST server causes either REST server or Ambari Web to fail	41
4.3.12. Problem: Multiple Ambari Agent processes are running, causing re-register	41
4.3.13. Problem: Some graphs do not show a complete hour of data until the cluster has been running for an hour	42
4.3.14. Problem: After performing a cluster install the Nagios server is not started	42
4.3.15. Problem: A service with a customized service user is not appearing properly in Ambari Web	42
4.3.16. Problem: Updated configuration changes are not pushed to client/gateway nodes	43
5. Appendix: Upgrading Ambari Server to 1.4.4	44
6. Appendix: Upgrading the HDP Stack from 1.3.2 or later to 2.0.6	48
6.1. Preparing for the Upgrade	48
6.2. Setting Up the Ambari Repository	50
6.3. Upgrading to Ambari 1.4.4	51
6.4. Upgrading the Stack	52
6.4.1. Prepare for the Stack Upgrade	52
6.4.2. Upgrade the Stack	54
6.4.3. Add YARN/MR2 and Update Configurations	56
6.4.4. Complete the Stack Upgrade	58
7. Appendix: Hadoop 2.x - Configuring Ports	65
7.1. HDFS Ports	65

7.2. MapReduce Ports	66
7.3. YARN Ports	66
7.4. Hive Ports	66
7.5. HBase Ports	67
7.6. ZooKeeper Ports	68
7.7. WebHCat Port	68
7.8. Ganglia Ports	69
7.9. MySQL Port	69
7.10. Ambari Ports	69
7.11. Nagios Ports	70
8. Appendix: NameNode High Availability	71
8.1. Setting Up NameNode High Availability	71
8.1.1. Rolling Back NameNode HA	78
III. Hadoop 1.x - Deploying, Configuring, and Upgrading Ambari	lxxxvi
9. Hadoop 1.x - Installing, Configuring, and Deploying the Cluster	87
9.1. Log into Apache Ambari	87
9.2. Welcome	87
9.3. Select Stack	87
9.4. Install Options	88
9.5. Confirm Hosts	89
9.6. Choose Services	89
9.7. Assign Masters	89
9.8. Assign Slaves and Clients	90
9.9. Customize Services	90
9.9.1. Service Users and Groups	91
9.9.2. Properties That Depend on Service Usernames/Groups	92
9.9.3. Recommended Memory Configurations for the MapReduce Service	92
9.10. Review	93
9.11. Install, Start and Test	93
9.12. Summary	94
10. Troubleshooting Ambari Deployments	95
10.1. Review Ambari Log Files	95
10.2. Quick Checks	95
10.3. Specific Issues	96
10.3.1. Problem: Browser crashed before Install Wizard completed	96
10.3.2. Problem: Install Wizard reports that the cluster install has failed	96
10.3.3. Problem: "Unable to create new native thread" exceptions in HDFS DataNode logs or those of any system daemon	97
10.3.4. Problem: The "yum install ambari-server" Command Fails	98
10.3.5. Problem: HDFS Smoke Test Fails	98
10.3.6. Problem: The HCatalog Daemon Metastore Smoke Test Fails.....	98
10.3.7. Problem: MySQL and Nagios fail to install on RightScale CentOS 5 images on EC2	99
10.3.8. Problem: Trouble starting Ambari on system reboot	99
10.3.9. Problem: Metrics and Host information display incorrectly in Ambari Web	100
10.3.10. Problem: On SUSE 11 Ambari Agent crashes within the first 24 hours	100

10.3.11. Problem: Attempting to Start HBase REST server causes either REST server or Ambari Web to fail	100
10.3.12. Problem: Multiple Ambari Agent processes are running, causing re-register	100
10.3.13. Problem: Some graphs do not show a complete hour of data until the cluster has been running for an hour	101
10.3.14. Problem: After performing a cluster install the Nagios server is not started	101
10.3.15. Problem: A service with a customized service user is not appearing properly in Ambari Web	101
10.3.16. Problem: Updated configuration changes are not pushed to client/gateway nodes	102
11. Appendix: Upgrading Ambari Server to 1.4.4	103
12. Appendix: Upgrading Ambari Server to 1.2.5	106
13. Appendix: Upgrading the HDP Stack to 1.3.3	109
13.1. Preparing for the Upgrade	109
13.2. Setting Up the Ambari Repository	110
13.3. Upgrading Ambari	111
13.4. Upgrading the Stack (from 1.2.* to 1.3.3)	113
13.5. Upgrading the Stack (from 1.3.0 to 1.3.3)	117
14. Appendix: Configuring Ports	120
14.1. HDFS Ports	120
14.2. MapReduce Ports	121
14.3. Hive Ports	121
14.4. HBase Ports	122
14.5. ZooKeeper Ports	122
14.6. WebHCat Port	123
14.7. Ganglia Ports	123
14.8. MySQL Port	123
14.9. Ambari Ports	124
14.10. Nagios Ports	124
15. Configuring RHEL HA for Hadoop 1.x	125
15.1. Deploy the scripts	125
15.2. Configure Ambari properties across the HA cluster	125
15.3. Troubleshooting RHEL HA	126
IV. Additional Tasks with Ambari	cxxviii
16. Appendix: Installing Ambari Agents Manually	129
16.1. RHEL/CentOS/Oracle Linux 5.x and 6.x	129
16.2. SLES	129
17. Appendix: Using Custom Hostnames	130
18. Appendix: Upgrading Operating Systems on an Ambari-based Hadoop Installation	131
19. Appendix: Moving the Ambari Server	132
19.1. Back up Current Data	132
19.2. Update Agents	132
19.3. Install the New Server and Populate the Databases	133
20. Appendix: Using Non-Default Databases	135
20.1. Hive/HCatalog	135
20.1.1. Troubleshooting Hive/HCatalog	136
20.2. Oozie	138
20.2.1. Troubleshooting Oozie	139

20.3. Ambari	140
20.3.1. Troubleshooting Ambari	142
21. Setting Up Kerberos for Use with Ambari	144
21.1. Setting Up Kerberos for Hadoop 2.x	144
21.1.1. Preparing Kerberos	144
21.1.2. Setting Up Hadoop Users	153
21.1.3. Enabling Kerberos Security	155
21.2. Setting Up Kerberos for Hadoop 1.x	156
21.2.1. Preparing Kerberos	157
21.2.2. Setting Up Hadoop Users	165
21.2.3. Enabling Kerberos Security	167

List of Tables

2.1. Download the repo	13
2.2. Ambari Server LDAP Properties	17
3.1. Service Users	33
3.2. Service Group	34
3.3. HDFS Settings: Advanced	34
3.4. MapReduce Settings: Advanced	34
6.1. Key Properties to Check	57
6.2. Properties to Modify	61
7.1. HDFS Ports	65
7.2. MapReduce Ports	66
7.3. YARN Ports	66
7.4. Hive Ports	67
7.5. HBase Ports	67
7.6. HBase Ports	68
7.7. WebHCat Port	69
7.8. Ganglia Ports	69
7.9. MySQL Port	69
7.10. Ambari Web	69
7.11. Nagios	70
8.1. Core-site.xml properties and values for NameNode HA on a cluster using Hue	78
8.2. Set Environment Variables	79
9.1. Service Users	91
9.2. Service Group	92
9.3. HDFS Settings: Advanced	92
9.4. MapReduce Settings: Advanced	92
14.1. HDFS Ports	120
14.2. MapReduce Ports	121
14.3. Hive Ports	121
14.4. HBase Ports	122
14.5. HBase Ports	122
14.6. WebHCat Port	123
14.7. Ganglia Ports	123
14.8. MySQL Port	124
14.9. Ambari Web	124
14.10. Ambari Web	124
15.1. Parameter Options for relocate_host_components.py	126
20.1. Hive Security Authorization Settings	137
21.1. Kerberos terminology	145
21.2. Service Principals	148
21.3. Ambari Principals	149
21.4. Service Keytab File Names	149
21.5. Kerberos terminology	158
21.6. Service Principals	160
21.7. Ambari Principals	161
21.8. Service Keytab File Names	162

Part I. Setting Up Ambari

This section describes installing and preparing Ambari so that it can deploy and manage your Hadoop installation. It includes:

- [Getting Ready](#)
- [Running the Ambari Installer](#)

Hortonworks Data Platform

Installing Hadoop Using Apache Ambari

(Aug 1, 2014)

1. Getting Ready

This section describes the information and materials you need to get ready to install Hadoop using the Apache Ambari Install Wizard. **Apache Ambari** provides an end-to-end management and monitoring application for Apache Hadoop. With Ambari, you can deploy and operate a complete Hadoop stack using a graphical user interface (GUI), manage configuration changes, monitor services, and create alerts for all the nodes in your cluster from a central point.

- [Determine Version Compatibility](#)
- [Understand the Basics](#)
- [Meet Minimum System Requirements](#)
- [Collect Information](#)
- [Prepare the Environment](#)
- [Optional: Configure Ambari for Local Repositories](#)

1.1. Determine Version Compatibility

Use this table to determine whether your Ambari and Hadoop stack versions are compatible.

Ambari	HDP 1.2.0	HDP 1.2.1.	HDP 1.3.0	HDP 1.3.2	HDP 1.3.3	HDP 2.0.6
1.4.4.23				X	X	X
1.4.3.38				X	X	X
1.4.2.104				X	X	X
1.4.1.61				X	X	X
1.4.1.25				X		X
1.2.5.17		X	X	X		
1.2.4.9	X	X	X			
1.2.3.7	X	X	X			
1.2.3.6	X	X				
1.2.2.5	X	X				
1.2.2.4	X	X				
1.2.2.3	X					
1.2.1.2	X					
1.2.0.1	X					

For more information about using Ambari and the Hadoop 1.x stack, see [Hadoop 1.x-Deploying, Configuring, and Upgrading Ambari](#).

For more information about using Ambari and the Hadoop 2.x stack, see [Hadoop 2.x-Deploying, Configuring, and Upgrading Ambari](#).

1.2. Understand the Basics

The Hortonworks Data Platform consists of three layers of components. A coordinated and tested set of these components is sometimes referred to as the Stack.

- **Core Hadoop 1:** The basic components of Apache Hadoop version 1.x.
 - **Hadoop Distributed File System (HDFS)** : A special purpose file system designed to provide high-throughput access to data in a highly distributed environment.
 - **MapReduce:** A framework for performing high volume distributed data processing using the MapReduce programming paradigm.
- **Core Hadoop 2:** The basic components of Apache Hadoop version 2.x.
 - **Hadoop Distributed File System (HDFS):** A special purpose file system designed to provide high-throughput access to data in a highly distributed environment.
 - **YARN:** A resource negotiator for managing high volume distributed data processing. Previously part of the first version of MapReduce.
 - **MapReduce 2 (MR2):** A set of client libraries for computation using the MapReduce programming paradigm and a History Server for logging job and task information. Previously part of the first version of MapReduce.
- **Essential Hadoop:** A set of Apache components designed to ease working with Core Hadoop.
 - **Apache Pig:** A platform for creating higher level data flow programs that can be compiled into sequences of MapReduce programs, using Pig Latin, the platform's native language.
 - **Apache Hive:** A tool for creating higher level SQL queries using HiveQL, the tool's native language, that can be compiled into sequences of MapReduce programs. Included with **Apache HCatalog**.
 - **Apache HCatalog:** A metadata abstraction layer that insulates users and scripts from how and where data is physically stored. Now part of **Apache Hive**. Includes **WebHCat**, which provides a set of REST APIs for HCatalog and related Hadoop components. Originally named **Templeton**.
 - **Apache HBase:** A distributed, column-oriented database that provides the ability to access and manipulate data randomly in the context of the large blocks that make up HDFS.
 - **Apache ZooKeeper:** A centralized tool for providing services to highly distributed systems. ZooKeeper is necessary for HBase installations.
- **Hadoop Support:** A set of components that allow you to monitor your Hadoop installation and to connect Hadoop with your larger compute environment.
 - **Apache Oozie:** A server based workflow engine optimized for running workflows that execute Hadoop jobs.

Running the current Oozie examples requires some reconfiguration from the standard Ambari installation. See [Using HDP for Workflow and Scheduling \(Oozie\)](#)

- **Apache Sqoop:** A component that provides a mechanism for moving data between Hadoop and external structured data stores. Can be integrated with Oozie workflows.
- **Apache Flume:** A log aggregator. This component must be installed manually. It is not supported in the context of Ambari at this time.

See [Installing and Configuring Flume](#) for more information.

- **Ganglia:** An Open Source tool for monitoring high-performance computing systems.
- **Nagios:** An Open Source tool for monitoring systems, services, and networks.

You must always install HDFS, but you can select components from the other layers based on your needs.

1.3. Meet Minimum System Requirements

To run Hadoop, your system must meet minimum requirements.

- [Hardware Recommendations](#)
- [Operating Systems Requirements](#)
- [Browser Requirements](#)
- [Software Requirements](#)
- [JDK Requirements](#)
- [Database Requirements](#)
- [File System Partitioning Recommendations](#)

1.3.1. Hardware Recommendations

There is no single hardware requirement set for installing Hadoop.

For more information on the parameters that may affect your installation, see [Hardware Recommendations For Apache Hadoop](#).

1.3.2. Operating Systems Requirements

The following operating systems are supported:

- Red Hat Enterprise Linux (RHEL) v5.x or 6.x (64-bit)
- CentOS v5.x or 6.x (64-bit)
- Oracle Linux v5.x or 6.x (64-bit)

- SUSE Linux Enterprise Server (SLES) 11, SP1 (64-bit)



Important

The installer pulls many packages from the base OS repos. If you do not have a complete set of base OS repos available to all your machines at the time of installation you may run into issues.

If you encounter problems with base OS repos being unavailable, please contact your system administrator to arrange for these additional repos to be proxied or mirrored. For more information see [Optional: Configure the Local Repositories](#)

1.3.3. Browser Requirements

The Ambari Install Wizard runs as a browser-based Web app. You must have a machine capable of running a graphical browser to use this tool. The supported browsers are:

- Windows (Vista, 7)
 - Internet Explorer 9.0 and higher (for Vista + Windows 7)
 - Firefox latest stable release
 - Safari latest stable release
 - Google Chrome latest stable release
- Mac OS X (10.6 or later)
 - Firefox latest stable release
 - Safari latest stable release
 - Google Chrome latest stable release
- Linux (RHEL, CentOS, SLES, Oracle Linux)
 - Firefox latest stable release
 - Google Chrome latest stable release

1.3.4. Software Requirements

On each of your hosts:

- yum and rpm (RHEL/CentOS/Oracle Linux)
- zypper (SLES)
- scp, curl, and wget
- python (2.6 or later)



Important

The Python version shipped with SUSE 11, 2.6.0-8.12.2, has a critical bug that may cause the Ambari Agent to fail within the first 24 hours. If you are installing on SUSE 11, please update all your hosts to Python version 2.6.8-0.15.1.

1.3.5. JDK Requirements

The following Java runtimes are supported:

- Oracle JDK 1.6.0_31 64-bit
- Oracle JDK 1.7 64-bit
- OpenJDK 7 64-bit

1.3.6. Database Requirements

Hive/HCatalog, Oozie, and Ambari all require their own internal databases.

- Hive/HCatalog: By default uses an Ambari-installed MySQL 5.x instance. With appropriate preparation, you can also use an existing MySQL 5.x or Oracle 11g r2 instance. See [Using Non-Default Databases](#) for more information on using existing instances.
- Oozie: By default uses an Ambari-installed Derby instance. With appropriate preparation, you can also use an existing MySQL 5.x or Oracle 11g r2 instance. See [Using Non-Default Databases](#) for more information on using existing instances.
- Ambari: By default uses an Ambari-installed PostgreSQL 8.x instance. With appropriate preparation, you can also use an existing Oracle 11g r2, or MySQL 5.x instance. See [Using Non-Default Databases](#) for more information on using existing instances.

1.3.7. File System Partitioning Recommendations

For information on setting up file system partitions on master and slave nodes in a HDP cluster, see [File System Partitioning Recommendations](#).

1.4. Collect Information

To deploy your Hadoop installation, you need to collect the following information:

- The fully qualified domain name (FQDN) for each host in your system, and which components you want to set up on which host. The Ambari install wizard *does not* support using IP addresses. You can use `hostname -f` to check for the FQDN if you do not know it.



Note

While it is possible to deploy all of Hadoop on a single host, this is appropriate only for initial evaluation. In general you should use at least three hosts: one master host and two slaves.

- The base directories you want to use as mount points for storing:
 - NameNode data
 - DataNodes data
 - Secondary NameNode data
 - Oozie data
 - MapReduce data (Hadoop version 1.x)
 - YARN data (Hadoop version 2.x)
 - ZooKeeper data, if you install ZooKeeper
 - Various log, pid, and db files, depending on your install type

1.5. Prepare the Environment

To deploy your Hadoop instance, you need to prepare your deploy environment:

- [Check Existing Installs](#)
- [Set up Password-less SSH](#)
- [Set up Users and Groups](#)
- [Enable NTP on the Cluster](#)
- [Check DNS](#)
- [Configure iptables](#)
- [Disable SELinux, PackageKit and Check umask Value](#)

1.5.1. Check Existing Installs

Ambari automatically installs the correct versions of the files that are necessary for Ambari and Hadoop to run. Versions other than the ones that Ambari uses can cause problems in running the installer, so remove any existing installs that do not match the following lists.

	RHEL/CentOS/Oracle Linux v5	RHEL/CentOS/Oracle Linux v6	SLES 11
Ambari Server	<ul style="list-style-type: none"> • libffi 3.0.5-1.el5 • python26 2.6.8-2.el5 • python26-libs 2.6.8-2.el5 • postgresql 8.4.13-1.el6_3 • postgresql-libs 8.4.13-1.el6_3 • postgresql-server 8.4.13-1.el6_3 	<ul style="list-style-type: none"> • postgresql 8.4.13-1.el6_3 • postgresql-libs 8.4.13-1.el6_3 • postgresql-server 8.4.13-1.el6_3 	<ul style="list-style-type: none"> • postgresql 8.3.5-1 • postgresql-server 8.3.5-1 • postgresql-libs 8.3.5-1
Ambari Agent ^a	<ul style="list-style-type: none"> • libffi 3.0.5-1.el5 	None	None

	RHEL/CentOS/Oracle Linux v5	RHEL/CentOS/Oracle Linux v6	SLES 11
	<ul style="list-style-type: none"> python26 2.6.8-2.el5 python26-libs 2.6.8-2.el5 		
Nagios Server ^b	<ul style="list-style-type: none"> nagios 3.5.0-99 nagios-devel 3.5.0-99 nagios-www 3.5.0-99 nagios-plugins 1.4.9-1 	<ul style="list-style-type: none"> nagios 3.5.0-99 nagios-devel 3.5.0-99 nagios-www 3.5.0-99 nagios-plugins 1.4.9-1 	<ul style="list-style-type: none"> nagios 3.5.0-99 nagios-devel 3.5.0-99 nagios-www 3.5.0-99 nagios-plugins 1.4.9-1
Ganglia Server ^c	<ul style="list-style-type: none"> ganglia-gmetad 3.5.0-99 ganglia-devel 3.5.0-99 libganglia 3.5.0-99 ganglia-web 3.5.7-99 rrdtool 1.4.5-1.el5 	<ul style="list-style-type: none"> ganglia-gmetad 3.5.0-99 ganglia-devel 3.5.0-99 libganglia 3.5.0-99 ganglia-web 3.5.7-99 rrdtool 1.4.5-1.el6 	<ul style="list-style-type: none"> ganglia-gmetad 3.5.0-99 ganglia-devel 3.5.0-99 libganglia 3.5.0-99 ganglia-web 3.5.7-99 rrdtool 1.4.5-4.5.1
Ganglia Monitor ^d	<ul style="list-style-type: none"> ganglia-gmond 3.5.0-99 libganglia 3.5.0-99 	<ul style="list-style-type: none"> ganglia-gmond 3.5.0-99 libganglia 3.5.0-99 	<ul style="list-style-type: none"> ganglia-gmond 3.5.0-99 libganglia 3.5.0-99

^aInstalled on each host in your cluster. Communicates with the Ambari Server to execute commands.

^bThe host that runs the Nagios server

^cThe host that runs the Ganglia Server

^dInstalled on each host in the cluster. Sends metrics data to the Ganglia Collector.

1.5.2. Set Up Password-less SSH

To have Ambari Server automatically install Ambari Agents in all your cluster hosts, you must set up password-less SSH connections between the main installation (Ambari Server) host and all other machines. The Ambari Server host acts as the client and uses the key-pair to access the other hosts in the cluster to install the Ambari Agent.



Note

You can choose to install the Agents on each cluster host manually. In this case you do not need to setup SSH. See [Appendix: Installing Ambari Agents Manually](#) for more information.

1. Generate public and private SSH keys on the Ambari Server host

```
ssh-keygen
```

2. Copy the SSH Public Key (id_rsa.pub) to the root account on your target hosts.

```
.ssh/id_rsa
.ssh/id_rsa.pub
```

3. Add the SSH Public Key to the authorized_keys file on your target hosts.

```
cat id_rsa.pub >> authorized_keys
```

4. Depending on your version of SSH, you may need to set permissions on the .ssh directory (to 700) and the authorized_keys file in that directory (to 600) on the target hosts.

```
chmod 700 ~/.ssh
```



```
chmod 600 ~/.ssh/authorized_keys
```

5. From the Ambari Server, make sure you can connect to each host in the cluster using SSH.

```
ssh root@{remote.target.host}
```

You may see this warning. This happens on your first connection and is normal.

```
Are you sure you want to continue connecting (yes/no)?
```

6. Retain a copy of the SSH Private Key on the machine from which you will run the web-based Ambari Install Wizard.



Note

It is possible to use a non-root SSH account, if that account can execute `sudo` without entering a password.

1.5.3. Set up Users and Groups

The Ambari cluster installer automatically creates user and group accounts for you. Ambari preserves any existing user and group accounts, and uses these accounts when configuring Hadoop services. User and group creation applies to user/group accounts on the local operating system and to LDAP/AD accounts.

To set up custom accounts before running the Ambari installer, see [Service Users and Groups \(for the 1.x stack\)](#) or [Service Users and Groups \(for the 2.x stack\)](#) for more information about customizing service users and groups.

1.5.4. Enable NTP on the Cluster and on the Browser Host

The clocks of all the nodes in your cluster and the machine that runs the browser through which you access Ambari Web must be able to synchronize with each other.

1.5.5. Check DNS

All hosts in your system must be configured for DNS and Reverse DNS.

If you are unable to configure DNS and Reverse DNS, you must edit the hosts file on every host in your cluster to contain the address of each of your hosts and to set the Fully Qualified Domain Name hostname of each of those hosts. The following instructions cover basic hostname network setup for generic Linux hosts. Different versions and flavors of Linux might require slightly different commands. Please refer to your specific operating system documentation for the specific details for your system.

1.5.5.1. Edit the Host File

1. Using a text editor, open the hosts file on every host in your cluster. For example:

```
vi /etc/hosts
```

2. Add a line for each host in your cluster. The line should consist of the IP address and the FQDN. For example:

```
1.2.3.4 fully.qualified.domain.name
```



Note

Do **not** remove the following two lines from your host file, or various programs that require network functionality may fail.

```
127.0.0.1 localhost.localdomain localhost  
::1 localhost6.localdomain6 localhost6
```

1.5.5.2. Set the Hostname

1. Use the "hostname" command to set the hostname on each host in your cluster. For example:

```
hostname fully.qualified.domain.name
```

2. Confirm that the hostname is set by running the following command:

```
hostname -f
```

This should return the name you just set.

1.5.5.3. Edit the Network Configuration File

1. Using a text editor, open the network configuration file on every host. This file is used to set the desired network configuration for each host. For example:

```
vi /etc/sysconfig/network
```

2. Modify the HOSTNAME property to set the fully.qualified.domain.name.

```
NETWORKING=yes  
NETWORKING_IPV6=yes  
HOSTNAME=fully.qualified.domain.name
```

1.5.6. Configuring iptables

For Ambari to communicate during setup with the hosts it deploys to and manages, certain ports must be open and available. The easiest way to do this is to temporarily disable iptables.

```
chkconfig iptables off  
/etc/init.d/iptables stop
```

You can restart iptables after setup is complete.

If the security protocols at your installation do not allow you to disable iptables, you can proceed with them on, as long as all of the relevant ports are open and available. If you plan to run with them enabled, see [Configuring Ports \(for the 1.x stack\)](#) or [Configuring Ports \(for the 2.x stack\)](#) for more information on the necessary ports per component.

During the Ambari Server setup process, Ambari checks to see if iptables is running. If it is, a warning prints to remind you to check that the necessary ports are open and available.

The **Host Confirm** step of the Cluster Install Wizard will also issue a warning for each host that has `iptables` running.



Important

If you leave `iptables` enabled and do not set up the necessary ports, the cluster installation will fail.

1.5.7. Disable SELinux and PackageKit and check the umask Value

1. SELinux must be temporarily disabled for the Ambari setup to function. Run the following command on each host in your cluster:

```
setenforce 0
```

2. On the RHEL/CentOS installation host, if PackageKit is installed, open `/etc/yum/pluginconf.d/refresh-packagekit.conf` with a text editor and make this change:

```
enabled=0
```



Note

PackageKit is not enabled by default on SLES. Unless you have specifically enabled it, you do not need to do this step.

3. Make sure `umask` is set to 022.

1.6. Optional: Configure Ambari for Local Repositories

If your cluster does **not** have access to the Internet, or you are creating a large cluster and you want to conserve bandwidth, you need to provide access to the bits using an alternative method.

1. Set up the local mirror repositories as needed for HDP and HDP Utils.

For more information on your options, see [Deploying HDP In Production Data Centers with Firewalls](#).

2. From the Ambari Server host, fetch the Ambari repository file or RPM package as described in [Set Up the Bits](#). You need a connection to the Internet for this step.

If you do not have a connection to the Internet for this machine, you should follow the instructions in [Deploying HDP In Production Data Centers with Firewalls](#) and be sure to perform the optional steps for setting up the Ambari local repository.

3. Be sure to remember the local mirror repository Base URLs for each operating system from the [Deploying HDP In Production Data Centers with Firewalls](#) steps. You will use these Base URLs during the cluster install.

For example, if your system includes hosts running CentOS 6, to point to the HDP 2.0.6 repositories, your local repository Base URL would look something like this:

```
http://{your.hosted.local.repository}/HDP-2.0.6/repos/centos6
```



Important

The appropriate relative path depends on how you have setup your local repositories. If you have mixed operating systems in your cluster (for example, CentOS 6 and RHEL 6), you must configure the repositories for both OSes - centos6 and redhat6. If you do not, some hosts in your cluster will not be able to retrieve the software packages for their operating system.

4. Configure which JDK to use and how the JDK will be downloaded and installed.

If you have not already installed the JDK on all hosts, and plan to use Oracle JDK 1.6, download [jdk-6u31-linux-x64.bin](#) to `/var/lib/ambari-server/resources`.

- If you plan to use a JDK other than Oracle JDK 1.6, you **must** install the JDK on each host in your cluster and use the `-j` flag when running `ambari-server setup`. See [JDK Requirements](#) for more information on supported JDKs.
- If you have already installed the JDK on all your hosts, you **must** use the `-j` flag when running `ambari-server setup`



Note

See [Setup Options](#) for more information on the `-j` flag.

2. Running the Ambari Installer

This section describes the process for installing Apache Ambari. Ambari manages installing and deploying Hadoop.

2.1. Set Up the Bits

1. Log into the machine that is to serve as the Ambari Server as `root`. You may login and `sudo` as `su` if this is what your environment requires. This machine is the main installation host.
2. Download the the Ambari repository file and copy it to your `repos.d`.

Table 2.1. Download the repo

Platform	Access
RHEL, CentOS, and Oracle Linux 5	<pre>wget http://public-repo-1.hortonworks.com/ambari/centos5/1.x/updates/1.4.4.23/ambari.repo</pre> <pre>cp ambari.repo /etc/yum.repos.d</pre>
RHEL, CentOS and Oracle Linux 6	<pre>wget http://public-repo-1.hortonworks.com/ambari/centos6/1.x/updates/1.4.4.23/ambari.repo</pre> <pre>cp ambari.repo /etc/yum.repos.d</pre>
SLES 11	<pre>wget http://public-repo-1.hortonworks.com/ambari/suse11/1.x/updates/1.4.4.23/ambari.repo</pre> <pre>cp ambari.repo /etc/zypp/repos.d</pre>



Note

If your cluster does not have access to the Internet, or you are creating a large cluster and you want to conserve bandwidth, you need to provide access to the bits using an alternative method. For more information, see [Optional: Configure the Local Repositories](#) section.

When you have the software, continue your installation based on your base platform.

2.1.1. RHEL/CentOS/Oracle Linux 5.x

1. Confirm the repository is configured by checking the repo list.

```
yum repolist
```

You should see the Ambari and HDP utilities repositories in the list. The version values vary depending the installation.

```
repo id          repo name
| AMBARI-1.x      | Ambari 1.x
| HDP-UTILS-1.1.0.16 | Hortonworks Data Platform Utils
```

2. Install the Ambari bits using `yum`. This also installs PostgreSQL.

```
yum install ambari-server
```

2.1.2. RHEL/CentOS/Oracle Linux 6.x

1. Confirm the repository is configured by checking the repo list.

```
yum repolist
```

You should see the Ambari and HDP utilities repositories in the list. The version values vary depending the installation.

```
repo id          repo name
| AMBARI-1.x     | Ambari 1.x
| HDP-UTILS-1.1.0.16 | Hortonworks Data Platform Utils
```

2. Install the Ambari bits using yum. This also installs PostgreSQL.

```
yum install ambari-server
```

2.1.3. SLES 11

1. Verify that `php_curl` is installed before you begin.

```
zypper se 'php_curl*'
```

If `php_curl` is not installed, install it:

```
zypper in php_curl
```

2. Confirm the downloaded repository is configured by checking the repo list.

```
zypper repos
```

You should see the Ambari and HDP utilities in the list. The version values vary depending the installation.

```
# | Alias          | Name
1 | AMBARI.dev-1.x | Ambari 1.x
2 | HDP-UTILS-1.1.0.16 | Hortonworks Data Platform Utils
```

3. Install the Ambari bits using zypper. This also installs PostgreSQL.

```
zypper install ambari-server
```

2.2. Set Up the Server

The `ambari-server` command manages the setup process. Follow the prompts:

```
ambari-server setup
```

1. If you have *not* temporarily disabled SELinux, you may get a warning. Enter `y` to continue.
2. By default, Ambari Server runs under `root`. If you want to create a different user to run the Ambari Server instead, or to assign a previously created user, select `y` at **Customize user account for ambari-server daemon** and give the prompt the username you want to use.

3. If you have not temporarily disabled `iptables` you may get a warning. Enter `y` to continue. See [Configuring Ports for \(2.x\)](#) or [\(1.x\)](#) for more information on the ports that must be open and accessible.
4. Agree to the Oracle JDK license when asked. You must accept this license to be able to download the necessary JDK from Oracle. The JDK is installed during the deploy phase.



Note

By default, Ambari Server setup will download and install Oracle JDK 1.6. If you plan to download this JDK and install on all your hosts, or plan to use a different version of the JDK, skip this step and see [Setup Options](#) for more information

5. At **Enter advanced database configuration**:

- To use the default PostgreSQL database, named `ambari`, with the default username and password (`ambari/bigdata`), enter 1.



Important

If you are using an existing Oracle or MySQL database instance, you must prepare the instance using the steps detailed in [Using Non-Default Databases](#) **before** running the installer.

- To use an existing Oracle 11g r2 instance, and select your own database name, username and password for that database, enter 2.

Select the database you want to use and provide any information required by the prompts, including hostname, port, Service Name or SID, username, and password.

- To use an existing MySQL 5.x database, and select your own database name, username and password for that database, enter 3.

Select the database you want to use and provide any information required by the prompts, including hostname, port, database name, username, and password.

6. Setup completes.



Note

If your host accesses the Internet through a proxy server, you must configure Ambari Server to use this proxy server. See [Configure Ambari Server for Internet Proxy](#) for more information.

2.2.1. Setup Options

The following table describes options frequently used for Ambari Server setup.

Option	Description
<code>-j</code> <code>-java-home</code>	Specifies the <code>JAVA_HOME</code> path to use on the Ambari Server and all hosts in the cluster. By default when you do

Option	Description
	<p>not specify this option, Setup downloads the Oracle JDK 1.6 binary to <code>/var/lib/ambari-server/resources</code> on the Ambari Server and installs the JDK to <code>/usr/jdk64</code>.</p> <p>Use this option when you plan to use a JDK other than the default Oracle JDK 1.6. See JDK Requirements for more information on the supported JDKs. If you are using an alternate JDK, you must manually install the JDK on all hosts and specify the Java Home path during Ambari Server setup.</p> <p>This path must be valid on all hosts. For example.</p> <pre>ambari-server setup -j /usr/java/default</pre>
-s --silent	Setup runs silently. Accepts all default prompt values.
-v --verbose	Prints verbose info and warning messages to the console during Setup.
-c --jce-policy	Use specified <code>jce_policy</code> . Must be valid on the Ambari Server host.
-i --jdk-location	Use specified JDK file in local filesystem instead of downloading
-g --debug	Start Ambari Server in debug mode

2.3. Optional: Set Up LDAP or Active Directory Authentication

By default Ambari uses an internal database as the user store for authentication and authorization. If you want to add LDAP or Active Directory (AD) external authentication in addition for Ambari Web, you need to collect the following information and run a special setup command. Ambari Server must not be running when you execute this command. An LDAP client must be installed on the Ambari Server host.



Important

Ambari Server should not be running when you do this: either make the edits before you start Ambari Server the first time or bring the server down to make the edits.

- The following table details the properties and values you need to know to set up LDAP authentication.



Note

If you are going to set `bindAnonymously` to `false` (the default), you need to make sure you have an LDAP Manager name and password set up. If you are going to use SSL, you need to make sure you have already set up your certificate and keys.

Table 2.2. Ambari Server LDAP Properties

Property	Values	Description
authentication.ldap.primaryUrl	server:port	The hostname and port for the LDAP or AD server. Example: my.ldap.server:389
authentication.ldap.secondaryUrl	server:port	The hostname and port for the secondary LDAP or AD server. Example: my.secondary.ldap.server:389 This is an optional value.
authentication.ldap.useSSL	true or false	If true, use SSL when connecting to the LDAP or AD server.
authentication.ldap.usernameAttribute	[LDAP attribute]	The attribute for username Example: uid
authentication.ldap.baseDn	[Distinguished Name]	The root Distinguished Name to search in the directory for users. Example: ou=people,dc=hadoop,dc=apache,dc=org
authentication.ldap.bindAnonymously	true or false	If true, bind to the LDAP or AD server anonymously
authentication.ldap.managerDn	[Full Distinguished Name]	If Bind anonymous is set to false, the Distinguished Name ("DN") for the manager. Example: uid=hdfs,ou=people,dc=hadoop,dc=apache,dc=org
authentication.ldap.managerPassword	[password]	If Bind anonymous is set to false, the password for the manager

**Note**

Ambari does not set up LDAP automatically. You must set up LDAP manually.

2. If the LDAPS server certificate is signed by a trusted Certificate Authority, there is no need to import the certificate into Ambari so this section does not apply to you. If the LDAPS server certificate is self-signed, or is signed by an unrecognized certificate authority such as an internal certificate authority, you must import the certificate and create a keystore file. The following example creates a keystore file at `/keys/ldaps-keystore.jks`, but you can create it anywhere in the file system:
 - a. On the Ambari server:
 - b. `mkdir /keys`
 - c. `$JAVA_HOME/bin/keytool -import -trustcacerts -alias root -file $PATH_TO_YOUR_LDAPS_CERT -keystore /keys/ldaps-keystore.jks`

d. Set a password when prompted. You will use this during `ambari-server setup-ldap`. Run the LDAP setup command and answer the prompts with the information you collected above:

```
ambari-server setup-ldap
```

- a. At the **Primary URL*** prompt, enter the server URL and port you collected above. Prompts marked with an asterisk are required values.
- b. At the **Secondary URL** prompt, enter the secondary server URL and port. This is optional value.
- c. At the **Use SSL*** prompt, enter your selection. **If using LDAPS, enter true.**
- d. At the **User name attribute*** prompt, enter your selection. The default value is `uid`.
- e. At the **Base DN*** prompt, enter your selection.
- f. At the **Bind anonymously*** prompt, enter your selection.
- g. At the **Manager DN*** prompt, enter your selection if you have set `bind.Anonymously` to `false`.
- h. At the **Enter the Manager Password*** , enter the password for your LDAP manager.
- i. If you set **Use SSL*** = `true` in step 3, the following prompt appears: **Do you want to provide custom TrustStore for Ambari?**

Consider the following options and respond as appropriate.

- **More secure option:** If using a self-signed certificate that you do not want imported to the existing JDK keystore, enter **y**.

For example, you want this certificate used only by Ambari, not by any other applications run by JDK on the same host.

If you choose this option, additional prompts appear. Respond to the additional prompts as follows:

- i. At the **TrustStore type** prompt, enter **jks**.
 - ii. At the **Path to TrustStore file** prompt, enter **/keys/ldaps-keystore.jks** (or the actual path to your keystore file).
 - iii. At the **Password for TrustStore** prompt, enter the password that you defined for the keystore.
- **Less secure option:** If using a self-signed certificate that you want to import and store in the existing, default JDK keystore, enter **n**.
 - i. Convert the SSL certificate to X.509 format, if necessary, by executing the following command:

```
openssl x509 -in slapd.pem -out slapd.crt
```

Where `slapd.crt` is the path to the X.509 certificate.

- ii. Import the SSL certificate to the existing keystore, for example the default jre certificates storage, using the following instruction:

```
/usr/jdk64/jdk1.7.0_45/bin/keytool -import -trustcacerts -file slapd.  
crt -keystore  
/usr/jdk64/jdk1.7.0_45/jre/lib/security/cacerts
```

Where Ambari is set up to use JDK 1.7. Therefore, the certificate must be imported in the JDK 7 keystore.

- j. Review your settings and if they are correct, select y.
- k. Start or restart the Server

```
ambari-server restart
```

Initially the users you have enabled all have Ambari User privileges. Ambari Users can read metrics, view service status and configuration, and browse job information. For these new users to be able to start or stop services, modify configurations, and run smoke tests, they need to be Admins. To make this change, use Ambari Web **Admin -> Users -> Edit**.

2.4. Optional: Set Up Security for Ambari

There are four ways you can increase the security settings for your Ambari server installation.



Note

If you plan to configure your cluster for Kerberos, you may use the **Setup Ambari kerberos JAAS configuration** option, which is described in [Setting up Kerberos for Use with Ambari](#).

- [Set Up HTTPS for Ambari Server](#)
- [Set Up HTTPS for Ganglia](#)
- [Set Up HTTPS for Nagios](#)
- [Encrypt Database and LDAP Passwords](#)

2.4.1. Set Up HTTPS for Ambari Server

If you want to limit access to the Ambari Server to HTTPS connections, you need to provide a certificate. While it is possible to use a self-signed certificate for initial trials, they are not suitable for production environments. After your certificate is in place, you must run a special setup command.



Important

Ambari Server should not be running when you do this. Either make these changes before you start Ambari the first time, or bring the server down before running the setup command.

1. Log into the Ambari Server host.
2. Locate your certificate. If you want to create a temporary self-signed certificate, use this as an example:

```
openssl genrsa -out $wserver.key 2048
openssl req -new -key $wserver.key -out $wserver.csr
openssl x509 -req -days 365 -in $wserver.csr -signkey $wserver.key -out
$wserver.crt
```

Where `$wserver` is the Ambari Server hostname.



Important

The certificate you use must be PEM-encoded, not DER-encoded. If you attempt to use a DER-encoded certificate, you see this error:

```
unable to load certificate
140109766494024:error:0906D06C:PEM routines:PEM_read_bio:no start
line:pem_lib.c
:698:Expecting: TRUSTED CERTIFICATE
```

You can convert a DER-encoded certificate to a PEM-encoded certificate using the following command:

```
openssl x509 -in cert.crt -inform der -outform pem -out cert.pem
```

where `cert.crt` is the DER-encoded certificate and `cert.pem` is the resulting PEM-encoded certificate.

3. Run the special setup command and answer the prompts

```
ambari-server setup-security
```

- a. Select 1 for **Enable HTTPS for Ambari server**.
- b. Respond `y` to **Do you want to configure HTTPS?**
- c. Select the port you want to use for SSL. Default is 8443.
- d. Provide the path to your certificate and your private key. For example, put your certificate and private key in `/etc/ambari-server/certs` with `root` as the owner or the non-root user you designated during Ambari Server setup for the `ambari-server` daemon.
- e. Provide the password for the private key.
- f. Start or restart the Server

```
ambari-server restart
```

2.4.2. Set Up HTTPS for Ganglia

If you want Ganglia to use HTTPS instead of the default HTTP to communicate with Ambari Server, use the following instructions.



Important

The servers should not be running when you do this: either make the edits before you start Ambari Server the first time or bring the servers down to make the edits.

1. Set up the Ganglia server.

- a. Log into the Ganglia server host.
- b. Create a self-signed certificate on the Ganglia server host. For example:

```
openssl genrsa -out $gserver.key 2048
openssl req -new -key $gserver.key -out $gserver.csr
openssl x509 -req -days 365 -in $gserver.csr -signkey $gserver.key -out
  $gserver.crt
```

Where `$gserver` is the Ganglia server hostname.

- c. Install SSL on the Ganglia server host.

```
yum install mod_ssl
```

- d. Edit the SSL configuration file on the Ganglia server host.

- i. Using a text editor, open:

```
/etc/httpd/conf.d/ssl.conf
```

- ii. Add lines setting the certificate and key file names to the files you created [above \[21\]](#). For example:

```
SSLCertificateFile    $gserver.crt
SSLCertificateKeyFile $gserver.key
```

- e. Disable HTTP access (optional)

- i. Using a text editor, open:

```
/etc/httpd/conf/httpd.conf
```

- ii. Comment out the port 80 listener:

```
# Listen 80
```

- f. Restart the `httpd` service on the Ganglia server host.

```
service httpd restart
```

2. Set up and restart the Ambari Server.

- a. Log into the Ambari Server.

- b. Run the special setup command and answer the prompts.

```
ambari-server setup-security
```

- i. Select 2 for **Enable HTTPS for Ganglia service**.
 - ii. Respond **y** to **Do you want to configure HTTPS for Ganglia service**.
 - iii. Enter your TrustStore type. Your options are `jks`, `jceks`, or `pks12`.
 - iv. Enter the path to your TrustStore file.
 - v. Enter the password for your TrustStore and then re-enter to confirm. The password must be at least 6 characters long.
 - vi. Enter the path to the Ganglia server certificate file.
- c. Start or restart the Server

```
ambari-server restart
```

2.4.3. Set Up HTTPS for Nagios

If you want Nagios to use HTTPS instead of HTTP (the default), use the following instructions.



Important

The servers should not be running when you do this: either make the edits before you start Ambari Server the first time or bring the servers down to make the edits.

1. Set up the Nagios server.
 - a. Log into the Nagios server host.
 - b. Create a self-signed certificate on the Nagios server host. For example:

```
openssl genrsa -out $nserver.key 2048
openssl req -new -key $nserver.key -out $nserver.csr
openssl x509 -req -days 365 -in $nserver.csr -signkey $nserver.key -out
$nserver.crt
```

Where `$nserver` is the Nagios server hostname.

- c. Install SSL on the Nagios server host.

```
yum install mod_ssl
```

- d. Edit the SSL configuration file on the Nagios server host.

- i. Using a text editor, open:

```
/etc/httpd/conf.d/ssl.conf
```

- ii. Add lines setting the certificate and key file names to the files you created [previously \[22\]](#). For example:

```
SSLCertificateFile    $nserver.crt
SSLCertificateKeyFile $nserver.key
```

- e. Disable HTTP access (optional)

- i. Using a text editor, open:

```
/etc/httpd/conf/httpd.conf
```

- ii. Comment out the port 80 listener:

```
# Listen 80
```

- f. Restart the `httpd` service on the Nagios server host.

```
service httpd restart
```

2. Set up and restart the Ambari Server.

- a. Log into the Ambari Server.

- b. Run the special setup command and answer the prompts.

```
ambari-server setup-security
```

- i. Select 2 for **Enable HTTPS for Nagios service**.
- ii. Respond `y` to **Do you want to configure HTTPS for Nagios?**
- iii. Enter your TrustStore type. Your options are `jks`, `jceks`, or `pks12`.
- iv. Enter the path to your TrustStore file.
- v. Enter the password for your TrustStore and then re-enter to confirm. The password must be at least 6 characters long.
- vi. Enter the path to the Nagios server certificate file.

- c. Start or restart the Server

```
ambari-server restart
```

2.4.4. Optional: Encrypt Database and LDAP Passwords

By default the passwords for Ambari's database and for access to the LDAP server are stored in a plain text configuration file. To have those passwords encrypted, you need to run a special setup command.



Important

Ambari Server should not be running when you do this: either make the edits before you start Ambari Server the first time or bring the server down to make the edits.

1. On the Ambari Server, run the special setup command and answer the prompts:

```
ambari-server setup-security
```

- a. Select 4 for **Encrypt passwords stored in ambari.properties file**.
- b. Provide a master key for encrypting the passwords. You are prompted to enter the key twice for accuracy.



Important

If your passwords are encrypted, you need access to the master key to start Ambari Server.

- c. You have three options for maintaining the master key:
 - At the **Persist** prompt, select `y`. This stores the key in a file on the server.
 - Create an environment variable `AMBARI_SECURITY_MASTER_KEY` and set it to the key.
 - Provide the key manually at the prompt on server startup.
- d. Start or restart the Server

```
ambari-server restart
```

2.4.4.1. Reset Encryption

There may be situations in which you want to:

- [Remove encryption entirely](#)
- [Change the current master key](#), either because the key has been forgotten or because you want to change the current key as a part of a security routine.



Important

Ambari Server should not be running when you do this.

2.4.4.1.1. Remove Encryption Entirely

To reset Ambari database and LDAP passwords to a completely unencrypted state:

1. On the Ambari host, open `/etc/ambari-server/conf/ambari.properties` with a text editor and set this property

```
security.passwords.encryption.enabled=false
```

2. Delete `/var/lib/ambari-server/keys/credentials.jceks`
3. Delete `/var/lib/ambari-server/keys/master`
4. You must now reset the database password and, if necessary, the LDAP password. Run `ambari-server setup` and `ambari-server setup-ldap` again.

2.4.4.1.2. Change the Current Master Key

To change the master key:

- If you know the current master key or if the current master key has been persisted:

1. Re-run the encryption setup command and follow the prompts.

```
ambari-server setup-security
```

- a. Select 4 for **Encrypt passwords stored in ambari.properties file**.
- b. Enter the current master key when prompted if necessary (if it is not persisted or set as an environment variable).
- c. At the **Do you want to reset Master Key** prompt, enter **yes**.
- d. At the prompt, enter the new master key and confirm.

- If you do **not** know the current master key:

1. Remove encryption entirely, as described [here](#).
2. Re-run `ambari-server setup-security` as described [here](#).
3. Start or restart the Ambari Server.

```
ambari-server restart
```

2.5. Optional: Set Up Two-Way SSL Between Ambari Server and Ambari Agents

Two-way SSL provides a way to encrypt communication between Ambari Server and Ambari Agents. By default Ambari ships with Two-way SSL disabled. To enable Two-way SSL:



Important

Ambari Server should not be running when you do this: either make the edits before you start Ambari Server the first time or bring the server down to make the edits.

1. On the Ambari Server host, open `/etc/ambari-server/conf/ambari.properties` with a text editor.
2. Add the following property:

```
security.server.two_way_ssl = true
```

3. Start or restart the Ambari Server.

```
ambari-server restart
```

The Agent certificates are downloaded automatically during Agent Registration.

2.6. Optional: Change the Ambari Server Port

By default Ambari uses port 8080 for access to Ambari Web and the REST API. If you want to change the port number, you need to edit the Ambari properties file.



Important

Ambari Server should not be running when you do this: either make the edits before you start Ambari Server the first time or bring the server down to make the edits.

1. On the Ambari Server host, open `/etc/ambari-server/conf/ambari.properties` with a text editor.
2. Add the client API port property and set it to your desired port value:

```
client.api.port=<port_number>
```

3. Start (or re-start) the Ambari Server. You can now access Ambari Web via the newly configured port:

```
http://{your.ambari.server}:<port_number>
```

2.7. Optional: Configure Ambari Server for Internet Proxy

If you plan to use the public repositories for installing the Stack, Ambari Server must have Internet access to confirm access to the repositories and validate the repositories. If your machine requires use of an Internet Proxy, you must configure Ambari Server to use the proxy.

1. On the Ambari Server, edit `/var/lib/ambari-server/ambari-env.sh`.
2. Add `"-Dhttp.proxyHost=myproxyhost-Dhttp.proxyPort=1234"` to the `AMBARI_JVM_ARGS`.
3. Restart the Ambari Server to pick up this change.



Note

If you plan to use local repositories, see [Optional: Configure Ambari for Local Repositories](#), configuring Ambari to use a proxy server and have Internet access is not required. The Ambari Server must have access to your local repositories.

2.8. Start the Ambari Server

- To start the Ambari Server:

```
ambari-server start
```

- To check the Ambari Server processes:

```
ps -ef | grep Ambari
```

- To stop the Ambari Server:

```
ambari-server stop
```

Part II. Hadoop 2.x - Deploying, Configuring, and Upgrading Ambari

This section describes setting up Hadoop 2.x. It includes:

- [Installing, Configuring, and Deploying the Cluster for Hadoop 2.x](#)
- [Troubleshooting Ambari Deployments for Hadoop 2.x](#)
- [Appendix: Upgrading Ambari Server to 1.4.4](#)
- [Appendix: Upgrading the HDP Stack from 1.3.2. to 2.0.6](#)
- [Appendix: Configuring Ports for Hadoop 2.x](#)

Hortonworks Data Platform

Installing Hadoop Using Apache Ambari

(Aug 1, 2014)

3. Hadoop 2.x - Installing, Configuring, and Deploying the Cluster

This section describes using the Ambari install wizard in your browser to complete your installation, configuration and deployment of Hadoop.

3.1. Log into Apache Ambari

Once you have started the Ambari service, you can access the Ambari Install Wizard through your browser.

1. Point your browser to `http://{your.ambari.server}:8080`.
2. Log in to the Ambari Server using the default username/password: `admin/admin`. You can change this later to whatever you want.


3.2. Welcome

The first step creates the cluster name.

1. At the **Welcome** page, type a name for the cluster you want to create in the text box. No whitespaces or special characters can be used in the name.
2. Click the **Next** button.

3.3. Select Stack

The Service Stack (or simply the Stack) is a coordinated and tested set of Hadoop components. Use the radio button to select the Stack version you want to install. To install a Hadoop 2 stack, select HDP 2.0.6, in Stacks.



The screenshot shows the 'Select Stack' step of the Ambari installation wizard. At the top, there is a heading 'Select Stack' and a light blue instruction box that says 'Please select the service stack that you want to use to install your Hadoop cluster.' Below this, under the heading 'Stacks', there are three radio button options: 'HDP 2.0.6' (which is selected), 'HDP 1.3.3', and 'HDP 1.3.2'. Below the radio buttons is a grey bar with a right-pointing arrow and the text 'Advanced Repository Options'. At the bottom of the form, there are two buttons: a grey '← Back' button on the left and a green 'Next →' button on the right.

Under Advanced Repository Options, you can specify the Base URLs of your local repositories for each Operating System you plan to use in your cluster. You should have configured the Base URLs for your local repositories in [Optional: Configure Ambari for Local Repositories](#).

▼ **Advanced Repository Options**

Customize the repository Base URLs for downloading the Stack software packages. If your hosts do not have access to the internet, you will have to create a local mirror of the Stack repository that is accessible by all hosts and use those Base URLs here.

Important: When using local mirror repositories, you only need to provide Base URLs for the Operating System you are installing for your Stack. Uncheck all other repositories.

OS	Base URL
Red Hat 5	
<input checked="" type="checkbox"/> CentOS 5	<input type="text" value="http://public-repo-1.hortonworks.com/HDP/centos5/2.x/"/>
<input type="checkbox"/> Oracle Linux 5	
Red Hat 6	
<input checked="" type="checkbox"/> CentOS 6	<input type="text" value="http://public-repo-1.hortonworks.com/HDP/centos6/2.x/"/>
<input type="checkbox"/> Oracle Linux 6	
<input checked="" type="checkbox"/> SLES 11	<input type="text" value="http://public-repo-1.hortonworks.com/HDP/suse11/2.x/u"/>
<input type="checkbox"/> SUSE 11	

3.4. Install Options

In order to build up the cluster, the install wizard needs to know general information about how you want to set it up. You need to supply the FQDN of each of your hosts. The wizard also needs to access the private key file you created in [Set Up Password-less SSH](#). It uses these to locate all the hosts in the system and to access and interact with them securely.

1. Use the **Target Hosts** text box to enter your list of host names, one per line. You can use ranges inside brackets to indicate larger sets of hosts. For example, for host01.domain through host10.domain use `host[01-10].domain`



Note

If you are deploying on EC2, use the **internal Private DNS** hostnames.

2. If you want to let Ambari automatically install the Ambari Agent on all your hosts using SSH, select **Provide your SSH Private Key** and either use the **Choose File** button in the **Host Registration Information** section to find the private key file that matches the public key you installed earlier on all your hosts or cut and paste the key into the text box manually.



Note

If you are using IE 9, the **Choose File** button may not appear. Use the text box to cut and paste your private key manually.

Fill in the username for the SSH key you have selected. If you do not want to use `root`, you must provide the username for an account that can execute `sudo` without entering a password.

3. If you do not want Ambari to automatically install the Ambari Agents, select **Perform manual registration**. See [Appendix: Installing Ambari Agents Manually](#) for more information.
4. Click the **Register and Confirm** button to continue.

3.5. Confirm Hosts

This screen lets you confirm that Ambari has located the correct hosts for your cluster and to check those hosts to make sure they have the correct directories, packages, and processes to continue the install.

If any hosts were selected in error, you can remove them by selecting the appropriate checkboxes and clicking the grey **Remove Selected** button. To remove a single host, click the small white **Remove** button in the Action column.

At the bottom of the screen, you may notice a yellow box that indicates some warnings were encountered during the check process. For example, your host may have already had a copy of `wget` or `curl`. Click **Click here to see the warnings** to see a list of what was checked and what caused the warning. On the same page you can get access to a python script that can help you clear any issues you may encounter and let you run **Rerun Checks**.

When you are satisfied with the list of hosts, click **Next**.

3.6. Choose Services

Hortonworks Data Platform is made up of a number of services. You must at a minimum install HDFS, but you can decide which of the other services you want to install. See [Understand the Basics](#) for more information on your options.

1. Select **all** to preselect all items or **minimum** to preselect only HDFS.
2. Use the checkboxes to unselect (if you have used **all**) or select (if you have used **minimum**) to arrive at your desired list of components.



Note

If you want to use Ambari for monitoring your cluster, make sure you select **Nagios** and **Ganglia**. If you do not select them, you get a warning popup when you finish this section. If you are using other monitoring tools, you can ignore the warning.

3. When you have made your selections, click **Next**.

3.7. Assign Masters

The Ambari install wizard attempts to assign the master nodes for various services you have selected to appropriate hosts in your cluster. The right column shows the current service

assignments by host, with the hostname and its number of CPU cores and amount of RAM indicated.

1. To change locations, click the dropdown list next to the service in the left column and select the appropriate host.
2. To remove a ZooKeeper instance, click the green minus icon next to the host address you want to remove.
3. When you are satisfied with the assignments, click the **Next** button.

3.8. Assign Slaves and Clients

The Ambari install wizard attempts to assign the slave components (DataNodes, NodeManagers, and RegionServers) to appropriate hosts in your cluster. It also attempts to select hosts for installing the appropriate set of clients.

1. Use **all** or **none** to select all of the hosts in the column or none of the hosts, respectively.

If a host has a red asterisk next to it, that host is also running one or more master components. Hover your mouse over the asterisk to see which master components are on that host.

2. Fine-tune your selections by using the checkboxes next to specific hosts.



Note

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 attempt to start it on the same host as the Ambari server, however, you need to start it with the `-p` option, as its default port is 8080 and that conflicts with the Ambari Web default port.

```
/usr/lib/hbase/bin/hbase-daemon.sh start rest -p  
<custom_port_number>
```

3. When you are satisfied with your assignments, click the **Next** button.

3.9. Customize Services

The **Customize Services** screen presents you with a set of tabs that let you manage configuration settings for Hadoop components. The wizard attempts to set reasonable defaults for each of the options here, but you can use this set of tabs to tweak those settings. and you are strongly encouraged to do so, as your requirements may be slightly different. Pay particular attention to the directories suggested by the installer.



Note

In **HDFS Services Configs General**, make sure to enter an integer value, in bytes, that sets the HDFS maximum edit log size for checkpointing. A typical value is 500000000.

Hover over each of the properties to see a brief description of what it does. The number of tabs you see is based on the type of installation you have decided to do. In a complete installation there are ten groups of configuration properties and other related options, such as database settings for Hive/HCat and Oozie, and admin name/password and alert email for Nagios.

The install wizard sets reasonable defaults for all properties except for those related to databases in the Hive and the Oozie tabs, and two related properties in the Nagios tab. These four are marked in red and are the only ones you *must* set yourself.



Note

If you decide to use an existing database instance for Hive/HCatalog or for Oozie, you must have completed the preparations described in [Using Non-Default Databases](#) prior to running the install wizard.

Click the name of the group in each tab to expand and collapse the display.

3.9.1. Service Users and Groups

The individual services in Hadoop are each run under the ownership of a corresponding Unix account. These accounts are known as service users. These service users belong to a special Unix group. In addition there is a special service user for running smoke tests on components during installation and on-demand using the Management Header in the **Services** View of the Ambari Web GUI. Any of these users and groups can be customized using the **Misc** tab of the **Customize Services** step.

If you choose to customize names, Ambari checks to see if these custom accounts already exist. If they do not exist, Ambari creates them. The default accounts are always created during installation whether or not custom accounts are specified. These default accounts are not used and can be removed post-install.



Note

All new service user accounts, and any existing user accounts used as service users, must have a UID ≥ 1000 .

Table 3.1. Service Users

Service	Component	Default User Account
HDFS	NameNode	hdfs
	SecondaryNameNode	
	DataNode	
MapReduce2	HistoryServer	mapred
YARN	NodeManager	yarn
	ResourceManager	
Hive	Hive Metastore	hive
	HiveServer2	
HCat	HCatalog Server	hcat
WebHCat	WebHCat Server	hcat

Service	Component	Default User Account
Oozie	Oozie Server	oozie
HBase	MasterServer	hbase
	RegionServer	
ZooKeeper	ZooKeeper	zookeeper
Ganglia	Ganglia Server	nobody
	Ganglia Collectors	
Nagios	Nagios Server	nagios ^a
Smoke Test ^b	All	ambari-qa

^aIf you plan to use an existing user account named “nagios”, that “nagios” account must either be in a group named “nagios” or you must customize the Nagios Group.

^bThe Smoke Test user performs smoke tests against cluster services as part of the install process. It also can perform these on-demand from the Ambari Web GUI.

Table 3.2. Service Group

Service	Components	Default Group Account
All	All	hadoop
Nagios	Nagios Server	nagios
Ganglia	Ganglia Server	nobody
	Ganglia Monitor	

3.9.2. Properties That Depend on Service Usernames/Groups

Some properties must be set to match specific service usernames or service groups. If you have set up non-default, customized service usernames for the HDFS or HBase service or the Hadoop group name, you must edit the following properties:

Table 3.3. HDFS Settings: Advanced

Property Name	Value
dfs.permissions.superusergroup	The same as the HDFS username. The default is "hdfs"
dfs.cluster.administrators	A single space followed by the HDFS username.
dfs.block.local-path-access.user	The HBase username. The default is "hbase".

Table 3.4. MapReduce Settings: Advanced

Property Name	Value
mapreduce.cluster.administrators	A single space followed by the Hadoop group name.

3.10. Review

The assignments you have made are displayed. Check to make sure everything is correct. If you need to make changes, use the left navigation bar to return to the appropriate screen.

To print your information for later reference, click **Print**.

When you are satisfied with your choices, click the **Deploy** button.

3.11. Install, Start and Test

The progress of the install is shown on the screen. Each component is installed and started and a simple test is run on the component. You are given an overall status on the process in the progress bar at the top of the screen and a host by host status in the main section.

To see specific information on what tasks have been completed per host, click the link in the **Message** column for the appropriate host. In the **Tasks** pop-up, click the individual task to see the related log files. You can select filter conditions by using the **Show** dropdown list. To see a larger version of the log contents, click the **Open** icon or to copy the contents to the clipboard, use the **Copy** icon.

Depending on which components you are installing, the entire process may take 40 or more minutes. Please be patient.

When **Successfully installed and started the services** appears, click **Next**.

3.12. Summary

The Summary page gives you a summary of the accomplished tasks. Click **Complete**. You are taken to the Ambari Web GUI.

4. Hadoop 2.x - Troubleshooting Ambari Deployments

The following information can help you troubleshoot issues you may run into with your Ambari-based installation.

4.1. Review 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 a specific number, for example N. Pay particular attention to the following three files:

- site-N.pp - the puppet file corresponding to a specific task.
- output-N.txt - output from puppet file execution.
- errors-N.txt - error messages.

4.2. Quick Checks

- Make sure all the appropriate services are running. If you have access to Ambari Web, use the **Services View** to check the status of each component. If you do not have access to Manage Services, you must start and stop the services manually.
- If the first HDFS `put` command fails to replicate the block, the clocks in the nodes may not be synchronized. Make sure that Network Time Protocol (NTP) is enabled for your cluster.
- If HBase does not start, check if its slaves are running on 64-bit JVMs. Ambari requires that all hosts must run on 64-bit machines.
- Make sure `umask` is set to 0022.
- Make sure the HCatalog host can access the MySQL server. From a shell try:

```
mysql -h $FQDN_for_MySQL_server -u $FQDN_for_HCatalog_Server -p
```

You will need to provide the password you set up for Hive/HCatalog during the installation process.

- Make sure MySQL is running. By default, MySQL server does not start automatically on reboot.

To set auto-start on boot, from a shell, type:

```
chkconfig --level 35 mysql on
```

To then start the service manually from a shell, type:

```
service mysqld start
```

4.3. Specific Issues

The following are common issues you might encounter.

4.3.1. Problem: Browser crashed before Install Wizard completed

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

4.3.1.1. Solution

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

- The browser closes prior to hitting 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 the **Deploy** button has launched the **Install, Start, and Test** screen

Re-launch the same browser and continue the process or use a different browser and re-login. You are returned to the **Install, Start, and Test** screen.

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

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

4.3.2.1. 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 nav 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** and proceed to the Monitoring **Dashboard**.
 2. Use the **Services View** to make your changes.
 3. Re-start the service using the **Mangement Header**.
- 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 the browser's cache.
 4. Re-run the entire Install Wizard.

4.3.3. Problem: “Unable to create new native thread” exceptions in HDFS DataNode logs or those of any system daemon

If your `nproc` limit is incorrectly configured, the smoke tests fail and you see an error similar to this in the DataNode logs:

```
INFO org.apache.hadoop.hdfs.DFSClient: Exception
increaseBlockOutputStream java.io.EOFException
INFO org.apache.hadoop.hdfs.DFSClient: Abandoning block
blk_-6935524980745310745_139190
```

4.3.3.1. Solution:

In certain recent Linux distributions (like RHEL/Centos/Oracle Linux 6.x), the default value of `nproc` is lower than the value required if you are deploying the HBase service. To change this value:

1. Using a text editor, open `/etc/security/limits.d/90-nproc.conf` and change the `nproc` limit to approximately 32000. For more information, see [ulimit and nproc recommendations for HBase servers](#).

2. Restart the HBase server.

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

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

4.3.4.1. Solution:

You may have incompatible versions of some software components in your environment. Check the list in [Check Existing Installs](#) and make any necessary changes. Also make sure you are running a [Supported Operating System](#)

4.3.5. 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
```

4.3.5.1. Solution:

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

Restart the installation process.

4.3.6. Problem: The HCatalog Daemon Metastore Smoke Test Fails

If the HCatalog smoke test fails, this is displayed in your console:

```
Metastore startup failed, see /var/log/hcatalog/hcat.err
```

4.3.6.1. Solution:

1. Log into the HCatalog node in your cluster
2. Open `/var/log/hcatalog/hcat.err` or `/var/log/hive/hive.log` (one of the two will exist depending on the installation) with a text editor
3. In the file, see if there is a `MySQL Unknown Host Exception` like this:

```
at java.lang.reflect.Method.invoke (Method.java:597)
at org.apache.hadoop.util.Runjar.main (runjar.java:156)
Caused by: java.net.UnknownHostException:mysql.host.com
at java.net.InetAddress.getAllByName (InetAddress.java:1157)
```

This exception can be thrown if you are using a previously existing MySQL instance and you have incorrectly identified the hostname during the installation process. When you do the reinstall, make sure this name is correct.

4. In the file, see if there is an `ERROR Failed initializing database` entry like this:

```
11/12/29 20:52:04 ERROR DataNucleus.Plugin: Bundle
org.eclipse.jdt.core required
11/12/29 20:52:04 ERROR DataStore.Schema: Failed initialising
database
```

This exception can be thrown if you are using a previously existing MySQL instance and you have incorrectly identified the username/password during the installation process. It can also occur when the user you specify does not have adequate privileges on the database. When you do the reinstall, make sure this username/password is correct and that the user has adequate privilege.

5. Restart the installation process.

4.3.7. Problem: MySQL and Nagios fail to install on RightScale CentOS 5 images on EC2

When using a RightScale CentOS 5 AMI on Amazon EC2, in certain cases MySQL and Nagios will fail to install. The MySQL failure is due to a conflict with the pre-installed MySQL and the use of the RightScale EPEL repository (error "Could not find package mysql-server"). Nagios fails to install due to conflicts with the RightScale php-common library.

4.3.7.1. Solution:

On the machines that will host MySQL and Nagios as part of your Hadoop cluster, perform the following:

1. Remove the existing MySQL server

```
yum erase MySQL-server-community
```

2. Install MySQL server with a disabled RightScale EPEL repository

```
yum install mysql-server --disable-repo=rightscales-epel
```

3. Remove the php-common library

```
yum erase php-common-5.2.4-RightScale.x86
```

4.3.8. Problem: Trouble starting Ambari on system reboot

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

4.3.8.1. 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
```

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

Charts appear incorrectly or not at all despite being available in the native Ganglia interface or Host health status is displayed incorrectly.

4.3.9.1. 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.

4.3.10. 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 bug that causes this crash.

4.3.10.1. Solution:

Upgrade to Python version 2.6.8-0.15.1

4.3.11. 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.

4.3.11.1. 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>
```

4.3.12. 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.

4.3.12.1. 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
```

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

When a cluster is first started, some graphs, like **Services View -> HDFS** and **Services View -> MapReduce**, do not plot a complete hour of data, instead showing data only for the length of time the service has been running. Other graphs display the run of a complete hour.

4.3.13.1. Solution

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

4.3.14. Problem: After performing a cluster install the Nagios server is not started

The Nagios server is not started after a cluster install and you are unable to manage it from Ambari Web.

4.3.14.1. Solution

1. Log into the Nagios server host.
2. Confirm that the Nagios server is not running. From a shell:

```
ps -ef | grep nagios
```

You should not see a Nagios process running.

3. Start the Nagios process manually. From a shell:

```
service nagios start
```

4. The server starts. You should be able to see that started state reflected in Ambari Web. You can now manage (start/stop) Nagios from Ambari Web.

4.3.15. 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`.

4.3.15.1. Solution

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

4.3.16. Problem: Updated configuration changes are not pushed to client/gateway nodes

Currently configuration changes are only pushed to daemon running nodes, so any changes are not automatically pushed to client only nodes such as gateway nodes.

4.3.16.1. Solution

Copy the files to the client nodes manually.

5. Appendix: Upgrading Ambari Server to 1.4.4

This procedure upgrades Ambari Server from version 1.2.5 and above to 1.4.4. If your current Ambari Server version is 1.2.4 or below, you must [upgrade the Ambari Server version to 1.2.5](#) before upgrading to version 1.4.4. Upgrading the Ambari Server version does not change the underlying Hadoop Stack.



Note

You must know the location of the Nagios server for Step 9. Use the **Services View-> Summary** panel to locate the host on which it is running.

1. Stop the Ambari Server and all Ambari Agents. From the Ambari Server host:

```
ambari-server stop
```

From each Ambari Agent host:

```
ambari-agent stop
```

2. Get the new Ambari bits. Use `wget` to fetch the repository file and replace the old repo file with the new repo file on every host.

- Fetch the new repo file:

For RHEL/CentOS 5/Oracle Linux 5

```
wget http://public-repo-1.hortonworks.com/ambari/centos5/1.x/updates/1.4.4.23/ambari.repo
```

For RHEL/CentOS 6/Oracle Linux 6

```
wget http://public-repo-1.hortonworks.com/ambari/centos6/1.x/updates/1.4.4.23/ambari.repo
```

For SLES 11

```
wget http://public-repo-1.hortonworks.com/ambari/suse11/1.x/updates/1.4.4.23/ambari.repo
```



Important

Check your current directory before you download the new repo file to make sure that there are no previous versions of the file. If you do not, and a previous version exists, the new download will be saved with a numeric extension such as `ambari.repo.1`. Make sure that the version you copy is the new version.

- Replace the old repo file with the new repo file.

For RHEL/CentOS 5/Oracle Linux 5

```
cp ambari.repo /etc/yum.repos.d/ambari.repo
```

For RHEL/CentOS 6/Oracle Linux 6

```
cp ambari.repo /etc/yum.repos.d/ambari.repo
```

For SLES 11

```
cp ambari.repo /etc/zypp/repos.d/ambari.repo
```



Note

If your cluster does not have access to the Internet, set up a local repository with this data before you continue. See [Configure the Local Repositories](#) for more information.

3. Upgrade Ambari Server.



Note

Ambari Server no longer automatically turns `iptables` off. Check your installation setup to make sure that you are not relying on this function. After you have upgraded the server you must either disable `iptables` manually or make sure that you have all the appropriate ports available. For more information on the ports that must be open and available, see [Configuring Ports for Hadoop 2.x](#)

From the Ambari Server host:

- RHEL/CentOS

```
yum clean all
yum upgrade ambari-server ambari-log4j
```

- SLES

```
zypper clean
zypper up ambari-server ambari-log4j
```

4. Check for upgrade success:

- As the process runs, the console should produce output similar, although not identical, to this:

```
Setting up Upgrade Process
Resolving Dependencies
--> Running transaction check
---> Package ambari-agent.x86_64 0:1.2.2.3-1 will be updated
---> Package ambari-agent.x86_64 0:1.2.2.4-1 will be updated ...
---> Package ambari-agent.x86_64 0:1.2.2.5-1 will be an update ...
```

After the process is complete, check each host to make sure the new 1.4.4 files have been installed:

```
rpm -qa | grep ambari
```

- If the upgrade fails, the console should produce output similar to this:

```
Setting up Upgrade Process
No Packages marked for Update
```

5. Check to see if you have a folder named `/etc/ambari-server/conf.save`. If you do, rename it back:

```
mv /etc/ambari-server/conf.save /etc/ambari-server/conf
```

6. Upgrade the Ambari Server schema. From the Ambari Server host:

```
ambari-server upgrade
```

7. Upgrade the Ambari Agent on all hosts. From each Ambari Agent host:

- RHEL/CentOS

```
yum upgrade ambari-agent ambari-log4j
```

- SLES

```
zypper up ambari-agent ambari-log4j
```



Note

If you get a warning that begins "There are some running programs that use files deleted by recent upgrade" you can ignore it.

8. Check to see if you have a file named `/etc/ambari-agent/conf.save` on each Agent host. If you do, rename it back. On each Agent host:

```
mv /etc/ambari-agent/conf.save /etc/ambari-agent/conf
```

9. Upgrade the Nagios addons package. On the Nagios host:

- RHEL/CentOS

```
yum upgrade hdp_mon_nagios_addons
```

- SLES

```
zypper up hdp_mon_nagios_addons
```

10. Start the Server and the Agents on all hosts. From the Server host:

```
ambari-server start
```

From each Agent host:

```
ambari-agent start
```

11. Open **Ambari Web**. Point your browser to `http://{your.ambari.server}:8080`



Important

You need to refresh your browser so that it loads the new version of the code. Hold the Shift key down while clicking the refresh button on the

browser. If you have problems, clear your browser cache manually and restart Ambari Server.

Use the Admin name and password you have set up to log in.

12.Re-start Nagios and Ganglia services. In **Ambari Web**.

- a. Go to the **Services View** and select each service.
- b. Use the **Management Header** to stop and re-start the service.

6. Appendix: Upgrading the HDP Stack from 1.3.2 or later to 2.0.6

The stack is the coordinated set of Hadoop components that you have installed. If you have a current instance of the 1.3.2 or later stack that was installed and managed by Ambari that you want to upgrade to the new 2.0.6 version of the stack and to also upgrade to the 1.4.4 version of Ambari Server and Agents, use the following instructions. This insures that the upgraded stack can still be managed by Ambari.



Note

If you have turned on Security (Kerberos) for your installation, you should turn it off before the upgrade. On Ambari Web->Admin view->Security-> click **Disable Security**. You can re-enable after the upgrade.

If you are upgrading from any other 1.x version of the stack, you must upgrade to 1.3.2 or later before you can upgrade to 2.0.6. See [Upgrading the HDP Stack to 1.3.3](#) for more information. Upgrades from previous versions of 2.x are not supported.



Note

If you have already upgraded to Ambari Server 1.4.4 and just want to upgrade the HDP stack, you can skip [Section 2](#) and [Section 3](#).

6.1. Preparing for the Upgrade

Use the following steps to prepare your system for the upgrade.

1. If you are upgrading Ambari as well as the stack, you must know the location of the Nagios servers for that process. Use the **Services->Nagios-> Summary** panel to locate the hosts on which they are running.
2. Use the **Services** view on the **Ambari Web** UI to stop all services, including all clients, running on HDFS. Do **not** stop HDFS yet.
3. Finalize any prior upgrade if you have not done so already.

```
su $HDFSUSER
hadoop namenode -finalize
```

4. Create the following logs and other files.

Because the upgrade to 2.0.6 includes a version upgrade of HDFS, creating these logs allows you to check the integrity of the file system post upgrade.

- a. Run `fsck` with the following flags and send the results to a log. The resulting file contains a complete block map of the file system. You use this log later to confirm the upgrade.

```
su $HDFS_USER
```



```
hadoop fsck / -files -blocks -locations > /tmp/dfs-old-fsck-1.log
```

where `$HDFS_USER` is the HDFS Service user (by default, `hdfs`).

- b. Capture the complete namespace of the filesystem. (The following command does a recursive listing of the root file system.)

```
su $HDFS_USER
hadoop dfs -lsr / > /tmp/dfs-old-lsr-1.log
```

where `$HDFS_USER` is the HDFS Service user (by default, `hdfs`).

- c. Create a list of all the DataNodes in the cluster.

```
su $HDFS_USER
hadoop dfsadmin -report > /tmp/dfs-old-report-1.log
```

where `$HDFS_USER` is the HDFS Service user (by default, `hdfs`).

- d. Optional: copy all or unrecoverable only data stored in HDFS to a local file system or to a backup instance of HDFS.
 - e. Optional: create the logs again and check to make sure the results are identical.
5. Save the namespace. You must be the HDFS service user to do this and you must put the cluster in Safe Mode.



Important

This is a **critical** step. If you do not do this step before you do the upgrade, the NameNode will **not** start afterwards.

```
su $HDFS_USER
hadoop dfsadmin -safemode enter
hadoop dfsadmin -saveNamespace
```

6. Copy the following checkpoint files into a backup directory. You can find the directory by using the **Services View** in the UI. Select the **HDFS** service, the **Configs** tab, in the Namenode section, look up the property **NameNode Directories**. It will be on your NameNode host.
 - `dfs.name.dir/edits`
 - `dfs.name.dir/image/fsimage`
 - `dfs.name.dir/current/fsimage`
7. On the JobTracker host, copy `/etc/hadoop/conf` to a backup directory.



Note

If you have deployed a custom version of `capacity-scheduler.xml` and `mapred-queue-acls.xml`, after the upgrade you will need to use Ambari Web to edit the default Capacity Scheduler. Select Services view -> **YARN** -> **Configs** -> **Scheduler** -> **Capacity Scheduler**.



Important

Fair Scheduler is not supported for use with HDP 2.x.

8. Store the layoutVersion for the NameNode. Make a copy of the file at `$dfs.name.dir/current/VERSION` where `$dfs.name.dir` is the value of the config parameter `NameNode directories`. This file will be used later to verify that the layout version is upgraded.
9. Stop HDFS. Make sure all services in the cluster are completely stopped.
10. If you are upgrading Hive, back up the Hive database.
11. Stop Ambari Server. On the Server host:

```
ambari-server stop
```

12. Stop Ambari Agents. On each host:

```
ambari-agent stop
```

6.2. Setting Up the Ambari Repository

This process prepares the updated repository.

1. Check to see if you have a `conf.save` directory for Ambari server and agents. If you do, move them to a back-up location:

```
mv /etc/ambari-server/conf.save/ /etc/ambari-server/conf.save.bak
```

```
mv /etc/ambari-agent/conf.save/ /etc/ambari-agent/conf.save.bak
```

2. Get the new Ambari bits. Use `wget` to fetch the repository file and replace the old repo file with the new repo file on every host.



Important

Check your current directory before you download the new repo file to make sure that there are no previous versions of the file. If you do not, and a previous version exists, the new download is saved with a numeric extension such as `ambari.repo.1`. Make sure that the version you copy is the new version.

- For RHEL/CentOS/Oracle Linux 5

```
wget http://public-repo-1.hortonworks.com/ambari/centos5/1.x/updates/1.4.4.23/ambari.repo
cp ambari.repo /etc/yum.repos.d/ambari.repo
```

- For RHEL/CentOS/Oracle Linux 6

```
wget http://public-repo-1.hortonworks.com/ambari/centos6/1.x/updates/1.4.4.23/ambari.repo
cp ambari.repo /etc/yum.repos.d/ambari.repo
```

- For SLES 11

```
wget http://public-repo-1.hortonworks.com/ambari/suse11/1.x/updates/1.4.4.23/ambari.repo
cp ambari.repo /etc/zypp/repos.d/ambari.repo
```



Note

If your cluster does not have access to the Internet, you need to set up a local repository with this data before you continue. See [Configure the Local Repositories](#) for more information.

6.3. Upgrading to Ambari 1.4.4

This process upgrades Ambari Server, Ambari Agents, Ganglia, and Nagios.



Note

Ambari Server no longer automatically turns `iptables` off. Check your installation setup to make sure that you are not relying on this function. After you have upgraded the server you must either disable `iptables` manually or make sure that you have all the appropriate ports available. For more information on the ports that must be open and available, see [Configuring Ports for Hadoop 2.x](#)

1. Upgrade Ambari Server. From the Ambari Server host:

- RHEL/CentOS/Oracle Linux

```
yum clean all
yum upgrade ambari-server ambari-log4j
```

- SLES

```
zypper clean
zypper up ambari-server ambari-log4j
```

2. Check for upgrade success:

- As the process runs, the console should produce output similar, although not identical, to this:

```
Setting up Upgrade Process
Resolving Dependencies
--> Running transaction check
---> Package ambari-agent.x86_64 0:1.2.2.3-1 will be updated
---> Package ambari-agent.x86_64 0:1.2.2.4-1 will be updated ...
---> Package ambari-agent.x86_64 0:1.2.2.5-1 will be an update ...
```

After the process is complete, check each host to make sure the new 1.4.4 files have been installed:

```
rpm -qa | grep ambari
```

- If the upgrade fails, the console should produce output similar to this:

```
Setting up Upgrade Process
```

```
No Packages marked for Update
```

3. Check to see if you have a folder named `/etc/ambari-server/conf.save`. If you do, rename it back:

```
mv /etc/ambari-server/conf.save /etc/ambari-server/conf
```

4. Upgrade the Ambari Server schema. From the Ambari Server host:

```
ambari-server upgrade
```

5. Upgrade the Ambari Agent on all hosts. From each Ambari Agent host:

- RHEL/CentOS/Oracle Linux

```
yum upgrade ambari-agent ambari-log4j
```

- SLES

```
zypper up ambari-agent ambari-log4j
```



Note

If you get a warning that begins "There are some running programs that use files deleted by recent upgrade" you can ignore it.

6. Check to see if you have a folder named `/etc/ambari-agent/conf.save` on each Agent host. If you do, rename it back. On each Agent host:

```
mv /etc/ambari-agent/conf.save /etc/ambari-agent/conf
```

7. Upgrade the Nagios addons:

- RHEL/CentOS/Oracle Linux

```
yum upgrade hdp_mon_nagios_addons
```

- SLES

```
zypper up hdp_mon_nagios_addons
```

6.4. Upgrading the Stack

This stack upgrade involves removing the HDP 1.x version of MapReduce and replacing it with the HDP 2.x YARN and MapReduce2 components. This process is somewhat long and complex. To help you, a Python script is provided to automate some of the upgrade steps. It is available at `/var/lib/ambari-server/resources/scripts/UpgradeHelper_HDP2.py` on the Ambari Server host. The script can be executed on any host that can communicate with Ambari Server. It requires Python 2.6 or higher.

6.4.1. Prepare for the Stack Upgrade

1. Make sure that you have saved the namespace. You should have done this [here](#). The upgrade will fail if you do not save the namespace. If you have not saved the namespace yet:

a. Restart Ambari Server and Ambari Agents.

b. Restart HDFS **only**.

c. On the NameNode host:

```
su $HDFS_USER
hadoop dfsadmin -safemode enter
hadoop dfsadmin -saveNamespace
```

d. Stop the HDFS service and wait for it to be fully stopped.

e. Stop the Ambari Server and Ambari Agents.

2. Prepare for the upgrade:

a. Create an "Upgrade Folder", for example `/work/upgrade_hdp_2`, on a host that can communicate with Ambari Server. The Ambari Server host would be a suitable candidate.

b. Copy the upgrade script to the Upgrade Folder. The script is available here: `/var/lib/ambari-server/resources/scripts/UpgradeHelper_HDP2.py` on the Ambari Server host.

c. Make sure that Python is available on the host and that the version is 2.6 or higher:

```
python --version
```



Note

For RHEL/Centos/Oracle Linux 5, you **must** use Python 2.6.

3. Start Ambari Server only. On the Ambari Server host:

```
ambari-server start
```

4. Back up current configuration settings and the component host mappings from MapReduce:

a. Go to the Upgrade Folder.

b. Execute the `backup-configs` action:

```
python UpgradeHelper_HDP2.py ---hostname $HOSTNAME --user $USERNAME --password $PASSWORD --clustername $CLUSTERNAME backup-configs
```

Where

- `$HOSTNAME` is the name of the Ambari Server host
- `$USERNAME` is the admin user for Ambari Server
- `$PASSWORD` is the password for the admin user
- `$CLUSTERNAME` is the name of the cluster

This step produces a set of files named `TYPE_TAG`, where `TYPE` is the configuration type and `TAG` is the tag. These files contain copies of the various configuration settings for the current (pre-upgrade) cluster. You can use these files as a reference later.

- c. Execute the `save-mr-mapping` action:

```
python UpgradeHelper_HDP2.py --hostname $HOSTNAME --user $USERNAME --password $PASSWORD --clustertype $CLUSTERTYPE save-mr-mapping
```

This step produces a file named `mr_mapping` that stores the host level mapping of MapReduce components such as MapReduce JobTracker/TaskTracker/Client.

5. Delete all the MapReduce server components installed on the cluster.

- a. If you are not already there, go to the Upgrade Folder.

- b. Execute the `delete-mr` action.

```
python UpgradeHelper_HDP2.py --hostname $HOSTNAME --user $USERNAME --password $PASSWORD --clustertype $CLUSTERTYPE delete-mr
```

Optionally, execute the delete script with the `-n` option to view, verify, and validate API calls, if necessary.



Note

Running the delete script with the `-n` option exposes API calls but does not remove installed components. Use the `-n` option for validation purposes only.

- c. The script asks you to confirm that you have executed the `save-mr-mapping` action and that you have a file named `mr_mapping` in the Upgrade Folder.

6.4.2. Upgrade the Stack

1. Stop Ambari Server. On the Ambari Server host:

```
ambari-server stop
```

2. Update the stack version in the Server database, depending on if you are using a local repository:



Important

Make sure you **delete the old MapReduce** version **before** you run `upgradestack`.

```
ambari-server upgradestack HDP-2.0.6
```

3. Upgrade the HDP repository on all hosts and replace the old repo file with the new file:



Important

The file you download is named `hdp.repo`. To function properly in the system, it must be named `HDP.repo`. Once you have completed the "mv" of the new repo file to the `repos.d` folder, make sure there is no file named `hdp.repo` anywhere in your `repos.d` folder.

- For RHEL/CentOS/Oracle Linux 5

```
wget http://public-repo-1.hortonworks.com/HDP/centos5/2.x/updates/2.0.6.1/hdp.repo
mv hdp.repo /etc/yum.repos.d/HDP.repo
```

- For RHEL/CentOS/Oracle Linux 6

```
wget http://public-repo-1.hortonworks.com/HDP/centos6/2.x/updates/2.0.6.1/hdp.repo
mv hdp.repo /etc/yum.repos.d/HDP.repo
```

- For SLES 11

```
wget http://public-repo-1.hortonworks.com/HDP/suse11/2.x/updates/2.0.6.1/hdp.repo
mv hdp.repo /etc/zypp/repos.d/HDP.repo
```

4. Upgrade the stack on all Agent hosts. Skip any components your installation does not use:

- For RHEL/CentOS/Oracle Linux

- a. Remove remaining MapReduce components on all hosts:

```
yum erase hadoop-pipes hadoop-sbin hadoop-native
```

- b. Upgrade the following components:

```
yum upgrade "collectd*" "epel-release*" "gccxml*" "pig*" "hadoop*"
"sqoop*" "zookeeper*" "hbase*" "hive*" "hcatalog*" "webhcat-tar*"
"oozie*" hdp_mon_nagios_addons
```

- c. Check to see that the components in that list are upgraded.

```
yum list installed | grep HDP-$old-stack-version-number
```

None of the components from that list should appear in the returned list.

- For SLES

- a. Remove remaining MapReduce components on all hosts:

```
zypper remove hadoop-pipes hadoop-sbin hadoop-native
```

- b. Upgrade the following components:

```
zypper up "collectd*" "epel-release*" "gccxml*" "pig*" "hadoop*"
"sqoop*" "zookeeper*" "hbase*" "hive*" "hcatalog*" "webhcat-tar*"
"oozie*" hdp_mon_nagios_addons
```

- c. To verify that components were upgraded, execute:

```
rpm -qa | grep hadoop, rpm -qa | grep hive and rpm -qa | grep hcatalog
```

- d. If components were not upgraded, upgrade them as follows:

```
yast --update hadoop hcatalog hive
```

6.4.3. Add YARN/MR2 and Update Configurations

1. Start the Ambari Server. On the Server host:

```
ambari-server start
```

2. **If using a local repository**, update the URL for the local repository. On the Server host:

```
curl -H "X-Requested-By: ambari" -X PUT -u admin:$ADMIN_PASSWD http://
/$AMBARI_SERVER_HOST:8080/api/v1/stacks2/HDP/versions/$HDP_VERSION/
operatingSystems/$OS_TYPE/repositories/$HDP_STACK_ID -d '{"Repositories":
{"base_url": "$LOCAL_REPO_URL", "verify_base_url": true}}'
```

For example,

```
curl -H "X-Requested-By: ambari" -X PUT -u admin:admin http://
localhost:8080/api/v1/stacks2/HDP/versions/2.0.6/operatingSystems/centos6/
repositories/HDP-2.0.6 -d '{"Repositories": {"base_url": "http://private-
repo.xyz.com/HDP/centos6/2.x/updates/2.0.6.1", "verify_base_url": true}}'
```



Note

This update verifies the base URL, by default. Optionally, to update this URL without verification, set "verify_base_url" to false.

3. Start each Ambari Agent. On all Agent hosts:

```
ambari-agent start
```

4. After the Server and all Agents are running, log into Ambari Web. Do a hard refresh on your browser to make sure you are displaying the updated GUI. Make sure all hosts are healthy and all services are in Stopped state.

5. Add YARN and MapReduce2 services:

- a. If you are not already there, go to the Upgrade Folder.

- b. Execute the `add-yarn-mr2` action:

```
python UpgradeHelper_HDP2.py --hostname $HOSTNAME --user $USERNAME --
password $PASSWORD --clustername $CLUSTERNAME add-yarn-mr2
```

If desired, you can use the `-n` option to see the API calls as they are being made so that you can verify them.

6. Update the respective configurations:

- a. If you are not already there, go to the Upgrade Folder.

- b. Execute the `update-configs` action:

```
python UpgradeHelper_HDP2.py --hostname $HOSTNAME --user $USERNAME --
password $PASSWORD --clustertype $CLUSTERTYPE update-configs
```

7. Update individual configuration settings as needed. On the Ambari Server, use `/var/lib/ambari-server/resources/scripts/configs.sh` to inspect and update the configuration properties.



Note

`configs.sh` creates temporary files. We recommend that you run `configs.sh` as root or as a user having write permission on the local folder.

- a. Get configuration details:

```
configs.sh get $HOSTNAME $CLUSTERTYPE $CONFIGURATION-TYPE
```

For example:

```
configs.sh get localhost myclustertype global
```

- b. Evaluate each property value returned and modify as needed:

```
configs.sh set $HOSTNAME $CLUSTERTYPE $CONFIGURATION-TYPE "property name"
"new value"
```

For example:

```
configs.sh set localhost myclustertype global yarn_log_dir_prefix "/apps/
logs/yarn"
```

- c. Remove properties that are not needed:

```
configs.sh delete $HOSTNAME $CLUSTERTYPE $CONFIGURATION-TYPE "property
name"
```

For example:

```
configs.sh delete localhost myclustertype global dfs.client-write-packet-
size
```

Table 6.1. Key Properties to Check

Configuration Type	Property	Description
global	yarn_log_dir_prefix	The location for the YARN logs
global	yarn_pid_dir_prefix	The location for the YARN pid files
global	yarn_user	The YARN user
-	-	-
yarn-site	yarn.nodemanager.local-dirs	The location for container logs
yarn-site	yarn.nodemanager.log-dirs	The directories for localized files



Note

Make any necessary modifications **before** starting the services.

- d. Install the YARN and MapReduce2 services:

- i. If you are not already there, go to the Upgrade Folder.
- ii. Execute the `install-yarn-mr2` action:

```
python UpgradeHelper_HDP2.py --hostname $HOSTNAME --user $USERNAME --
password $PASSWORD --clustertype $CLUSTERTYPE install-yarn-mr2
```



Note

This is a two step process. You can use the Ambari Web GUI to monitor the progress. Both steps must be **complete** before you continue to the next step.

6.4.4. Complete the Stack Upgrade

1. Because the file system version has now changed you must start the NameNode manually. On the NameNode host:

```
su -l $HDFS_USER -c "export HADOOP_LIBEXEC_DIR=/usr/lib/hadoop/libexec && /
usr/lib/hadoop/sbin/hadoop-daemon.sh start namenode -upgrade"
```

Depending on the size of your system, this step may take up to 10 minutes.

2. This upgrade can take a long time depending on the number of files you have. You can use `tail` to monitor the log file so that you can track your progress:

```
tail -f /var/log/$HDFS_LOG_FOLDER/hadoop-hdfs-namenode-$HOSTNAME.log
```

Look for lines that confirm the upgrade is complete, one line per name directory, such as `Upgrade of /hadoop/hdfs/namenode is complete`. You can also look for `Registered FSNamesystem State MBean` which follows the upgrade of all name directories.

3. Prepare the NameNode to work with Ambari:
 - a. Open the Ambari Web GUI. If it has been open throughout the process, do a hard reset on your browser to force a reload.
 - b. On the Services view, click **HDFS** to open the HDFS service.
 - c. Click **View Host** to open the NameNode host details page.
 - d. Use the dropdown menu to stop the NameNode.
 - e. On the Services view, restart the HDFS service. Make sure it passes the ServiceCheck. It is now under Ambari's control.
4. After the DataNodes are started, HDFS exits safemode. To monitor the status:

```
sudo su -l $HDFS_USER -c "hdfs dfsadmin -safemode get"
```

Depending on the size of your system, this may take up to 10 minutes or so. When HDFS exits safemode, this is displayed as a response to the command:

```
Safe mode is OFF
```

5. Make sure that the HDFS upgrade was successful. Go through steps 2 and 3 in [Section 9.1](#) to create new versions of the logs and reports. Substitute "new" for "old" in the file names as necessary

6. Compare the old and new versions of the following:

- `dfs-old-fsck-1.log` versus `dfs-new-fsck-1.log`.

The files should be identical unless the `hadoop fsck` reporting format has changed in the new version.

- `dfs-old-lsr-1.log` versus `dfs-new-lsr-1.log`.

The files should be identical unless the the format of `hadoop fs -lsr` reporting or the data structures have changed in the new version.

- `dfs-old-report-1.log` versus `fs-new-report-1.log`

Make sure all DataNodes previously belonging to the cluster are up and running.

7. Use the Ambari Web Services view to start YARN.

8. Use the Ambari Web Services view to start MapReduce2.

9. Upgrade HBase:

a. Make sure that all HBase components - RegionServers and HBase Master - are stopped.

b. Use the Ambari Web Services view, start the ZooKeeper service. Wait until the ZK service is up and running.

c. On the HBase Master host, make these configuration changes:

i. If `dfs.domain.socket.path` isn't set, either set `dfs.domain.socket.path` or set the property `dfs.client.read.shortcircuit` to `false` in `HBASE_CONFDIR/hbase-site.xml`.

ii. In the configuration file, find the value of the `hbase.tmp.dir` property and make sure that the directory exists and is readable and writeable for the HBase service user and group.

```
chown -R $HBASE_USER:$HADOOP_GROUP $HBASE.TMP.DIR
```

iii. Go to the Upgrade Folder and check in the saved global configuration file named `global_<TAG>` for the value of the property `hbase_pid_dir` and `hbase_log_dir`. Make sure that the directories are readable and writeable for the HBase service user and group.

```
chown -R $HBASE_USER:$HADOOP_GROUP $hbase_pid_dir
chown -R $HBASE_USER:$HADOOP_GROUP $hbase_log_dir
```

Do this on **every** host where a RegionServer is installed as well as on the HBase Master host.

- iv. Check for HFiles in V1 format. HBase 0.96.0 discontinues support for HFileV1. Before the actual upgrade, run the following command with HBase0.96 binaries to check if there are HFiles in V1 format while 0.94 cluster is running:

```
hbase upgrade -check
```

HFileV1 was a common format prior to HBase 0.94. You may see output similar to:

```
Tables Processed:
hdfs://localhost:41020/myHBase/.META.
hdfs://localhost:41020/myHBase/usertable
hdfs://localhost:41020/myHBase/TestTable
hdfs://localhost:41020/myHBase/t

Count of HFileV1: 2
HFileV1:
hdfs://localhost:41020/myHBase/usertable/
fa02dac1f38d03577bd0f7e666f12812/family/249450144068442524
hdfs://localhost:41020/myHBase/usertable/
ecdd3eaae2d2fcf8184ac025555bb2af/family/249450144068442512

Count of corrupted files: 1
Corrupted Files:
hdfs://localhost:41020/myHBase/usertable/
fa02dac1f38d03577bd0f7e666f12812/family/1
Count of Regions with HFileV1: 2
Regions to Major Compact:
hdfs://localhost:41020/myHBase/usertable/
fa02dac1f38d03577bd0f7e666f12812
hdfs://localhost:41020/myHBase/usertable/
ecdd3eaae2d2fcf8184ac025555bb2af
```

When you run the upgrade check, if "Count of HFileV1" returns any files, start the hbase shell to use major compaction for regions that have HFileV1 format. For example in the sample output above, you must compact the fa02dac1f38d03577bd0f7e666f12812 and ecdd3eaae2d2fcf8184ac025555bb2af regions.

- v. Upgrade HBase. You must be the HBase service user.

```
sudo su -l $HBASE_USER -c "hbase upgrade -execute"
```

Make sure that the output contains the string "Successfully completed Znode upgrade".

- vi. Use the Services view to start the HBase service. Make sure that Service Check passes.

10. Upgrade Oozie:

- On the Services view, make sure YARN and MapReduce2 are running.
- Make sure that the Oozie service is stopped.
- Upgrade Oozie. You must be the Oozie service user. On the Oozie host:

```
sudo su -l $OOZIE_USER -c "/usr/lib/oozie/bin/ooziedb.sh upgrade -run"
```

Make sure that the output contains the string "Oozie DB has been upgrade to Oozie version '*OOZIE Build Version*'".

- d. Prepare the WAR file:



Note

The Oozie server must be **not** running for this step. If you get the message "ERROR: Stop Oozie first", it means the script still thinks it's running. Check, and if needed, remove the process id (pid) file indicated in the output.

```
/usr/lib/oozie/bin/oozie-setup.sh prepare-war
```

Make sure that the output contains the string "New Oozie WAR file with added".

- e. Modify the following configuration properties in `oozie-site.xml`. On the Ambari Server, use `/var/lib/ambari-server/resources/scripts/configs.sh` to inspect and update the configuration properties as described [here \[57\]](#).

Table 6.2. Properties to Modify

Action	Property Name	Property Value
Add	oozie.service.URIHandlerService.uri.handlers	org.apache.oozie.dependency.FSURIHandler, org.apache.oozie.dependency.HCatURIHandler
Add	oozie.service.coord.push.check.requeue.interval	30000
Add	oozie.services.ext	org.apache.oozie.service.PartitionDependencyManagerService, org.apache.oozie.service.HCatAccessorService
Add/Modify	oozie.service.SchemaService.wfext.schemas	shell-action-0.1.xsd, email-action-0.1.xsd, hive-action-0.2.xsd, sqoop-action-0.2.xsd, ssh-action-0.1.xsd, distcp-action-0.1.xsd, shell-action-0.2.xsd, oozie-sla-0.1.xsd, oozie-sla-0.2.xsd ^a

^aUse this list if you have not modified the default Ambari values. If you have added custom schemas, make sure they exist after the modification. The schemas being added here are `shell-action-0.2.xsd`, `oozie-sla-0.1.xsd`, and `oozie-sla-0.2.xsd`. You can add these to your existing list.

- f. Replace the content of `/user/oozie/share` in HDFS. On the Oozie server host:

- i. Extract the Oozie sharelib into a `tmp` folder.

```
mkdir -p /tmp/oozie_tmp
cp /usr/lib/oozie/oozie-sharelib.tar.gz /tmp/oozie_tmp
cd /tmp/oozie_tmp
tar xzvf oozie-sharelib.tar.gz
```

- ii. Back up the `/user/oozie/share` folder in HDFS and then delete it. If you have any custom files in this folder back them up separately and then add them back after the share folder is updated.

```
su -l $HDFS_USR -c "$hdfs dfs -copyToLocal /user/oozie/share /tmp/oozie_tmp/oozie_share_backup"
su -l $HDFS_USR -c "$hdfs dfs -rm -r /user/oozie/share"
```

- iii. Add the latest share libs that you extracted in step 1. After you have added the files, modify ownership and acl.

```
su -l $HDFS_USR -c "hdfs dfs -copyFromLocal /tmp/oozie_tmp/share /user/oozie/."
su -l $HDFS_USR -c "hdfs dfs -chown -R oozie:hadoop /user/oozie"
su -l $HDFS_USR -c "hdfs dfs -chmod -R 755 /user/oozie"
```

- g. Use the Services view to start the Oozie service. Make sure that ServiceCheck passes for Oozie.

11. Update webhcat.

- a. Modify the webhcat-site config type.

On the Ambari server, use `/var/lib/ambari-server/resources/scripts/configs.sh` to modify configuration properties in `templeton.storage.class`:

```
configs.sh set $HOSTNAME $CLUSTERNAME $CONFIGURATION-TYPE $PROPERTY-NAME
$PROPERTY-VALUE
For example: configs.sh set <yourhostname> <yourclustername> webhcat-
site "templeton.storage.class" "org.apache.hive.hcatalog.templeton.tool.
ZooKeeperStorage"
```

- b. Update the Pig and Hive tar bundles, by updating the following files:

- `/apps/webhcat/pig.tar.gz`
- `/apps/webhcat/hive.tar.gz`



Note

You will find these files on a host where webhcat is installed.

For example, to update a *.tar.gz file:

- i. Move the file to a local directory.

```
su -l $HCAT_USR -c "hadoop --config /etc/hadoop/conf fs -copyToLocal /
apps/webhcat/*.tar.gz ${local_backup_dir}"
```

- ii. Remove the old file.

```
su -l $HCAT_USR -c "hadoop --config /etc/hadoop/conf fs -rm /apps/
webhcat/*.tar.gz"
```

- iii. Copy the new file.

```
su -l $HCAT_USR -c "hadoop --config /etc/hadoop/conf fs -
copyFromLocal /user/share/HDP-webhcat/*.tar.gz /apps/webhcat"
```

- c. Update `/app/webhcat/hadoop-streaming.jar` file.

- i. Move the file to a local directory.

```
su -l $HCAT_USR -c "hadoop --config /etc/hadoop/conf fs -copyToLocal /
apps/webhcat/hadoop-streaming*.jar ${local_backup_dir}"
```

- ii. Remove the old file.

```
su -l $HCAT_USR -c "hadoop --config /etc/hadoop/conf fs -rm /apps/
webhcat/hadoop-streaming*.jar"
```

- iii. Copy the new hadoop-streaming.jar file.

```
su -l $HCAT_USR -c "hadoop --config /etc/hadoop/conf fs -
copyFromLocal /user/lib/hadoop-mapreduce/hadoop-streaming*.jar /apps/
webhcat"
```

12. Make sure Ganglia no longer attempts to monitor JobTracker.

- a. Make sure Ganglia is stopped.
- b. Log into the host where JobTracker was installed (and where ResourceManager is installed after the upgrade).
- c. Backup the folder `/etc/ganglia/hdp/HDPJobTracker`.
- d. Remove the folder `/etc/ganglia/hdp/HDPJobTracker`.
- e. Remove the folder `$ganglia_runtime_dir/HDPJobTracker`.



Note

For the value of `$ganglia_runtime_dir`, in the Upgrade Folder, check the saved global configuration file `global_<$TAG>`.

13. Use the Services view to start the remaining services back up.

14. The upgrade is now fully functional but not yet finalized. Using the `finalize` command removes the previous version of the NameNode and DataNode's storage directories.



Important

After the upgrade is finalized, the system cannot be rolled back. Usually this step is not taken until a thorough testing of the upgrade has been performed.

The upgrade must be finalized before another upgrade can be performed.



Note

Directories used by Hadoop 1 services set in `/etc/hadoop/conf/taskcontroller.cfg` are not automatically deleted after upgrade. Administrators can choose to delete these directories after the upgrade.

To finalize the upgrade:

```
sudo su -l $HDFS_USER -c "hadoop dfsadmin -finalizeUpgrade"
```

where `$HDFS_USER` is the HDFS Service user (by default, `hdfs`).

7. Appendix: Hadoop 2.x - Configuring Ports

The tables below specify which ports must be opened for which ecosystem components to communicate with each other. Make sure the appropriate ports are opened before you install Hadoop.

- [HDFS Ports](#)
- [MapReduce Ports](#)
- [YARN Ports](#)
- [Hive Ports](#)
- [HBase Ports](#)
- [ZooKeeper Ports](#)
- [WebHCat Port](#)
- [Ganglia Ports](#)
- [MySQL Port](#)
- [Ambari Ports](#)
- [Nagios Port](#)

7.1. HDFS Ports

The following table lists the default ports used by the various HDFS services.

Table 7.1. HDFS Ports

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
NameNode WebUI	Master Node hosts (NameNode and any back-up NameNodes)	50070	http	Web UI to look at current status of HDFS, explore file system	Yes (Typically admins, Dev/ Support teams)	<code>dfs.namnode.http-address</code>
		50470	https	Secure http service		<code>dfs.namenode.https-address</code>
NameNode metadata service	Master Node hosts (NameNode and any back-up NameNodes)	8020/9000	IPC	File system metadata operations	Yes (All clients who directly need to interact with the HDFS)	Embedded in URI specified by <code>fs.defaultFS</code>
DataNode	All Slave Node hosts	50075	http	DataNode WebUI to	Yes (Typically admins, Dev/ Support teams)	<code>dfs.datanode.http.address</code>

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
				access the status, logs etc.		
		50475	https	Secure http service		dfs.datanode.https.address
		50010		Data transfer		dfs.datanode.address
		0.0.0.0:8010	IPC	Metadata operations	No	dfs.datanode.ipc.address
Secondary NameNode	Secondary NameNode and any backup Secondary NameNode hosts	50090	http	Checkpoint for NameNode metadata	No	dfs.namenode.secondary.http.address

7.2. MapReduce Ports

The following table lists the default port used by the History Server WebUI.

Table 7.2. MapReduce Ports

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
History Server WebUI		19888	http	Web UI for Job History	Yes	mapreduce.jobhistory.webapp.address

7.3. YARN Ports

The following table lists the default ports used by YARN.

Table 7.3. YARN Ports

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
YARN Resource Manager		8025				yarn.resourcemanager.resource-tracker.address
YARN RM Admin		8141		The address of the RM admin interface		yarn.resourcemanager.admin.address
Container Manager		0.0.0.0:45454		The address of the container manager in the NameNode		yarn.nodemanager.address
Applications Manager		8050		The address of the applications manager interface in the RM.		yarn.resourcemanager.address

7.4. Hive Ports

The following table lists the default ports used by the various Hive services.



Note

Neither of these services is used in a standard HDP installation.

Table 7.4. Hive Ports

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
Hive Server2	Hive Server machine (Usually a utility machine)	10000	thrift	Service for programatically (Thrift/JDBC) connecting to Hive	Yes (Clients who need to connect to Hive either programatically or through UI SQL tools that use JDBC)	ENV Variable HIVE_PORT
Hive Metastore		9083	thrift	Service for accesing metadata about Hive tables and partitions.*	Yes (Clients that run Hive, Pig and potentially M/R jobs that use HCatalog)	hive.metastore.uris

* To change the metastore port, use this hive command: `hive --service metastore -p port_number`

7.5. HBase Ports

The following table lists the default ports used by the various HBase services.

Table 7.5. HBase Ports

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
HMaster	Master Node hosts (HBase Master Node and any back-up HBase Master node)	60000			Yes	hbase.master.port
HMaster Info Web UI	Master Node hosts (HBase master Node and back up HBase Master node if any)	60010	http	The port for the HBase-Master web UI. Set to -1 if you do not want the info server to run.	Yes	hbase.master.info.port
Region Server	All Slave Node hosts	60020			Yes (Typically admins, dev/support teams)	hbase.regionserver.port
Region Server	All Slave Node hosts	60030	http		Yes (Typically admins, dev/support teams)	hbase.regionserver.info.port
HBase REST Server (optional)	All REST Servers	8080	http	The port used by HBase Rest Servers. REST servers are optional, and	Yes	hbase.rest.port

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
				not installed by default		
HBase REST Server Web UI(optional)	All REST Servers	8085	http	The port used by HBase Rest Servers web UI. REST servers are optional, and not installed by default	Yes (Typically admins, dev/support teams)	hbase.rest.info.port
HBase Thrift Server (optional)	All Thrift Servers	9090		The port used by HBase Thrift Servers. Thrift servers are optional, and not installed by default	Yes	hbase.zookeeper.property.clientPort
HBase Thrift Server Web UI (optional)	All Thrift Servers	9095		The port used by HBase Thrift Servers web UI. Thrift servers are optional, and not installed by default	Yes (Typically admins, dev/support teams)	hbase.thrift.info.port

7.6. ZooKeeper Ports

The following table lists the default ports used by the various ZooKeeper services.

Table 7.6. HBase Ports

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
ZooKeeper Server	All ZooKeeper Node hosts	2888		Port used by ZooKeeper peers to talk to each other. See here for more information.	No	hbase.zookeeper.peerport
ZooKeeper Server	All ZooKeeper Node hosts	3888		Port used by ZooKeeper peers to talk to each other. See here for more information.	No	hbase.zookeeper.leaderport
ZooKeeper Server	All ZooKeeper Hosts	2181		Property from ZooKeeper's config <code>zoo.cfg</code> . The port at which the clients will connect.	Yes	hbase.zookeeper.property.clientPort

7.7. WebHCat Port

The following table lists the default port used by the WebHCat service.

Table 7.7. WebHCat Port

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
WebHCat Server	Any utility machine	50111	http	Web API on top of HCatalog and other Hadoop services	Yes	templeton.port

7.8. Ganglia Ports

The following table lists the default ports used by the various Ganglia services.

Table 7.8. Ganglia Ports

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
Ganglia Server	Ganglia server host	8660/61/62/63		For metric (gmond) collectors	No	
Ganglia Monitor	All Slave Node hosts	8660		For monitoring (gmond) agents	No	
Ganglia Server	Ganglia server host	8651		For ganglia gmetad		
Ganglia Web	Ganglia server host		http ^a			

^aSee [Optional: Set Up HTTPS for Ganglia](#) for instructions on enabling HTTPS.

7.9. MySQL Port

The following table lists the default port used by the MySQL service.

Table 7.9. MySQL Port

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
MySQL	MySQL database server host	3306				

7.10. Ambari Ports

The following table lists the default ports used by Ambari.

Table 7.10. Ambari Web

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
Ambari Server	Ambari Server host	8080 ^a	http ^b	Interface to Ambari Web and Ambari REST API	No	

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
Ambari Server	Ambari Server host	8440	https	Handshake Port for Ambari Agents to Ambari Server	No	
Ambari Server	Ambari Server host	8441	https	Registration and Heartbeat Port for Ambari Agents to Ambari Server	No	

^aSee [Optional: Change the Ambari Server Port](#) for instructions on changing the default port.

^bSee [Optional: Set Up HTTPS for Ambari Web](#) for instructions on enabling HTTPS.

7.11. Nagios Ports

The following table lists the default port used by Nagios.

Table 7.11. Nagios

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
Nagios Server	Nagios server host	80	http ^a	Nagios Web UI	No	

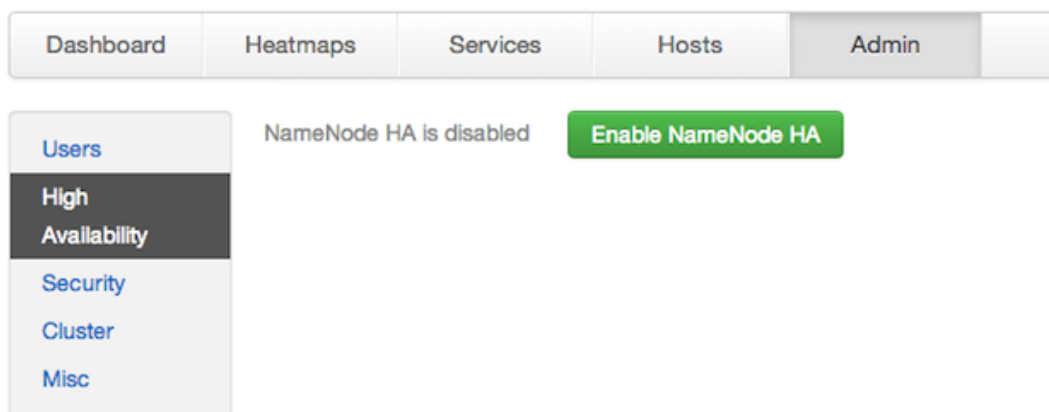
^aSee [Optional: Set Up HTTPS for Nagios](#) for instructions on enabling HTTPS.

8. Appendix: NameNode High Availability

Use these instructions to set up NameNode HA using Ambari Web.

8.1. Setting Up NameNode High Availability

On Ambari Web, go to the Admin view. Select **High Availability** in the left nav bar.



1. Check to make sure you have at least three hosts in your cluster and are running at least three ZooKeeper servers.
2. Click **Enable NameNode HA** and follow the **Enable NameNode HA Wizard**. The wizard describes a set of automated and manual steps you must take to set up NameNode high availability.
3. **Get Started**: This step gives you an overview of the process and allows you to select a Nameservice ID. You use this Nameservice ID instead of the NameNode FQDN once HA has been set up. Click **Next** to proceed.

Enable NameNode HA Wizard

**ENABLE NAMEDNODE
HA WIZARD**

Get Started

Select Hosts

Review

Create Checkpoint

Configure Components

Initialize JournalNodes

Start Components

Initialize Metadata

Finalize HA Setup

Get Started

This wizard will walk you through enabling NameNode HA on your cluster. Once enabled, you will be running a Standby NameNode in addition to the active NameNode. This allows for an Active-Standby NameNode configuration that allows for a failover.

The process to enable HA involves a combination of **automated steps** (that are handled by the wizard) and **manual steps** (that you must perform as instructed by the wizard).

You should plan a cluster maintenance window and prepare for downtime when enabling NameNode HA.

If you have HBase running, please exit this wizard and stop HBase before proceeding.

Nameservice ID:

4. **Select Hosts:** Select a host for the additional NameNode and the JournalNodes. The wizard suggests options, but you can adjust using the dropdown lists. Click **Next** to proceed.

Select Hosts

Select a host that will be running the additional NameNode.
In addition, select the hosts to run JournalNodes, which store NameNode edit logs in a fault tolerant manner.

Current NameNode:

Additional NameNode:

JournalNode:

JournalNode:

JournalNode:

c6401.ambari.apache.org (1.8 GB, 1 cores)

NameNode Nagios Server Ganglia Server

HBase Master ZooKeeper JournalNode

c6402.ambari.apache.org (1.8 GB, 1 cores)

SNameNode History Server

ResourceManager HiveServer2

Hive Metastore WebHCat Server

Oozie Server ZooKeeper JournalNode

NameNode

c6403.ambari.apache.org (1.8 GB, 1 cores)

ZooKeeper JournalNode

5. **Review:** Confirm your host selections and click **Next**.

Review

Confirm your host selections.

Current NameNode: c6401.ambari.apache.org

Secondary NameNode: c6402.ambari.apache.org - TO BE DELETED

Additional NameNode: c6402.ambari.apache.org + TO BE INSTALLED

JournalNode: c6401.ambari.apache.org + TO BE INSTALLED
 c6402.ambari.apache.org + TO BE INSTALLED
 c6403.ambari.apache.org + TO BE INSTALLED

6. **Create Checkpoints:** Follow the instructions in the step. You need to login to your **current** NameNode host to run the commands to put your NameNode into safe mode and create a checkpoint. When Ambari detects success, the message on the bottom of the window changes. Click **Next**.

Manual Steps Required: Create Checkpoint on NameNode

1. Login to the NameNode host `c6401.ambari.apache.org`.
2. Put the NameNode in safe mode (read-only-mode):

```
sudo su -l hdfs -c 'hdfs dfsadmin -safemode enter'
```
3. Once in Safe Mode, create a checkpoint:

```
sudo su -l hdfs -c 'hdfs dfsadmin -saveNamespace'
```
4. You will be able to proceed once Ambari detects that the NameNode is in Safe Mode and the checkpoint has been created successfully.


Checkpoint not created yet

Next →

7. **Configure Components:** The wizard configures your components, displaying progress bars to let you track the steps. Click **Next** to continue.

Configure Components

Please wait while the wizard configures the components.

- ⚙ Stop All Services  19%
- ⚙ Install Additional NameNode
- ⚙ Install JournalNodes
- ⚙ Start JournalNodes
- ⚙ Disable Secondary NameNode
- ⚙ Reconfigure HDFS

Next

8. **Initialize JournalNodes:** Follow the instructions in the step. You need to login to your **current** NameNode host to run the command to initialize the JournalNodes. When Ambari detects success, the message on the bottom of the window changes. Click **Next**.

Manual Steps Required: Initialize JournalNodes

1. Login to the NameNode host `c6401.ambari.apache.org`.
2. Initialize the JournalNodes by running:


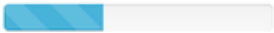

```
sudo su -l hdfs -c 'hdfs namenode -initializeSharedEdits'
```
3. You will be able to proceed once Ambari detects that the JournalNodes have been initialized successfully.

JournalNodes not initialized yet [Next →](#)

9. **Start Components:** The wizard starts the ZooKeeper servers and the NameNode, displaying progress bars to let you track the steps. Click **Next** to continue.

Start Components

Please wait while the wizard starts the components.

-  Start ZooKeeper Servers  37%
-  Start NameNode

[Next](#)

10. **Initialize Metadata:** Follow the instructions in the step. For this step you must login to both the **current** NameNode and the **additional** NameNode. Make sure you are logged into the correct host for each command. Click **Next** when you have completed the two commands. A **Confirmation** popup appears to remind you that you must do both steps. Click **OK** to confirm.

Manual Steps Required: Initialize NameNode HA Metadata

1. Login to the additional NameNode host `c6402.ambari.apache.org`.
2. Initialize the metadata for the additional NameNode by running:


```
sudo su -l hdfs -c 'hdfs namenode -bootstrapStandby'
```
3. Login to the NameNode host `c6401.ambari.apache.org`.
4. Initialize the metadata for NameNode automatic failover by running:








```
sudo su -l hdfs -c 'hdfs zkfc -formatZK'
```
5. Please proceed once you have completed the steps above.

Next →

11 Finalize HA Setup: The wizard the setup, displaying progress bars to let you track the steps. Click **Done** to finish the wizard. After the Ambari Web GUI reloads, you may see some alert notifications. Wait a few minutes until the services come back up. If necessary, restart any components using Ambari Web.

Finalize HA Setup

Please wait while the wizard finalizes the HA setup.

-  Start Additional NameNode 35%
-  Install Failover Controllers
-  Start Failover Controllers
-  Reconfigure HBase
-  Start All Services
-  Delete Secondary NameNode

Done



Note

Choose Services, then start Nagios, after completing all steps in HA wizard.

12If you are using Hive, you must manually change the Hive Metastore FS root to point to the Nameservice URI instead of the NameNode URI. You created the Nameservice ID in the [Get Started](#) step.

- a. Check the current FS root. On the Hive host:

```
hive --config /etc/hive/conf.server --service metatool -listFSRoot
```

The output might look like this:

```
Listing FS Roots..
hdfs://<namenode-host>:8020/apps/hive/warehouse
```

- b. Use this command to change the FS root:

```
$ hive --config /etc/hive/conf.server --service metatool -updateLocation
<new-location> <old-location>

For example, where the Nameservice ID is mycluster:
$ hive --config /etc/hive/conf.server --service metatool -updateLocation
hdfs://mycluster:8020/apps/hive/warehouse hdfs://c6401.ambari.apache.
org:8020/apps/hive/warehouse
```

The output might look like this:

```
Successfully updated the following locations..
Updated X records in SDS table
```

- 13.If you are using Oozie, you must use the Nameservice URI instead of the NameNode URI in your workflow files. For example, where the Nameservice ID is mycluster:

```
<workflow-app xmlns="uri:oozie:workflow:0.2" name="map-reduce-wf">
  <start to="mr-node"/>
  <action name="mr-node">
    <map-reduce>
      <job-tracker>${jobTracker}</job-tracker>
      <name-node>hdfs://mycluster</name-node>
```

- 14.If you are using Hue, to enable NameNode High Availability, you must use httpfs instead of webhdfs to communicate with name nodes inside the cluster. After successfully setting up NameNode High Availability:

- a. Install an httpfs server on any node in the cluster:

```
yum install hadoop-httpfs
```

- b. Ensure that Hue hosts and groups use the httpfs server.

For example, on the httpfs server host, add to httpfs-site.xml the following lines:

```
<property> <name>httpfs.proxyuser.hue.hosts</name> <value>*</value> </
property>
<property> <name>httpfs.proxyuser.hue.groups</name> <value>*</value> </
property>
```

- c. Ensure that groups and hosts in the cluster use the httpfs server. For example, Using Ambari, in **Services > HDFS > Configs** add to core-site.xml the following properties and values:

Table 8.1. Core-site.xml properties and values for NameNode HA on a cluster using Hue

Property	Value
hadoop.proxyuser.httpfs.groups	*
hadoop.proxyuser.httpfs.hosts	*

- d. Using Ambari, in **Services > HDFS** restart the HDFS service in your cluster.
- e. On the Hue host, configure Hue to use the httpfs server by editing hue.ini to include the following lines:

```
fs_defaultfs=hdfs://mycluster
webhdfs_url=http://{fqdn of httpfs server}:14000/webhdfs/v1/
```

- f. Restart the Hue service.

8.1.1. Rolling Back NameNode HA

To roll back NameNode HA to the previous non-HA state use the following step-by-step manual process. Some of the steps are optional depending on your installation.

1. [Stop HBase](#)
2. [Checkpoint the Active NameNode](#)
3. [Stop All Services](#)
4. [Prepare the Ambari Server Host for Rollback](#)
5. [Restore the HBase Configuration](#)
6. [Delete ZK Failover Controllers](#)
7. [Modify HDFS Configurations](#)
8. [Recreate the Secondary NameNode](#)
9. [Re-enable the Secondary NameNode](#)
10. [Delete All JournalNodes](#)
11. [Delete the Additional NameNode](#)
12. [Verify the HDFS Components](#)
13. [Start HDFS](#)

8.1.1.1. Stop HBase

1. From Ambari Web, go to the Services view and select HBase.
2. Click **Stop** on the Management Header.

3. Wait until HBase has stopped completely before continuing.

8.1.1.2. Checkpoint the Active NameNode

If HDFS has been in use **after** you enabled NameNode HA, but you wish to revert back to a non-HA state, you must checkpoint HDFS state before proceeding with the rollback.

If the **Enable NameNode HA** wizard failed and you need to revert back, you can skip this step and move on to [Stop All Services](#).

- If Kerberos security has **not** been enabled on the cluster:

On the Active NameNode host, execute the following commands to save the namespace. You must be the HDFS service user (*\$HDFS_USER*) to do this.

```
sudo su -l $HDFS_USER -c 'hdfs dfsadmin -safemode enter'
sudo su -l $HDFS_USER -c 'hdfs dfsadmin -saveNamespace'
```

- If Kerberos security has been enabled on the cluster:

```
sudo su -l $HDFS_USER -c 'kinit -kt /etc/security/keytabs/nn.service.keytab
nn/$HOSTNAME@$REALM;hdfs dfsadmin -safemode enter'
sudo su -l $HDFS_USER -c 'kinit -kt /etc/security/keytabs/nn.service.keytab
nn/$HOSTNAME@$REALM;hdfs dfsadmin -saveNamespace'
```

Where *\$HDFS_USER* is the HDFS service user, *\$HOSTNAME* is the Active NameNode hostname, and *\$REALM* is your Kerberos realm.

8.1.1.3. Stop All Services

Use the Services view in Ambari Web and click **Stop All** in the Services navigation panel. You must wait until all the services are completely stopped.

8.1.1.4. Prepare the Ambari Server Host for Rollback

Log into the Ambari server host and set the following environment variables to prepare for the rollback procedure:

Table 8.2. Set Environment Variables

Variable	Value
export AMBARI_USER=AMBARI_USERNAME	Substitute the value of the administrative user for Ambari Web. The default value is admin.
export AMBARI_PW=AMBARI_PASSWORD	Substitute the value of the administrative password for Ambari Web. The default value is admin.
export AMBARI_PORT=AMBARI_PORT	Substitute the Ambari Web port. The default value is 8080.
export AMBARI_PROTO=AMBARI_PROTOCOL	Substitute the value of the protocol for connecting to Ambari Web. Options are http or https. The default value is http.
export CLUSTER_NAME=CLUSTER_NAME	Substitute the name of your cluster, set during the Ambari Install Wizard process. For example: mycluster.
export NAMENODE_HOSTNAME=NN_HOSTNAME	Substitute the FQDN of the host for the non-HA NameNode. For example: nn01.mycompany.com.

Variable	Value
export ADDITIONAL_NAMENODE_HOSTNAME= <i>ANN_HOSTNAME</i>	Substitute the FQDN of the host for the additional NameNode in your HA setup.
export SECONDARY_NAMENODE_HOSTNAME= <i>SNN_HOSTNAME</i>	Substitute the FQDN of the host for the Secondary NameNode for the non-HA setup.
export JOURNALNODE1_HOSTNAME= <i>JOUR1_HOSTNAME</i>	Substitute the FQDN of the host for the first Journal Node.
export JOURNALNODE2_HOSTNAME= <i>JOUR2_HOSTNAME</i>	Substitute the FQDN of the host for the second Journal Node.
export JOURNALNODE3_HOSTNAME= <i>JOUR3_HOSTNAME</i>	Substitute the FQDN of the host for the third Journal Node.



Important

Double check that these environment variables are set correctly.

8.1.1.5. Restore the HBase Configuration

If you have installed HBase, you may need to restore a configuration to its pre-HA state.

1. To check if your current HBase configuration needs to be restored, on the Ambari Server host:

```
/var/lib/ambari-server/resources/scripts/configs.sh -u $AMBARI_USER -p $AMBARI_PW -port $AMBARI_PORT get localhost $CLUSTER_NAME hbase-site
```

Where the environment variables you set up before substitute for the variable names.

Look for the configuration property `hbase.rootdir`. If the value is set to the NameService ID you set up using the **Enable NameNode HA** wizard, you need to revert the `hbase-site` configuration set up back to non-HA values. If it points instead to a specific NameNode host, it does not need to be rolled back and you can go on to [Delete ZK Failover Controllers](#).

For example:

```
"hbase.rootdir": "hdfs://name-service-id:8020/apps/hbase/data"
The hbase.rootdir property points to the NameService ID and the value needs to be rolled back

"hbase.rootdir": "hdfs://nn01.mycompany.com:8020/apps/hbase/data"
The hbase.rootdir property points to a specific NameNode host and not a NameService ID. This does not need to be rolled back.
```

2. If you need to roll back the `hbase.rootdir` value, on the Ambari Server host, use the `config.sh` script to make the necessary change:

```
/var/lib/ambari-server/resources/scripts/configs.sh -u $AMBARI_USER -p $AMBARI_PW -port $AMBARI_PORT set localhost $CLUSTER_NAME hbase-site hbase.rootdir hdfs://${NAMENODE_HOSTNAME}:8020/apps/hbase/data
```

Where the environment variables you set up before substitute for the variable names.

3. Verify that the `hbase.rootdir` property has been restored properly. On the Ambari Server host:

```
/var/lib/ambari-server/resources/scripts/configs.sh -u $AMBARI_USER -p $AMBARI_PW -port $AMBARI_PORT get localhost $CLUSTER_NAME hbase-site
```


The `hbase.rootdir` property should now be set to the NameNode hostname, not the NameService ID.

8.1.1.6. Delete ZK Failover Controllers

You may need to delete ZK Failover Controllers.

1. To check if you need to delete ZK Failover Controllers, on the Ambari Server host:

```
curl -u ${AMBARI_USER}:${AMBARI_PW} -H "X-Requested-By: ambari" -i -X GET
${AMBARI_PROTO}://localhost:${AMBARI_PORT}/api/v1/clusters/${CLUSTER_NAME}/
host_components?HostRoles/component_name=ZKFC
```

If this returns an empty `items` array, you can go on to [Modify HDFS Configuration](#). Otherwise you must use the DELETE commands below.

2. To delete all ZK Failover Controllers, on the Ambari Server host:

```
curl -u ${AMBARI_USER}:${AMBARI_PW} -H "X-Requested-By: ambari" -i -X DELETE
${AMBARI_PROTO}://localhost:${AMBARI_PORT}/api/v1/clusters/${CLUSTER_NAME}/
hosts/${NAMENODE_HOSTNAME}/host_components/ZKFC
curl -u ${AMBARI_USER}:${AMBARI_PW} -H "X-Requested-By: ambari" -i -X DELETE
${AMBARI_PROTO}://localhost:${AMBARI_PORT}/api/v1/clusters/${CLUSTER_NAME}/
hosts/${ADDITIONAL_NAMENODE_HOSTNAME}/host_components/ZKFC
```

3. Verify that the ZK Failover Controllers have been deleted. On the Ambari Server host:

```
curl -u ${AMBARI_USER}:${AMBARI_PW} -H "X-Requested-By: ambari" -i
${AMBARI_PROTO}://localhost:${AMBARI_PORT}/api/v1/clusters/${CLUSTER_NAME}/
host_components?HostRoles/component_name=ZKFC
```

This command should return an empty `items` array.

8.1.1.7. Modify HDFS Configurations

You may need to modify your `hdfs-site` configuration and/or your `core-site` configuration.

1. To check if you need to modify your `hdfs-site` configuration, on the Ambari Server host:

```
/var/lib/ambari-server/resources/scripts/configs.sh -u $AMBARI_USER -p
$AMBARI_PW -port $AMBARI_PORT get localhost $CLUSTER_NAME hdfs-site
```

If you see **any** of the following properties, you must delete them from your `hdfs-site` configuration.

- `dfs.nameservices`
- `dfs.client.failover.proxy.provider.${NAMESERVICE_ID}`
- `dfs.ha.namenodes.${NAMESERVICE_ID}`
- `dfs.ha.fencing.methods`
- `dfs.ha.automatic-failover.enabled`

- `dfs.namenode.http-address.${NAMESERVICE_ID}.nn1`
- `dfs.namenode.http-address.${NAMESERVICE_ID}.nn2`
- `dfs.namenode.rpc-address.${NAMESERVICE_ID}.nn1`
- `dfs.namenode.rpc-address.${NAMESERVICE_ID}.nn2`
- `dfs.namenode.shared.edits.dir`
- `dfs.journalnode.edits.dir`
- `dfs.journalnode.http-address`
- `dfs.journalnode.kerberos.internal.spnego.principal`
- `dfs.journalnode.kerberos.principal`
- `dfs.journalnode.keytab.file`

Where `${NAMESERVICE_ID}` is the NameService ID you created when you ran the **Enable NameNode HA** wizard.

2. To delete these properties, execute the following for **each property** you found. On the Ambari Server host:

```
/var/lib/ambari-server/resources/scripts/configs.sh -u $AMBARI_USER -p $AMBARI_PW -port $AMBARI_PORT delete localhost $CLUSTER_NAME hdfs-site property_name
```

Where you replace *property_name* with the name of **each** of the properties to be deleted.

3. Verify that all of the properties have been deleted. On the Ambari Server host:

```
/var/lib/ambari-server/resources/scripts/configs.sh -u $AMBARI_USER -p $AMBARI_PW -port $AMBARI_PORT get localhost $CLUSTER_NAME hdfs-site
```

None of the properties listed above should be present.

4. To check if you need to modify your `core-site` configuration, on the Ambari Server host:

```
/var/lib/ambari-server/resources/scripts/configs.sh -u $AMBARI_USER -p $AMBARI_PW -port $AMBARI_PORT get localhost $CLUSTER_NAME core-site
```

5. If you see the property `ha.zookeeper.quorum`, it must be deleted. On the Ambari Server host:

```
/var/lib/ambari-server/resources/scripts/configs.sh -u $AMBARI_USER -p $AMBARI_PW -port $AMBARI_PORT delete localhost $CLUSTER_NAME core-site ha.zookeeper.quorum
```

6. If the property `fs.defaultFS` is set to the NameService ID, it must be reverted back to its non-HA value. For example:

```
"fs.defaultFS" : "hdfs://name-service-id"
```

The property `fs.defaultFS` needs to be modified as it points to a NameService ID

```
"fs.defaultFS" : "hdfs://nn01.mycompany.com"
```

The property `fs.defaultFS` does not need to be changed as it points to a specific NameNode and not a NameService ID

7. To revert the property `fs.defaultFS` to the NameNode host value, on the Ambari Server host:

```
/var/lib/ambari-server/resources/scripts/configs.sh -u $AMBARI_USER -p $AMBARI_PW -port $AMBARI_PORT set localhost $CLUSTER_NAME core-site fs.defaultFS hdfs://${NAMENODE_HOSTNAME}
```

8. Verify that the `core-site` properties are now properly set. On the Ambari Server host:

```
/var/lib/ambari-server/resources/scripts/configs.sh -u $AMBARI_USER -p $AMBARI_PW -port $AMBARI_PORT get localhost $CLUSTER_NAME core-site
```

The property `fs.defaultFS` should be set to point to the NameNode host and the property `ha.zookeeper.quorum` should not be there.

8.1.1.8. Recreate the Secondary NameNode

You may need to recreate your Secondary NameNode.

1. To check to see if you need to recreate the Secondary NameNode, on the Ambari Server host:

```
curl -u ${AMBARI_USER}:${AMBARI_PW} -H "X-Requested-By: ambari" -i -X GET ${AMBARI_PROTO}://localhost:${AMBARI_PORT}/api/v1/clusters/${CLUSTER_NAME}/host_components?HostRoles/component_name=SECONDARY_NAMENODE
```

If this returns an empty `items` array, you must recreate your Secondary NameNode. Otherwise you can go on to [Re-enable Secondary NameNode](#).

2. Recreate your Secondary NameNode. On the Ambari Server host:

```
curl -u ${AMBARI_USER}:${AMBARI_PW} -H "X-Requested-By: ambari" -i -X POST -d '{"host_components" : [{"HostRoles": {"component_name": "SECONDARY_NAMENODE"}}]}' ${AMBARI_PROTO}://localhost:${AMBARI_PORT}/api/v1/clusters/${CLUSTER_NAME}/hosts?Hosts/host_name=${SECONDARY_NAMENODE_HOSTNAME}
```

3. Verify that the Secondary NameNode now exists. On the Ambari Server host:

```
curl -u ${AMBARI_USER}:${AMBARI_PW} -H "X-Requested-By: ambari" -i -X GET ${AMBARI_PROTO}://localhost:${AMBARI_PORT}/api/v1/clusters/${CLUSTER_NAME}/host_components?HostRoles/component_name=SECONDARY_NAMENODE
```

This should return a non-empty `items` array containing the Secondary NameNode.

8.1.1.9. Re-enable the Secondary NameNode

To re-enable the Secondary NameNode, on the Ambari Server host:

```
curl -u ${AMBARI_USER}:${AMBARI_PW} -H "X-Requested-By: ambari" -i -X PUT -d '{"RequestInfo":{"context":"Enable Secondary NameNode"},"Body":{"HostRoles":
```

```
{"state":"INSTALLED"}}}' ${AMBARI_PROTO}://localhost:${AMBARI_PORT}/api/v1/
clusters/${CLUSTER_NAME}/hosts/${SECONDARY_NAMENODE_HOSTNAME}/host_components/
SECONDARY_NAMENODE
```

- If this returns 200, go to [Delete All JournalNodes](#).
- If this returns 202, wait a few minutes and run the following on the Ambari Server host:

```
curl -u ${AMBARI_USER}:${AMBARI_PW} -H "X-Requested-By: ambari" -i
-X GET "${AMBARI_PROTO}://localhost:${AMBARI_PORT}/api/v1/clusters/
${CLUSTER_NAME}/host_components?HostRoles/component_name=SECONDARY_NAMENODE&
fields=HostRoles/state"
```

When "state" : "INSTALLED" is in the response, go on to the next step.

8.1.1.10. Delete All JournalNodes

You may need to delete any JournalNodes.

1. To check to see if you need to delete JournalNodes, on the Ambari Server host:

```
curl -u ${AMBARI_USER}:${AMBARI_PW} -H "X-Requested-By: ambari" -i -X GET
${AMBARI_PROTO}://localhost:${AMBARI_PORT}/api/v1/clusters/${CLUSTER_NAME}/
host_components?HostRoles/component_name=JOURNALNODE
```

If this returns an empty items array, you can go on to [Delete Additional NameNode](#). Otherwise you must delete the JournalNodes.

2. To delete the JournalNodes, on the Ambari Server host:

```
curl -u ${AMBARI_USER}:${AMBARI_PW} -H "X-Requested-By: ambari" -i -X DELETE
${AMBARI_PROTO}://localhost:${AMBARI_PORT}/api/v1/clusters/${CLUSTER_NAME}/
hosts/${JOURNALNODE1_HOSTNAME}/host_components/JOURNALNODE

curl -u ${AMBARI_USER}:${AMBARI_PW} -H "X-Requested-By: ambari" -i -X DELETE
${AMBARI_PROTO}://localhost:${AMBARI_PORT}/api/v1/clusters/${CLUSTER_NAME}/
hosts/${JOURNALNODE2_HOSTNAME}/host_components/JOURNALNODE

curl -u ${AMBARI_USER}:${AMBARI_PW} -H "X-Requested-By: ambari" -i -X DELETE
${AMBARI_PROTO}://localhost:${AMBARI_PORT}/api/v1/clusters/${CLUSTER_NAME}/
hosts/${JOURNALNODE3_HOSTNAME}/host_components/JOURNALNODE
```

3. Verify that all the JournalNodes have been deleted. On the Ambari Server host:

```
curl -u ${AMBARI_USER}:${AMBARI_PW} -H "X-Requested-By: ambari" -i -X GET
${AMBARI_PROTO}://localhost:${AMBARI_PORT}/api/v1/clusters/${CLUSTER_NAME}/
host_components?HostRoles/component_name=JOURNALNODE
```

This should return an empty items array.

8.1.1.11. Delete the Additional NameNode

You may need to delete your Additional NameNode.

1. To check to see if you need to delete your Additional NameNode, on the Ambari Server host:

```
curl -u ${AMBARI_USER}:${AMBARI_PW} -H "X-Requested-By: ambari" -i -X GET  
${AMBARI_PROTO}://localhost:${AMBARI_PORT}/api/v1/clusters/${CLUSTER_NAME}/  
host_components?HostRoles/component_name=NAMENODE
```

If the `items` array contains two NameNodes, the Additional NameNode must be deleted.

2. To delete the Additional NameNode that was set up for HA, on the Ambari Server host:

```
curl -u ${AMBARI_USER}:${AMBARI_PW} -H "X-Requested-By: ambari" -i -X DELETE  
${AMBARI_PROTO}://localhost:${AMBARI_PORT}/api/v1/clusters/${CLUSTER_NAME}/  
hosts/${ADDITIONAL_NAMENODE_HOSTNAME}/host_components/NAMENODE
```

3. Verify that the Additional NameNode has been deleted:

```
curl -u ${AMBARI_USER}:${AMBARI_PW} -H "X-Requested-By: ambari" -i -X GET  
${AMBARI_PROTO}://localhost:${AMBARI_PORT}/api/v1/clusters/${CLUSTER_NAME}/  
host_components?HostRoles/component_name=NAMENODE
```

This should return an `items` array that shows only one NameNode.

8.1.1.12. Verify your HDFS Components

Make sure you have the correct components showing in HDFS.

1. In Ambari Web, go to the Services view and select HDFS from the Services navigation panel.
2. Check the Summary panel and make sure that the first three lines look like this:
 - NameNode
 - SNameNode
 - DataNodes

You should **not** see any line for JournalNodes.

8.1.1.13. Start HDFS

1. On the HDFS page of the Services view, click **Start** in the Management Header to start up HDFS. Wait until the service is fully started and has passed the service check.

If HDFS does not start, you may need to go through the previous steps again.
2. Start all the other services by using **Start All** in the Services navigation panel.

Part III. Hadoop 1.x - Deploying, Configuring, and Upgrading Ambari

This section describes setting up Hadoop 1.x. It includes:

- [Installing, Configuring, and Deploying the Cluster for Hadoop 1.x](#)
- [Troubleshooting Ambari Deployments for Hadoop 1.x](#)
- [Appendix: Upgrading Ambari Server to 1.4.4](#)
- [Appendix: Upgrading Ambari Server to 1.2.5](#)
- [Appendix: Upgrading the HDP Stack to 1.3.3](#)
- [Appendix: Configuring Ports for Hadoop 1.x](#)

Hortonworks Data Platform

Installing Hadoop Using Apache Ambari

(Aug 1, 2014)

9. Hadoop 1.x - Installing, Configuring, and Deploying the Cluster

This section describes using the Ambari install wizard in your browser to complete your installation, configuration and deployment of Hadoop.

9.1. Log into Apache Ambari

Once you have started the Ambari service, you can access the Ambari Install Wizard through your browser.

1. Point your browser to `http://{main.install.hostname}:8080`.
2. Log in to the Ambari Server using the default username/password: `admin/admin`. You can change this later to whatever you want.

9.2. Welcome

The first step creates the cluster name.

1. At the **Welcome** page, type a name for the cluster you want to create in the text box. No whitespaces or special characters can be used in the name.
2. Click the **Next** button.

9.3. Select Stack

The Service Stack (or simply the Stack) is a coordinated and tested set of Hadoop components. Use the radio button to select the Stack version you want to install. To install a Hadoop 1 stack, select HDP 1.3.3, in Stacks.

Select Stack

Please select the service stack that you want to use to install your Hadoop cluster.

Stacks

- HDP 2.0.6
- HDP 1.3.3
- HDP 1.3.2

▶ Advanced Repository Options

← Back Next →

Under Advanced Repository Options, you can specify the Base URLs of your local repositories for each Operating System you plan to use in your cluster. You should have configured the Base URLs for your local repositories in [Optional: Configure Ambari for Local Repositories](#).

▼ **Advanced Repository Options**

Customize the repository Base URLs for downloading the Stack software packages. If your hosts do not have access to the internet, you will have to create a local mirror of the Stack repository that is accessible by all hosts and use those Base URLs here.

Important: When using local mirror repositories, you only need to provide Base URLs for the Operating System you are installing for your Stack. Uncheck all other repositories.

OS	Base URL
Red Hat 5	
<input checked="" type="checkbox"/> CentOS 5	<input type="text" value="http://public-repo-1.hortonworks.com/HDP/centos5/2.x/"/>
<input type="checkbox"/> Oracle Linux 5	
Red Hat 6	
<input checked="" type="checkbox"/> CentOS 6	<input type="text" value="http://public-repo-1.hortonworks.com/HDP/centos6/2.x/"/>
<input type="checkbox"/> Oracle Linux 6	
<input checked="" type="checkbox"/> SLES 11	<input type="text" value="http://public-repo-1.hortonworks.com/HDP/suse11/2.x/u"/>
<input type="checkbox"/> SUSE 11	

9.4. Install Options

In order to build up the cluster, the install wizard needs to know general information about how you want to set up your cluster. You need to supply the FQDN of each of your hosts. The wizard also needs to access the private key file you created in [Set Up Password-less SSH](#). It uses these to locate all the hosts in the system and to access and interact with them securely.

1. Use the **Target Hosts** text box to enter your list of host names, one per line. You can use ranges inside brackets to indicate larger sets of hosts. For example, for host01.domain through host10.domain use `host[01-10].domain`



Note

If you are deploying on EC2, use the **internal Private DNS** hostnames.

2. If you want to let Ambari automatically install the Ambari Agent on all your hosts using SSH, select **Provide your SSH Private Key** and either use the **Choose File** button in the **Host Registration Information** section to find the private key file that matches the public key you installed earlier on all your hosts or cut and paste the key into the text box manually.



Note

If you are using IE 9, the **Choose File** button may not appear. Use the text box to cut and paste your private key manually.

Fill in the username for the SSH key you have selected. If you do not want to use `root`, you must provide the username for an account that can execute `sudo` without entering a password.

3. If you do not want Ambari to automatically install the Ambari Agents, select **Perform manual registration**. See [Appendix: Installing Ambari Agents Manually](#) for more information.
4. Click the **Register and Confirm** button to continue.

9.5. Confirm Hosts

This screen lets you confirm that Ambari has located the correct hosts for your cluster and to check those hosts to make sure they have the correct directories, packages, and processes to continue the install.

If any hosts were selected in error, you can remove them by selecting the appropriate checkboxes and clicking the grey **Remove Selected** button. To remove a single host, click the small white **Remove** button in the Action column.

At the bottom of the screen, you may notice a yellow box that indicates some warnings were encountered during the check process. For example, your host may have already had a copy of `wget` or `curl`. Click **Click here to see the warnings** to see a list of what was checked and what caused the warning. On the same page you can get access to a python script that can help you clear any issues you may encounter and let you run **Rerun Checks**.

When you are satisfied with the list of hosts, click **Next**.

9.6. Choose Services

Hortonworks Data Platform is made up of a number of services. You must at a minimum install HDFS, but you can decide which of the other services you want to install. See [Understand the Basics](#) for more information on your options.

1. Select **all** to preselect all items or **minimum** to preselect only HDFS.
2. Use the checkboxes to unselect (if you have used **all**) or select (if you have used **minimum**) to arrive at your desired list of components.



Note

If you want to use Ambari for monitoring your cluster, make sure you select **Nagios** and **Ganglia**. If you do not select them, you get a warning popup when you finish this section. If you are using other monitoring tools, you can ignore the warning.

3. When you have made your selections, click **Next**.

9.7. Assign Masters

The Ambari install wizard attempts to assign the master nodes for various services you have selected to appropriate hosts in your cluster. The right column shows the current service

assignments by host, with the hostname and its number of CPU cores and amount of RAM indicated.

1. To change locations, click the dropdown list next to the service in the left column and select the appropriate host.
2. To remove a ZooKeeper instance, click the green minus icon next to the host address you want to remove.
3. When you are satisfied with the assignments, click the **Next** button.

9.8. Assign Slaves and Clients

The Ambari install wizard attempts to assign the slave components (DataNodes, NodeManagers, and RegionServers) to appropriate hosts in your cluster. It also attempts to select hosts for installing the appropriate set of clients.

1. Use **all** or **none** to select all of the hosts in the column or none of the hosts, respectively.

If a host has a red asterisk next to it, that host is also running one or more master components. Hover your mouse over the asterisk to see which master components are on that host.

2. Fine-tune your selections by using the checkboxes next to specific hosts.



Note

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 attempt to start it on the same host as the Ambari server, however, you need to start it with the `-p` option, as its default port is 8080 and that conflicts with the Ambari Web default port.

```
/usr/lib/hbase/bin/hbase-daemon.sh start rest -p  
<custom_port_number>
```

3. When you are satisfied with your assignments, click the **Next** button.

9.9. Customize Services

The **Customize Services** screen presents you with a set of tabs that let you manage configuration settings for Hadoop components. The wizard attempts to set reasonable defaults for each of the options here, but you can use this set of tabs to tweak those settings. and you are strongly encouraged to do so, as your requirements may be slightly different. Pay particular attention to the directories suggested by the installer.



Note

In **HDFS Services Configs General**, make sure to enter an integer value, in bytes, that sets the HDFS maximum edit log size for checkpointing. A typical value is 500000000.

Hover your mouse over each of the properties to see a brief description of what it does. The number of tabs you see is based on the type of installation you have decided to do. In a complete installation there are nine groups of configuration properties and other related options, such as database settings for Hive and Oozie and admin name/password and alert email for Nagios.

The install wizard sets reasonable defaults for all properties except for those related to databases in the Hive tab and the Oozie tab, and two related properties in the Nagios tab. These four are marked in red and are the only ones you *must* set yourself.



Note

If you decide to use an existing database instance for Hive/HCatalog or for Oozie, you must have completed the preparations described in [Using Non-Default Databases](#) prior to running the install wizard.

Click the name of the group in each tab to expand and collapse the display.

9.9.1. Service Users and Groups

The individual services in Hadoop are each run under the ownership of a corresponding Unix account. These accounts are known as service users. These service users belong to a special Unix group. In addition there is a special service user for running smoke tests on components during installation and on-demand using the Management Header in the **Services** View of the Ambari Web GUI. Any of these users and groups can be customized using the **Misc** tab of the **Customize Services** step.

If you choose to customize names, Ambari checks to see if these custom accounts already exist. If they do not exist, Ambari creates them. The default accounts are always created during installation whether or not custom accounts are specified. These default accounts are not used and can be removed post-install.



Note

All new service user accounts, and any existing user accounts used as service users, must have a UID ≥ 1000 .

Table 9.1. Service Users

Service	Component	Default User Account
HDFS	NameNode	hdfs
	SecondaryNameNode	
	DataNode	
MapReduce	JobTracker	mapred
	HistoryServer	
	TaskTracker	
Hive	Hive Metastore	hive
	HiveServer2	
HCat	HCatalog Server	hcat

Service	Component	Default User Account
WebHCat	WebHCat Server	hcat
Oozie	Oozie Server	oozie
HBase	MasterServer	hbase
	RegionServer	
ZooKeeper	ZooKeeper	zookeeper
Ganglia	Ganglia Server	nobody
	Ganglia Collectors	
Nagios	Nagios Server	nagios ^a
Smoke Test ^b	All	ambari-qa

^aIf you plan to use an existing user account named "nagios", that "nagios" account must be in a group named "nagios". If you customize this account, that account will be created and put in a group "nagios".

^bThe Smoke Test user performs smoke tests against cluster services as part of the install process. It also can perform these on-demand from the Ambari Web GUI.

Table 9.2. Service Group

Service	Components	Default Group Account
All	All	hadoop

9.9.2. Properties That Depend on Service Usernames/Groups

Some properties must be set to match specific service usernames or service groups. If you have set up non-default, customized service usernames for the HDFS or HBase service or the Hadoop group name, you must edit the following properties:

Table 9.3. HDFS Settings: Advanced

Property Name	Value
dfs.permissions.supergroup	The same as the HDFS username. The default is "hdfs"
dfs.cluster.administrators	A single space followed by the HDFS username.
dfs.block.local-path-access.user	The HBase username. The default is "hbase".

Table 9.4. MapReduce Settings: Advanced

Property Name	Value
mapreduce.tasktracker.group	The Hadoop group name. The default is "hadoop".
mapreduce.cluster.administrators	A single space followed by the Hadoop group name.

9.9.3. Recommended Memory Configurations for the MapReduce Service

The following recommendations can help you determine appropriate memory configurations based on your usage scenario:

- Make sure that there is enough memory for all the processes. Remember that system processes take around 10% of the available memory.

- For co-deploying an HBase RegionServer and MapReduce service on the same node, reduce the RegionServer's heap size (use the **HBase Settings: RegionServer:** `HBase Region Servers maximum Java heap size` property to modify the RegionServer heap size).
- For co-deploying an HBase RegionServer and the MapReduce service on the same node, or for memory intensive MapReduce applications, modify the map and reduce slots as suggested in the following example:

EXAMPLE: For co-deploying an HBase RegionServer and the MapReduce service on a machine with 16GB of available memory, the following would be a recommended configuration:

2 GB: system processes

8 GB: MapReduce slots. 6 Map + 2 Reduce slots per 1 GB task

4 GB: HBase RegionServer

1 GB: TaskTracker

1 GB: DataNode

To change the number of Map and Reduce slots based on the memory requirements of your application, use the following properties:

MapReduce Settings: TaskTracker: `Number of Map slots per node`

MapReduce Settings: TaskTracker: `Number of Reduce slots per node`

9.10. Review

The assignments you have made are displayed. Check to make sure everything is correct. If you need to make changes, use the left navigation bar to return to the appropriate screen.

To print your information for later reference, click **Print**.

When you are satisfied with your choices, click the **Deploy** button.

9.11. Install, Start and Test

The progress of the install is shown on the screen. Each component is installed and started and a simple test is run on the component. You are given an overall status on the process in the progress bar at the top of the screen and a host by host status in the main section.

To see specific information on what tasks have been completed per host, click the link in the **Message** column for the appropriate host. In the **Tasks** pop-up, click the individual task to see the related log files. You can select filter conditions by using the **Show** dropdown list. To see a larger version of the log contents, click the **Open** icon or to copy the contents to the clipboard, use the **Copy** icon.

Depending on which components you are installing, the entire process may take 40 or more minutes. Please be patient.

When **Successfully installed and started the services** appears, click **Next**.

9.12. Summary

The Summary page gives you a summary of the accomplished tasks. Click **Complete**. You are taken to the Ambari Web GUI.

10. Troubleshooting Ambari Deployments

The following information can help you troubleshoot issues you may run into with your Ambari-based installation.

10.1. Review 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 a specific number, for example N. Pay particular attention to the following three files:

- site-N.pp - the puppet file corresponding to a specific task.
- output-N.txt - output from puppet file execution.
- errors-N.txt - error messages.

10.2. Quick Checks

- Make sure all the appropriate services are running. If you have access to Ambari Web, use the **Services View** to check the status of each component. If you do not have access to Manage Services, you must start and stop the services manually.
- If the first HDFS `put` command fails to replicate the block, the clocks in the nodes may not be synchronized. Make sure that Network Time Protocol (NTP) is enabled for your cluster.
- If HBase does not start, check if its slaves are running on 64-bit JVMs. Ambari requires that all hosts must run on 64-bit machines.
- Make sure `umask` is set to 0022.
- Make sure the HCatalog host can access the MySQL server. From a shell try:

```
mysql -h $FQDN_for_MySQL_server -u $FQDN_for_HCatalog_Server -p
```

You will need to provide the password you set up for Hive/HCatalog during the installation process.

- Make sure MySQL is running. By default, MySQL server does not start automatically on reboot.

To set auto-start on boot, from a shell, type:

```
chkconfig --level 35 mysql on
```

To then start the service manually from a shell, type:

```
service mysqld start
```

10.3. Specific Issues

The following are common issues you might encounter.

10.3.1. Problem: Browser crashed before Install Wizard completed

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

10.3.1.1. Solution

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

- The browser closes prior to hitting 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 the **Deploy** button has launched the **Install, Start, and Test** screen

Re-launch the same browser and continue the process or use a different browser and re-login. You are returned to the **Install, Start, and Test** screen.

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

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

10.3.2.1. 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 nav 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** and proceed to the Monitoring **Dashboard**.
 2. Use the **Services View** to make your changes.
 3. Re-start the service using the **Mangement Header**.
- 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 the browser's cache.
 4. Re-run the entire Install Wizard.

10.3.3. Problem: “Unable to create new native thread” exceptions in HDFS DataNode logs or those of any system daemon

If your `nproc` limit is incorrectly configured, the smoke tests fail and you see an error similar to this in the DataNode logs:

```
INFO org.apache.hadoop.hdfs.DFSClient: Exception
increaseBlockOutputStream java.io.EOFException
INFO org.apache.hadoop.hdfs.DFSClient: Abandoning block
blk_-6935524980745310745_139190
```

10.3.3.1. Solution:

In certain recent Linux distributions (like RHEL/Centos/Oracle Linux 6.x), the default value of `nproc` is lower than the value required if you are deploying the HBase service. To change this value:

1. Using a text editor, open `/etc/security/limits.d/90-nproc.conf` and change the `nproc` limit to approximately 32000. For more information, see [ulimit and nproc recommendations for HBase servers](#).
2. Restart the HBase server.

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

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

10.3.4.1. Solution:

You may have incompatible versions of some software components in your environment. Check the list in [Check Existing Installs](#) and make any necessary changes. Also make sure you are running a [Supported Operating System](#)

10.3.5. 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
```

10.3.5.1. Solution:

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

Restart the installation process.

10.3.6. Problem: The HCatalog Daemon Metastore Smoke Test Fails

If the HCatalog smoke test fails, this is displayed in your console:

```
Metastore startup failed, see /var/log/hcatalog/hcat.err
```

10.3.6.1. Solution:

1. Log into the HCatalog node in your cluster
2. Open `/var/log/hcatalog/hcat.err` or `/var/log/hive/hive.log` (one of the two will exist depending on the installation) with a text editor
3. In the file, see if there is a `MySQL Unknown Host Exception` like this:

```
at java.lang.reflect.Method.invoke (Method.java:597)
at org.apache.hadoop.util.Runjar.main (runjar.java:156)
Caused by: java.net.UnknownHostException:mysql.host.com
at java.net.InetAddress.getAllByName (InetAddress.java:1157)
```

This exception can be thrown if you are using a previously existing MySQL instance and you have incorrectly identified the hostname during the installation process. When you do the reinstall, make sure this name is correct.

4. In the file, see if there is an ERROR Failed initializing database entry like this:

```
11/12/29 20:52:04 ERROR DataNucleus.Plugin: Bundle
org.eclipse.jdt.core required
11/12/29 20:52:04 ERROR DataStore.Schema: Failed initialising
database
```

This exception can be thrown if you are using a previously existing MySQL instance and you have incorrectly identified the username/password during the installation process. It can also occur when the user you specify does not have adequate privileges on the database. When you do the reinstall, make sure this username/password is correct and that the user has adequate privilege.

5. Restart the installation process.

10.3.7. Problem: MySQL and Nagios fail to install on RightScale CentOS 5 images on EC2

When using a RightScale CentOS 5 AMI on Amazon EC2, in certain cases MySQL and Nagios will fail to install. The MySQL failure is due to a conflict with the pre-installed MySQL and the use of the RightScale EPEL repository (error "Could not find package mysql-server"). Nagios fails to install due to conflicts with the RightScale php-common library.

10.3.7.1. Solution:

On the machines that will host MySQL and Nagios as part of your Hadoop cluster, perform the following:

1. Remove the existing MySQL server

```
yum erase MySQL-server-community
```

2. Install MySQL server with a disabled RightScale EPEL repository

```
yum install mysql-server --disable-repo=rightscales-epel
```

3. Remove the php-common library

```
yum erase php-common-5.2.4-RightScale.x86
```

10.3.8. Problem: Trouble starting Ambari on system reboot

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

10.3.8.1. 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
```

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

Charts appear incorrectly or not at all despite being available in the native Ganglia interface or Host health status is displayed incorrectly.

10.3.9.1. 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.

10.3.10. 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 bug that causes this crash.

10.3.10.1. Solution:

Upgrade to Python version 2.6.8-0.15.1

10.3.11. 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.

10.3.11.1. 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>
```

10.3.12. 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.

10.3.12.1. 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
```

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

When a cluster is first started, some graphs, like **Services View -> HDFS** and **Services View -> MapReduce**, do not plot a complete hour of data, instead showing data only for the length of time the service has been running. Other graphs display the run of a complete hour.

10.3.13.1. Solution

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

10.3.14. Problem: After performing a cluster install the Nagios server is not started

The Nagios server is not started after a cluster install and you are unable to manage it from Ambari Web.

10.3.14.1. Solution

1. Log into the Nagios server host.
2. Confirm that the Nagios server is not running. From a shell:

```
ps -ef | grep nagios
```

You should not see a Nagios process running.

3. Start the Nagios process manually. From a shell:

```
service nagios start
```

4. The server starts. You should be able to see that started state reflected in Ambari Web. You can now manage (start/stop) Nagios from Ambari Web.

10.3.15. 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`.

10.3.15.1. Solution

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

10.3.16. Problem: Updated configuration changes are not pushed to client/gateway nodes

Currently configuration changes are only pushed to daemon running nodes, so any changes are not automatically pushed to client only nodes such as gateway nodes.

10.3.16.1. Solution

Copy the files to the client nodes manually.

11. Appendix: Upgrading Ambari Server to 1.4.4

This procedure upgrades Ambari Server from version 1.2.5 and above to 1.4.4. If your current Ambari Server version is 1.2.4 or below, you must [upgrade the Ambari Server version to 1.2.5](#) before upgrading to version 1.4.4. Upgrading the Ambari Server version does not change the underlying Hadoop Stack.



Note

You must know the location of the Nagios server for Step 9. Use the **Services View-> Summary** panel to locate the host on which it is running.

1. Stop the Ambari Server and all Ambari Agents. From the Ambari Server host:

```
ambari-server stop
```

From each Ambari Agent host:

```
ambari-agent stop
```

2. Get the new Ambari bits. Use `wget` to fetch the repository file and replace the old repo file with the new repo file on every host.

- Fetch the new repo file:

For RHEL/CentOS 5/Oracle Linux 5

```
wget http://public-repo-1.hortonworks.com/ambari/centos5/1.x/updates/1.4.4.23/ambari.repo
```

For RHEL/CentOS 6/Oracle Linux 6

```
wget http://public-repo-1.hortonworks.com/ambari/centos6/1.x/updates/1.4.4.23/ambari.repo
```

For SLES 11

```
wget http://public-repo-1.hortonworks.com/ambari/suse11/1.x/updates/1.4.4.23/ambari.repo
```



Important

Check your current directory before you download the new repo file to make sure that there are no previous versions of the file. If you do not, and a previous version exists, the new download will be saved with a numeric extension such as `ambari.repo.1`. Make sure that the version you copy is the new version.

- Replace the old repo file with the new repo file.

For RHEL/CentOS 5/Oracle Linux 5

```
cp ambari.repo /etc/yum.repos.d/ambari.repo
```

For RHEL/CentOS 6/Oracle Linux 6

```
cp ambari.repo /etc/yum.repos.d/ambari.repo
```

For SLES 11

```
cp ambari.repo /etc/zypp/repos.d/ambari.repo
```



Note

If your cluster does not have access to the Internet, set up a local repository with this data before you continue. See [Configure the Local Repositories](#) for more information.

3. Upgrade Ambari Server.

From the Ambari Server host:

- RHEL/CentOS

```
yum clean all
yum upgrade ambari-server ambari-log4j
```

- SLES

```
zypper clean
zypper up ambari-server ambari-log4j
```

4. Check for upgrade success:

- As the process runs, the console should produce output similar, although not identical, to this:

```
Setting up Upgrade Process
Resolving Dependencies
--> Running transaction check
---> Package ambari-agent.x86_64 0:1.2.2.3-1 will be updated
---> Package ambari-agent.x86_64 0:1.2.2.4-1 will be updated ...
---> Package ambari-agent.x86_64 0:1.2.2.5-1 will be an update ...
```

After the process is complete, check each host to make sure the new 1.4.4 files have been installed:

```
rpm -qa | grep ambari
```

- If the upgrade fails, the console should produce output similar to this:

```
Setting up Upgrade Process
No Packages marked for Update
```

5. Check to see if you have a folder named `/etc/ambari-server/conf.save`. If you do, rename it back:

```
mv /etc/ambari-server/conf.save /etc/ambari-server/conf
```

6. Upgrade the Ambari Server schema. From the Ambari Server host:

```
ambari-server upgrade
```


7. Upgrade the Ambari Agent on all hosts. From each Ambari Agent host:

- RHEL/CentOS

```
yum upgrade ambari-agent ambari-log4j
```

- SLES

```
zypper up ambari-agent ambari-log4j
```



Note

If you get a warning that begins "There are some running programs that use files deleted by recent upgrade" you can ignore it.

8. Check to see if you have a file named `/etc/ambari-agent/conf.save` on each Agent host. If you do, rename it back. On each Agent host:

```
mv /etc/ambari-agent/conf.save /etc/ambari-agent/conf
```

9. Upgrade the Nagios addons package. On the Nagios host:

- RHEL/CentOS

```
yum upgrade hdp_mon_nagios_addons
```

- SLES

```
zypper up hdp_mon_nagios_addons
```

10. Start the Server and the Agents on all hosts. From the Server host:

```
ambari-server start
```

From each Agent host:

```
ambari-agent start
```

11. Open **Ambari Web**. Point your browser to `http://{your.ambari.server}:8080`



Important

You need to refresh your browser so that it loads the new version of the code. Hold the Shift key down while clicking the refresh button on the browser. If you have problems, clear your browser cache manually and restart Ambari Server.

Use the Admin name and password you have set up to log in.

12. Re-start Nagios and Ganglia services. In **Ambari Web**.

- a. Go to the **Services View** and select each service.
- b. Use the **Management Header** to stop and re-start the service.

12. Appendix: Upgrading Ambari Server to 1.2.5

This process upgrades Ambari Server. It does not change the underlying Hadoop Stack. This is a twelve step manual process.



Note

You must know the location of the Nagios server for Step 9. Use the **Services View-> Summary** panel to locate the host on which it is running.

1. Stop the Ambari Server and all Ambari Agents. From the Ambari Server host:

```
ambari-server stop
```

From each Ambari Agent host:

```
ambari-agent stop
```

2. Get the new Ambari bits. Use `wget` to fetch the repository file and replace the old repo file with the new repo file on every host.

- Fetch the new repo file:

For RHEL/CentOS 5/Oracle Linux 5

```
wget http://public-repo-1.hortonworks.com/ambari/centos5/1.x/updates/1.2.5.17/ambari.repo
```

For RHEL/CentOS 6/Oracle Linux 6

```
wget http://public-repo-1.hortonworks.com/ambari/centos6/1.x/updates/1.2.5.17/ambari.repo
```

For SLES 11

```
wget http://public-repo-1.hortonworks.com/ambari/suse11/1.x/updates/1.2.5.17/ambari.repo
```



Important

Check your current directory before you download the new repo file to make sure that there are no previous versions of the file. If you do not, and a previous version exists, the new download will be saved with a numeric extension such as `ambari.repo.1`. Make sure that the version you copy is the new version.

- Replace the old repo file with the new repo file.

For RHEL/CentOS 5/Oracle Linux 5

```
cp ambari.repo /etc/yum.repos.d/ambari.repo
```

For RHEL/CentOS 6/Oracle Linux 6

```
cp ambari.repo /etc/yum.repos.d/ambari.repo
```

For SLES 11

```
cp ambari.repo /etc/zypp/repos.d/ambari.repo
```



Note

If your cluster does not have access to the Internet, set up a local repository with this data before you continue. See [Configure the Local Repositories](#) for more information.

3. Upgrade Ambari Server. From the Ambari Server host:

- RHEL/CentOS

```
yum clean all
yum upgrade ambari-server-1.2.5.17 ambari-log4j-1.2.5.17
```

- SLES

```
zypper clean
zypper up ambari-server-1.2.5.17 ambari-log4j-1.2.5.17
```

4. Check for upgrade success:

- As the process runs, the console should produce output similar, although not identical, to this:

```
Setting up Upgrade Process
Resolving Dependencies
--> Running transaction check
---> Package ambari-agent.x86_64 0:1.2.2.3-1 will be updated
---> Package ambari-agent.x86_64 0:1.2.2.4-1 will be updated ...
---> Package ambari-agent.x86_64 0:1.2.2.5-1 will be an update ...
```

After the process is complete, check each host to make sure the new 1.2.4 files have been installed:

```
rpm -qa | grep ambari
```

- If the upgrade fails, the console should produce output similar to this:

```
Setting up Upgrade Process
No Packages marked for Update
```

5. Check to see if you have a folder named `/etc/ambari-server/conf.save`. If you do, rename it back:

```
mv /etc/ambari-server/conf.save /etc/ambari-server/conf
```

6. Upgrade the Ambari Server schema. From the Ambari Server host:

```
ambari-server upgrade
```

7. Upgrade the Ambari Agent on all hosts. From each Ambari Agent host:

- RHEL/CentOS

```
yum upgrade ambari-agent ambari-log4j
```

- SLES

```
zypper up ambari-agent ambari-log4j
```



Note

If you get a warning that begins "There are some running programs that use files deleted by recent upgrade" you can ignore it.

8. Check to see if you have a file named `/etc/ambari-agent/conf.save` on each Agent host. If you do, rename it back. On each Agent host:

```
mv /etc/ambari-agent/conf.save /etc/ambari-agent/conf
```

9. Upgrade the Nagios and Ganglia addons package and restart. On the Nagios/Ganglia host:

- RHEL/CentOS

```
yum upgrade hdp_mon_nagios_addons hdp_mon_ganglia_addons  
service httpd restart
```

- SLES

```
zypper up hdp_mon_nagios_addons hdp_mon_ganglia_addons  
service apache2 restart
```

10. Start the Server and the Agents on all hosts. From the Server host:

```
ambari-server start
```

From each Agent host:

```
ambari-agent start
```

11. Open **Ambari Web**. Point your browser to `http://{your.ambari.server}:8080`



Important

You need to refresh your browser so that it loads the new version of the code. Hold the Shift key down while clicking the refresh button on the browser. If you have problems, clear your browser cache manually and restart Ambari Server.

Use the Admin name and password you have set up to log in.

12. Re-start the Ganglia, Nagios, and MapReduce services. In **Ambari Web**.

- a. Go to the **Services View** and select each service.
- b. Use the **Management Header** to stop and re-start each service.

13. Appendix: Upgrading the HDP Stack to 1.3.3

The stack is the coordinated set of Hadoop components that you have installed. If you have a current instance of the 1.2.0/1.2.1 stack that was installed and managed by Ambari that you want to upgrade to the current 1.3.3 version of the stack and to also upgrade to the 1.2.5 version of Ambari Server and Agents, use the following instructions. This insures that the upgraded stack can still be managed by Ambari.

If you are upgrading from the 1.3.0 stack to the 1.3.3 stack, use [Section 5: Upgrading the Stack \(from 1.3.0 to 1.3.3\)](#), not Section 4: Upgrading the Stack (from 1.2.* to 1.3.3).



Note

If you have already upgraded to Ambari Server 1.2.5 and just want to upgrade the HDP stack, you can skip Sections 9.2 and 9.3.

13.1. Preparing for the Upgrade

Use the following steps to prepare your system for the upgrade.

1. If you are upgrading Ambari as well as the stack, you must know the location of the Nagios and Ganglia servers for that process. Use the **Services->Nagios/Ganglia->Summary** panel to locate the hosts on which they are running.
2. Use the **Services** view on the **Ambari Web UI** to stop all services, including MapReduce and all clients, running on HDFS. Do **not** stop HDFS yet.
3. Create the following logs and other files.

Because the upgrade to 1.3.3 includes a version upgrade of HDFS, creating these logs allows you to check the integrity of the file system post upgrade. While this is not absolutely necessary, doing so is strongly encouraged.

- a. Run `fsck` with the following flags and send the results to a log. The resulting file contains a complete block map of the file system. You use this log later to confirm the upgrade.

```
su $HDFS_USER
hadoop fsck / -files -blocks -locations > /tmp/dfs-old-fsck-1.log
```

where `$HDFS_USER` is the HDFS Service user (by default, `hdfs`).

- b. Capture the complete namespace of the filesystem. (The following command does a recursive listing of the root file system.)

```
su $HDFS_USER
hadoop dfs -lsr / > /tmp/dfs-old-lsr-1.log
```

where `$HDFS_USER` is the HDFS Service user (by default, `hdfs`).

- c. Create a list of all the DataNodes in the cluster.

```
su $HDFS_USER
hadoop dfsadmin -report > /tmp/dfs-old-report-1.log
```

where `$HDFS_USER` is the HDFS Service user (by default, `hdfs`).

- d. Optional: copy all or unrecoverable only data stored in HDFS to a local file system or to a backup instance of HDFS.
 - e. Optional: create the logs again and check to make sure the results are identical.
4. Save the namespace. You must be the HDFS service user to do this and you must put the cluster in Safe Mode.

```
hadoop dfsadmin -safemode enter
hadoop dfsadmin -saveNamespace
```

5. Copy the following checkpoint files into a backup directory. You can find the directory by using the **Services View** in the UI. Select the **HDFS** service, the **Configs** tab, in the **Namenode** section, look up the property **NameNode Directories**. It will be on your **NameNode** host.

- `dfs.name.dir/edits` // depending on your system, may not exist
- `dfs.name.dir/image/fsimage`

6. Stop HDFS. Make sure all services in the cluster are completely stopped.
7. If you are upgrading Hive, back up the Hive database.
8. Stop Ambari Server. On the Server host:

```
ambari-server stop
```

9. Stop Ambari Agents. On each host:

```
ambari-agent stop
```

13.2. Setting Up the Ambari Repository

This process prepares the updated repository.

1. Check to see if you have a `conf.save` directory for Ambari server and agents. If you do, move them to a back-up location:

```
mv /etc/ambari-server/conf.save/ /etc/ambari-server/conf.save.bak
```

```
mv /etc/ambari-agent/conf.save/ /etc/ambari-agent/conf.save.bak
```

2. Get the new Ambari bits. Use `wget` to fetch the repository file and replace the old repo file with the new repo file on every host.



Important

Check your current directory before you download the new repo file to make sure that there are no previous versions of the file. If you do not, and a previous version exists, the new download is saved with a numeric extension

such as `ambari.repo.1`. Make sure that the version you copy is the new version.

- For RHEL/CentOS/Oracle Linux 5

```
wget http://public-repo-1.hortonworks.com/ambari/centos5/1.x/updates/1.4.4.23/ambari.repo
cp ambari.repo /etc/yum.repos.d/ambari.repo
```

- For RHEL/CentOS/Oracle Linux 6

```
wget http://public-repo-1.hortonworks.com/ambari/centos6/1.x/updates/1.4.4.23/ambari.repo
cp ambari.repo /etc/yum.repos.d/ambari.repo
```

- For SLES 11

```
wget http://public-repo-1.hortonworks.com/ambari/suse11/1.x/updates/1.4.4.23/ambari.repo
cp ambari.repo /etc/zypp/repos.d/ambari.repo
```



Note

If your cluster does not have access to the Internet, you need to set up a local repository with this data before you continue. See [Configure the Local Repositories](#) for more information.

13.3. Upgrading Ambari

This process upgrades Ambari Server, Ambari Agents, Ganglia, and Nagios.

1. Upgrade Ambari Server. From the Ambari Server host:

- RHEL/CentOS/Oracle Linux

```
yum clean all
yum upgrade ambari-server ambari-log4j
```

- SLES

```
zypper clean
zypper up ambari-server ambari-log4j
```

2. Check for upgrade success:

- As the process runs, the console should produce output similar, although not identical, to this:

```
Setting up Upgrade Process
Resolving Dependencies
--> Running transaction check
---> Package ambari-agent.x86_64 0:1.2.2.3-1 will be updated
---> Package ambari-agent.x86_64 0:1.2.2.4-1 will be updated ...
---> Package ambari-agent.x86_64 0:1.2.2.5-1 will be an update ...
```

After the process is complete, check each host to make sure the new 1.2.5 files have been installed:

```
rpm -qa | grep ambari
```

- If the upgrade fails, the console should produce output similar to this:

```
Setting up Upgrade Process  
No Packages marked for Update
```

3. Check to see if you have a folder named `/etc/ambari-server/conf.save`. If you do, rename it back:

```
mv /etc/ambari-server/conf.save /etc/ambari-server/conf
```

4. Upgrade the Ambari Server schema. From the Ambari Server host:

```
ambari-server upgrade
```

5. Upgrade the Ambari Agent on all hosts. From each Ambari Agent host:

- RHEL/CentOS/Oracle Linux

```
yum upgrade ambari-agent ambari-log4j
```

- SLES

```
zypper up ambari-agent ambari-log4j
```



Note

If you get a warning that begins "There are some running programs that use files deleted by recent upgrade" you can ignore it.

6. Check to see if you have a folder named `/etc/ambari-agent/conf.save` on each Agent host. If you do, rename it back. On each Agent host:

```
mv /etc/ambari-agent/conf.save /etc/ambari-agent/conf
```

7. Upgrade Ganglia and Nagios:

- Upgrade Ganglia Server. From the Ganglia Server host:

- RHEL/CentOS/Oracle Linux

```
yum upgrade ganglia-gmond ganglia-gmetad libganglia  
yum erase gweb hdp_mon_ganglia_addons  
yum install ganglia-web
```

- SLES

```
zypper up ganglia-gmond ganglia-gmetad libganglia  
zypper remove gweb hdp_mon_ganglia_addons  
zypper install ganglia-web
```

- Upgrade Ganglia Monitor. From every host that has Ganglia Monitor installed:

- RHEL/CentOS/Oracle Linux

```
yum upgrade ganglia-gmond libganglia
```


- SLES

```
zypper up ganglia-gmond libganglia
```

- Upgrade Nagios. From the Nagios Server host:

- RHEL/CentOS/Oracle Linux

```
yum upgrade nagios
yum upgrade hdp_mon_nagios_addons
yum erase nagios-plugins-1.4.15
yum install nagios-plugins-1.4.9
```

The 1.4.9 version of the plugin may already be installed. In this case, the second step is a no-op.

- SLES

```
zypper up nagios
zypper up hdp_mon_nagios_addons
zypper remove nagios-plugins-1.4.15
zypper install nagios-plugins-1.4.9
```

The 1.4.9 version of the plugin may already be installed. In this case, the second step is a no-op.

13.4. Upgrading the Stack (from 1.2.* to 1.3.3)

1. Update the stack version in the Server database, depending on if you are using a local repository:

```
ambari-server upgradestack HDP-1.3.3
```

2. Upgrade the HDP repository on all hosts and replace the old repo file with the new file:



Important

The file you download is named `hdp.repo`. To function properly in the system, it must be named `HDP.repo`. Once you have completed the "mv" of the new repo file to the `repos.d` folder, make sure there is no file named `hdp.repo` anywhere in your `repos.d` folder.

- For RHEL/CentOS/Oracle Linux 5

```
wget http://public-repo-1.hortonworks.com/HDP/centos5/1.x/updates/1.3.3.0/hdp.repo
mv hdp.repo /etc/yum.repos.d/HDP.repo
```

- For RHEL/CentOS/Oracle Linux 6

```
wget http://public-repo-1.hortonworks.com/HDP/centos6/1.x/updates/1.3.3.0/hdp.repo
mv hdp.repo /etc/yum.repos.d/HDP.repo
```

- For SLES 11

```
wget http://public-repo-1.hortonworks.com/HDP/suse11/1.x/updates/1.3.3.0/hdp.repo
mv hdp.repo /etc/zypp/repos.d/HDP.repo
```

3. Upgrade the stack on all Agent hosts. Skip any components your installation does not use:

- For RHEL/CentOS/Oracle Linux

a. Upgrade the following components:

```
yum upgrade "collectd*" "epel-release*" "gccxml*" "pig*" "hadoop*"
"sqoop*" "zookeeper*" "hbase*" "hive*" "hcatalog*" "webhcat-tar*"
hdp_mon_nagios_addons
```

b. Check to see that the components in that list are upgraded.

```
yum list installed | grep HDP-$old-stack-version-number
```

None of the components from that list should appear in the returned list.

c. Upgrade Oozie, if you are using Oozie:

```
rpm -e --nopostun oozie-$old_version_number
yum install oozie
```

You can get the value of `$old_version_number` from the output of the previous step.

d. Upgrade Oozie Client:

```
yum upgrade oozie-client
```

e. Upgrade ZooKeeper:

i. Check to see if ZooKeeper needs upgrading.

```
yum list installed | grep zookeeper
```

If the displayed version number is **not** 3.4.5.1.3.2.0, you need to upgrade.

ii. Because HBase depends on ZooKeeper, deleting the current version of ZooKeeper automatically deletes the current version of HBase. It must be re-installed. Check to see if HBase is currently installed.

```
yum list installed | grep hbase
```

iii. Delete the current version of ZooKeeper.

```
yum erase zookeeper
```

iv. Install ZooKeeper

```
yum install zookeeper
```

v. If you need to, re-install HBase.

```
yum install hbase
```

vi. Check to see if all components have been upgraded:

```
yum list installed | grep HDP-$old-stack-version-number
```

The only non-upgraded component you may see in this list is `extjs`, which does not need to be upgraded.

- For SLES

a. Upgrade the following components:

```
zypper up collectd epel-release* gccxml* pig* hadoop* sqoop* hive*  
hcatalog* webhcat-tar* hdp_mon_nagios_addons*  
yast --update hadoop hcatalog hive
```

b. Upgrade ZooKeeper and HBase:

```
zypper update zookeeper-3.4.5.1.3.2.0  
zypper remove zookeeper  
zypper se -s zookeeper
```

You should see ZooKeeper v3.4.5.1.3.2.0 in the output.

Install ZooKeeper v3.4.5.1.3.2.0:

```
zypper install zookeeper-3.4.5.1.3.2.0
```

This command also uninstalls HBase. Now use the following commands to install HBase:

```
zypper install hbase-0.94.6.1.3.2.0  
zypper update hbase
```

c. Upgrade Oozie:

```
rpm -e --nopostun oozie-$old_version_number  
zypper update oozie-3.3.2.1.3.2.0  
zypper remove oozie  
zypper se -s oozie
```

You should see Oozie v3.3.2.1.3.2.0 in the output.

Install Oozie v3.3.2.1.3.2.0:

```
zypper install oozie-3.3.2.1.3.2.0
```

4. Start the Ambari Server. On the Server host:

```
ambari-server start
```

5. Start each Ambari Agent. On all Agent hosts:

```
ambari-agent start
```

6. Because the file system version has now changed you must start the NameNode manually. On the NameNode host:

```
sudo su -l $HDFS_USER -c "/usr/lib/hadoop/bin/hadoop-daemon.sh start  
namenode -upgrade"
```

Depending on the size of your system, this step may take up to 10 minutes.

7. Track the status of the upgrade:

```
hadoop dfsadmin -upgradeProgress status
```

Continue tracking until you see:

```
Upgrade for version -44 has been completed.  
Upgrade is not finalized.
```



Note

You finalize the upgrade later.

8. Open the Ambari Web GUI. If you have continued to run the Ambari Web GUI, do a hard reset on your browser. Use **Services View** to start the HDFS service. This starts the SecondaryNameNode and the DataNodes.

9. After the DataNodes are started, HDFS exits safemode. To monitor the status:

```
hadoop dfsadmin -safemode get
```

Depending on the size of your system, this may take up to 10 minutes or so. When HDFS exits safemode, this is displayed as a response to the command:

```
Safe mode is OFF
```

10. Make sure that the HDFS upgrade was successful. Go through steps 2 and 3 in [Section 9.1](#) to create new versions of the logs and reports. Substitute "new" for "old" in the file names as necessary

11. Compare the old and new versions of the following:

- `dfs-old-fsck-1.log` versus `dfs-new-fsck-1.log`.

The files should be identical unless the `hadoop fsck` reporting format has changed in the new version.

- `dfs-old-lsr-1.log` versus `dfs-new-lsr-1.log`.

The files should be identical unless the the format of `hadoop fs -lsr` reporting or the data structures have changed in the new version.

- `dfs-old-report-1.log` versus `fs-new-report-1.log`

Make sure all DataNodes previously belonging to the cluster are up and running.

12. Use the Ambari Web **Services** view-> Services Navigation->**Start All** to start services back up.

13. The upgrade is now fully functional but not yet finalized. Using the `finalize` command removes the previous version of the NameNode and DataNode's storage directories.



Important

After the upgrade is finalized, the system cannot be rolled back. Usually this step is not taken until a thorough testing of the upgrade has been performed.

The upgrade must be finalized before another upgrade can be performed.



Note

Directories used by Hadoop 1 services set in `/etc/hadoop/conf/taskcontroller.cfg` are not automatically deleted after upgrade. Administrators can choose to delete these directories after the upgrade.

To finalize the upgrade:

```
sudo su -l $HDFS_USER -c "hadoop dfsadmin -finalizeUpgrade"
```

where `$HDFS_USER` is the HDFS Service user (by default, `hdfs`).

13.5. Upgrading the Stack (from 1.3.0 to 1.3.3)

1. Update the stack version in the Server database, depending on if you are using a local repository:

```
ambari-server upgradestack HDP-1.3.3
```

2. Upgrade the HDP repository on all hosts and replace the old repo file with the new file:



Important

The file you download is named `hdp.repo`. To function properly in the system, it must be named `HDP.repo`. After you have completed the "mv" of the new repo file to the `repos.d` folder, make sure there is no file named `hdp.repo` anywhere in your `repos.d` folder.

- For RHEL/CentOS/Oracle Linux 5

```
wget http://public-repo-1.hortonworks.com/HDP/centos5/1.x/updates/1.3.3.0/hdp.repo
mv hdp.repo /etc/yum.repos.d/HDP.repo
```

- For RHEL/CentOS/Oracle Linux 6

```
wget http://public-repo-1.hortonworks.com/HDP/centos6/1.x/updates/1.3.3.0/hdp.repo
mv hdp.repo /etc/yum.repos.d/HDP.repo
```

- For SLES 11

```
wget http://public-repo-1.hortonworks.com/HDP/suse11/1.x/updates/1.3.3.0/hdp.repo
mv hdp.repo /etc/zypp/repos.d/HDP.repo
```

- Upgrade the stack on all Agent hosts. Skip any components your installation does not use:

- For RHEL/CentOS/Oracle Linux
 - Upgrade the following components:

```
yum upgrade "collectd*" "epel-release*" "gccxml*" "pig*" "hadoop*"
"sqoop*" "zookeeper*" "hbase*" "hive*" "hcatalog*" "webhcat-tar*"
"oozie*" hdp_mon_nagios_addons
```

- Check to see if those components have been upgraded:

```
yum list installed | grep HDP-$old-stack-version-number
```

The only non-upgraded component you may see in this list is `extjs`, which does not need to be upgraded.

- For SLES
 - Upgrade the following components:

```
zypper up collectd gccxml* pig* hadoop* sqoop* hive* hcatalog* webhcat-
tar* zookeeper* oozie* hbase* hdp_mon_nagios_addons*
yast --update hadoop hcatalog hive
```

- Start the Ambari Server. On the Server host:

```
ambari-server start
```

- Start each Ambari Agent. On all Agent hosts:

```
ambari-agent start
```

- Because the file system version has now changed you must start the NameNode manually. On the NameNode host:

```
sudo su -l $HDFS_USER -c "/usr/lib/hadoop/bin/hadoop-daemon.sh start
namenode -upgrade"
```

Depending on the size of your system, this step may take up to 10 minutes.

- Track the status of the upgrade:

```
hadoop dfsadmin -upgradeProgress status
```

Continue tracking until you see:

```
Upgrade for version -44 has been completed.
Upgrade is not finalized.
```



Note

You finalize the upgrade later.

- Open the Ambari Web GUI. If you have continued to run the Ambari Web GUI, do a hard reset on your browser. Use **Services View** to start the HDFS service. This starts the SecondaryNameNode and the DataNodes.

9. After the DataNodes are started, HDFS exits safemode. To monitor the status:

```
hadoop dfsadmin -safemode get
```

Depending on the size of your system, this may take up to 10 minutes or so. When HDFS exits safemode, this is displayed as a response to the command:

```
Safe mode is OFF
```

10. Make sure that the HDFS upgrade succeeded. Go through steps 2 and 3 in [Section 9.1](#) to create new versions of the logs and reports. Substitute "new" for "old" in the file names as necessary

11. Compare the old and new versions of the following files:

- `dfs-old-fsck-1.log` versus `dfs-new-fsck-1.log`.

The files should be identical unless the `hadoop fsck` reporting format has changed in the new version.

- `dfs-old-lsr-1.log` versus `dfs-new-lsr-1.log`.

The files should be identical unless the the format of `hadoop fs -lsr` reporting or the data structures have changed in the new version.

- `dfs-old-report-1.log` versus `dfs-new-report-1.log`

Make sure all DataNodes previously belonging to the cluster are up and running.

12. Use the Ambari Web **Services** view-> Services Navigation->**Start All** to start services back up.

13. The upgrade is now fully functional but not yet finalized. Using the `finalize` command removes the previous version of the NameNode and DataNode's storage directories.



Important

Once the upgrade is finalized, the system cannot be rolled back. Usually this step is not taken until a thorough testing of the upgrade has been performed.

The upgrade must be finalized, however, before another upgrade can be performed.



Note

Directories used by Hadoop 1 services set in `/etc/hadoop/conf/taskcontroller.cfg` are not automatically deleted after upgrade. Administrators can choose to delete these directories after the upgrade.

To finalize the upgrade:

```
sudo su -l $HDFS_USER -c "hadoop dfsadmin -finalizeUpgrade"
```

where `$HDFS_USER` is the HDFS Service user (by default, `hdfs`).

14. Appendix: Configuring Ports

The tables below specify which ports must be opened for which ecosystem components to communicate with each other. Make sure the appropriate ports are opened before you install Hadoop.

- [HDFS Ports](#)
- [MapReduce Ports](#)
- [Hive Ports](#)
- [HBase Ports](#)
- [ZooKeeper Ports](#)
- [WebHCat Port](#)
- [Ganglia Ports](#)
- [MySQL Port](#)
- [Ambari Ports](#)
- [Nagios Port](#)

14.1. HDFS Ports

The following table lists the default ports used by the various HDFS services.

Table 14.1. HDFS Ports

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
NameNode WebUI	Master Node hosts (NameNode and any back-up NameNodes)	50070	http	Web UI to look at current status of HDFS, explore file system	Yes (Typically admins, Dev/ Support teams)	<code>dfs.http.address</code>
		50470	https	Secure http service		<code>dfs.https.address</code>
NameNode metadata service	Master Node hosts (NameNode and any back-up NameNodes)	8020/9000	IPC	File system metadata operations	Yes (All clients who directly need to interact with the HDFS)	Embedded in URI specified by <code>fs.default.name</code>
DataNode	All Slave Node hosts	50075	http	DataNode WebUI to access the status, logs etc.	Yes (Typically admins, Dev/ Support teams)	<code>dfs.datanode.http.address</code>
		50475	https	Secure http service		<code>dfs.datanode.https.address</code>
		50010		Data transfer		<code>dfs.datanode.address</code>
		50020	IPC	Metadata operations	No	<code>dfs.datanode.ipc.address</code>

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
Secondary NameNode	Secondary NameNode and any backup Secondary NameNode hosts	50090	http	Checkpoint for NameNode metadata	No	dfs.secondary.http.address

14.2. MapReduce Ports

The following table lists the default ports used by the various MapReduce services.

Table 14.2. MapReduce Ports

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
JobTracker WebUI	Master Node hosts (JobTracker Node and any back-up Job-Tracker node)	50030	http	Web UI for JobTracker	Yes	mapred.job.tracker.http.add
JobTracker	Master Node hosts (JobTracker Node)	8021	IPC	For job submissions	Yes (All clients who need to submit the MapReduce jobs including Hive, Hive server, Pig)	Embedded in URI specified by mapred.job.tracker
TaskTracker Web UI and Shuffle	All Slave Node hosts	50060	http	DataNode Web UI to access status, logs, etc.	Yes (Typically admins, Dev/ Support teams)	mapred.task.tracker.http.ad
History Server WebUI		51111	http	Web UI for Job History	Yes	mapreduce.history.server.ht

14.3. Hive Ports

The following table lists the default ports used by the various Hive services.



Note

Neither of these services is used in a standard HDP installation.

Table 14.3. Hive Ports

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
Hive Server2	Hive Server machine (Usually a utility machine)	10000	thrift	Service for programatically (Thrift/JDBC) connecting to Hive	Yes (Clients who need to connect to Hive either programatically or through UI SQL tools that use JDBC)	ENV Variable HIVE_PORT

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
Hive Metastore		9083	thrift	Service for accessing metadata about Hive tables and partitions.*	Yes (Clients that run Hive, Pig and potentially M/R jobs that use HCatalog)	hive.metastore.uris

* To change the metastore port, use this hive command: `hive --service metastore -p port_number`

14.4. HBase Ports

The following table lists the default ports used by the various HBase services.

Table 14.4. HBase Ports

Service	Servers	Default Ports Used	Protocol
HMaster	Master Node hosts (HBase Master Node and any back-up HBase Master node)	60000	
HMaster Info Web UI	Master Node hosts (HBase master Node and back up HBase Master node if any)	60010	http
Region Server	All Slave Node hosts	60020	
Region Server	All Slave Node hosts	60030	http
HBase REST Server (optional)	All REST Servers	8080	http
HBase REST Server Web UI(optional)	All REST Servers	8085	http
HBase Thrift Server (optional)	All Thrift Servers	9090	
HBase Thrift Server Web UI (optional)	All Thrift Servers	9095	

14.5. ZooKeeper Ports

The following table lists the default ports used by the various ZooKeeper services.

Table 14.5. HBase Ports

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
ZooKeeper Server	All ZooKeeper Node hosts	2888		Port used by ZooKeeper peers to talk to each other. See	No	hbase.zookeeper.peerport

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
				here for more information.		
ZooKeeper Server	All ZooKeeper Node hosts	3888		Port used by ZooKeeper peers to talk to each other. See here for more information.	No	<code>hbase.zookeeper.leaderport</code>
ZooKeeper Server	All ZooKeeper Hosts	2181		Property from ZooKeeper's config <code>zoo.cfg</code> . The port at which the clients will connect.	Yes	<code>hbase.zookeeper.property.cl</code>

14.6. WebHCat Port

The following table lists the default port used by the WebHCat service.

Table 14.6. WebHCat Port

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
WebHCat Server	Any utility machine	50111	http	Web API on top of HCatalog and other Hadoop services	Yes	<code>templeton.port</code>

14.7. Ganglia Ports

The following table lists the default ports used by the various Ganglia services.

Table 14.7. Ganglia Ports

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
Ganglia Server	Ganglia server host	8660/61/62/63		For metric (gmond) collectors	No	
Ganglia Monitor	All Slave Node hosts	8660		For monitoring (gmond) agents	No	
Ganglia Server	Ganglia server host	8651		For ganglia gmetad		
Ganglia Web	Ganglia server host		http ^a			

^aSee [Optional: Set Up HTTPS for Ganglia](#) for instructions on enabling HTTPS.

14.8. MySQL Port

The following table lists the default port used by the MySQL service.

Table 14.8. MySQL Port

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
MySQL	MySQL database server host	3306				

14.9. Ambari Ports

The following table lists the default ports used by Ambari.

Table 14.9. Ambari Web

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
Ambari Server	Ambari Server host	8080 ^a	http ^b	Interface to Ambari Web and Ambari REST API	No	
Ambari Server	Ambari Server host	8440	https	Handshake Port for Ambari Agents to Ambari Server	No	
Ambari Server	Ambari Server host	8441	https	Registration and Heartbeat Port for Ambari Agents to Ambari Server	No	

^aSee [Optional: Change the Ambari Server Port](#) for instructions on changing the default port.

^bSee [Optional: Set Up HTTPS for Ambari Web](#) for instructions on enabling HTTPS.

14.10. Nagios Ports

The following table lists the default port used by Ambari Web.

Table 14.10. Ambari Web

Service	Servers	Default Ports Used	Protocol	Description	Need End User Access?	Configuration Parameters
Nagios Server	Nagios server host	80	http ^a	Nagios Web UI	No	

^aSee [Optional: Set Up HTTPS for Nagios](#) for instructions on enabling HTTPS.

15. Configuring RHEL HA for Hadoop 1.x

Ambari supports High Availability of components such as NameNode or JobTracker in a HDP 1.x cluster running RHEL HA. After installing NameNode monitoring components on hosts in an HA cluster, as described in [HDP System Administration](#), configure Ambari to reassign any component on a failover host in the cluster, using the `host_relocate_component.py` script.

For example, if the host for the primary NameNode or JobTracker component fails, Ambari reassigns the primary NameNode or JobTracker component to the configured failover host, when you start or restart Ambari server.



Note

Make sure that NameNode is installed. For successful reassignment, a component must be in one of the following states:

- Maintenance (started or stopped)
- Unknown

15.1. Deploy the scripts

While the Ambari server and ambari agents are running on each host:

1. Download `relocate_host_component.py` from `/var/lib/ambari-server/resources/scripts` on the Ambari server to `/usr/bin/` on each failover host.
2. Download `hadoop.sh` from `/var/lib/ambari-server/resources/scripts` on the Ambari server and replace `hadoop.sh` in `/usr/share/cluster/` on each failover host.

15.2. Configure Ambari properties across the HA cluster

To enable Ambari to run `relocate_host_component.py`, edit the cluster configuration file on each failover host in the HA cluster, using a text editor.

In `/etc/cluster/cluster.conf`, set values for each of the following properties:

- `<server>=<ambari-hostname / ip>`
- `<port>=<8080>`
- `<protocol>=<http / https>`
- `<user>=<admin>`
- `<password>=<admin>`
- `<cluster>=<cluster-name>`

- `<output>=</var/log/ambari_relocate.log>`

For example, the Hadoop daemon section of `cluster.conf` on the NameNode localhost in an HA cluster will look like:

```
<hadoop
__independent_subtree="1" __max_restarts="10" __restart_expire_time="600"
name="NameNode Process"
daemon="namenode" boottime="10000" probetime="10000" stoptime="10000" url=
"http://10.0.0.30:50070/dfshealth.jsp"
pid="/var/run/hadoop/hdfs/hadoop-hdfs-namenode.pid" path="/"

ambariproperties="server=localhost,port=8080,protocol=http,user=admin,
password=admin,cluster=c1,output=/var/log/ambari_relocate.log"

/>
```

The `relocate_host_component.py` script reassigns components on failover of any host in the HA cluster, when you start or restart Ambari server.

15.3. Troubleshooting RHEL HA

1. Review errors in `/var/log/messages/`.
2. If the following error message appears:

```
abrtcd: Executable '/usr/bin/relocate_resources.py' doesn't belong to any
package and ProcessUnpackaged is set to 'no'
```

In `/etc/abrt/abrt-action-save-package-data.conf`, set `ProcessUnpackaged=Yes`.

3. If the scripts return Error status=exit code 3, make sure the following are true:
 - The ambari agent on the failover host is running.
 - Failover did not result from STOP HDFS or STOP NN/JT, using Ambari.

The following table lists and describes parameters for `relocate_host_components.py`.

Table 15.1. Parameter Options for `relocate_host_components.py`

Parameter	Value	Example	Description
-h,	na	-help	Display all parameter options.
-v,	na	-verbose	Increases output verbosity.
-s,	SERVER_HOSTNAME	-s,--host=SERVER_HOSTNAME	Ambari server host name (FQDN)
-p,	SERVER_PORT,	-port=SERVER_PORT	Ambari server port. [default: 8080]
-r,	PROTOCOL,	-protocol=PROTOCOL	Protocol for communicating with Ambari server (http/https) [default: http]
-c,	CLUSTER_NAME,	-cluster-name=CLUSTER_NAME	Ambari cluster to operate on.
-e,	SERVICE_NAME,	-service-name=SERVICE_NAME	Ambari Service to which the component belongs.
-m,	COMPONENT_NAME	-component-name=COMPONENT_NAME	Ambari Service Component to operate on.

Parameter	Value	Example	Description
-n,	NEW_HOSTNAME,	-new-host=NEW_HOSTNAME	New host to relocate the component to.
-a,	ACTION,	-action=ACTION	Script action. [default: relocate]
-o,	FILE,	-output-file=FILE	Output file. [default: /temp/ambari_reinstall_probe.out]
-u,	USERNAME,	-username=USERNAME	Ambari server admin user. [default: admin]
-w,	PASSWORD,	-password=PASSWORD	Ambari server admin password.
-d,	COMPONENT_NAME,	start-component	Start the component after reassignment.

Part IV. Additional Tasks with Ambari

This section describes additional tasks you may need to do with your Hadoop installation. It includes:

- [Installing Ambari Agents Manually](#)
- [Using Custom Hostnames](#)
- [Upgrading Operating Systems on an Ambari-based Hadoop Installation](#)
- [Moving the Ambari Server](#)
- [Using Non-Default Databases](#)
- [Setting Up Kerberos for Use with Ambari](#)

Hortonworks Data Platform

Installing Hadoop Using Apache Ambari

(Aug 1, 2014)

16. Appendix: Installing Ambari Agents Manually

In some situations you may decide you do not want to have the Ambari Install Wizard install and configure the Agent software on your cluster hosts automatically. In this case you can install the software manually.

Before you begin: on every host in your cluster download the HDP repository as described in [Set Up the Bits](#).

16.1. RHEL/CentOS/Oracle Linux 5.x and 6.x

1. Install the Ambari Agent

```
yum install ambari-agent
```

2. Using a text editor, Configure the Ambari Agent by editing the `ambari-agent.ini` file. For example:

```
vi /etc/ambari-agent/conf/ambari-agent.ini

[server]
hostname={your.ambari.server.hostname}
url_port=8440
secured_url_port=8441
```

3. Start the agent. The agent registers with the Server on start.

```
ambari-agent start
```

16.2. SLES

1. Install the Ambari Agent

```
zypper install ambari-agent
```

2. Configure the Ambari Agent by editing the `ambari-agent.ini` file.

```
vi /etc/ambari-agent/conf/ambari-agent.ini

[server]
hostname={your.ambari.server.hostname}
url_port=8440
secured_url_port=8441
```

3. Start the agent. The agent registers with the Server on start.

```
ambari-agent start
```

17. Appendix: Using Custom Hostnames

You can customize the agent registration hostname and the public hostname used for each host in Ambari.

To customize the name of each host in your cluster:

1. On the **Install Options** screen, select **Perform Manual Registration** for Ambari Agents.
2. Install the Agents manually, as described in [Installing Ambari Agents Manually](#).
3. To echo the customized name of the host to which the Ambari agent registers, for every host, create a script like the following example, named `/var/lib/ambari-agent/hostname.sh`.

```
#!/bin/sh
echo <ambari_hostname>
```

4. On every host, using a text editor, open `/etc/ambari-agent/conf/ambari-agent.ini`.

Add the following line to the agent section:

```
hostname_script=/var/lib/ambari-agent/hostname.sh
```

(where `/var/lib/ambari-agent/hostname.sh` is the name of your custom echo script)

5. To generate a public hostname for every host, create a script like the following example, named `/var/lib/ambari-agent/public_hostname.sh` to show the name for that host in the UI.

```
#!/bin/bash
hostname <-f>
```

6. On every host, using a text editor, open `/etc/ambari-agent/conf/ambari-agent.ini`.

Add the following line to the agent section:

```
public_hostname_script=/var/lib/ambari-agent/public_hostname.sh
```

7. Add the hostnames to `/etc/hosts` on all nodes.

18. Appendix: Upgrading Operating Systems on an Ambari-based Hadoop Installation

Ambari requires specific versions of the files for components that it uses. There are three steps you should take to make sure that these versions continue to be available:

- Disable automatic OS updates
- Do not update any HDP components such as MySQL, Ganglia, etc.
- If you must perform an OS update, do a manual kernel update only.

19. Appendix: Moving the Ambari Server

Use the following instructions to transfer the Ambari Server to a new host.



Note

These steps describe moving the Ambari Server when it uses the default PostgreSQL database. If you are using a non-default database for Ambari (such as Oracle), adjust the database backup, restore and stop/start procedures to match that database.

1. [Back up all current data from the original Ambari Server and MapReduce databases.](#)
2. [Update all Agents to point to the new Ambari Server.](#)
3. [Install the Server on a new host and populate databases with information from original Server.](#)

19.1. Back up Current Data

1. Stop the original Ambari Server.

```
ambari-server stop
```

2. Create a directory to hold the database backups.

```
cd /tmp
mkdir dbdumps
cd dbdumps/
```

3. Create the database backups.

```
pg_dump -U $AMBARI-SERVER-USERNAME ambari > ambari.sql Password: $AMBARI-
SERVER-PASSWORD
pg_dump -U $MAPRED-USERNAME ambarirca > ambarirca.sql Password: $MAPRED-
PASSWORD
```

Substitute the values you set up at installation for `$AMBARI-SERVER-USERNAME`, `$MAPRED-USERNAME`, `$AMBARI-SERVER-PASSWORD`, and `$MAPRED-PASSWORD`. Defaults are `ambari-server/bigdata` and `mapred/mapred`.

19.2. Update Agents

1. On each Agent node, stop the Agent.

```
ambari-agent stop
```

2. Remove old Agent certificates.

```
rm /var/lib/ambari-agent/keys/*
```

3. Using a text editor, edit `/etc/ambari-agent/conf/ambari-agent.ini` to point to the new host.

```
[server]
hostname=$NEW FULLY.QUALIFIED.DOMAIN.NAME
url_port=8440
secured_url_port=8441
```

19.3. Install the New Server and Populate the Databases

1. On the new host, install the Server as described in [Running the Installer](#), Sections 2.1 and 2.2.

2. Stop the Server so that you can copy the old database data to the new Server.

```
ambari-server stop
```

3. Restart the PostgreSQL instance.

```
service postgresql restart
```

4. Open the PostgreSQL interactive terminal.

```
su - postgres
psql
```

5. Using the interactive terminal, drop the databases created by the fresh install.

```
drop database ambari;
drop database ambarirca;
```

6. Check to make sure the databases have been dropped.

```
/list
```

The databases should not be listed.

7. Create new databases to hold the transferred data.

```
create database ambari;
create database ambarirca;
```

8. Exit the interactive terminal

```
^d
```

9. Copy the saved data from [Back up Current Data](#) to the new Server.

```
cd /tmp
scp -i <ssh-key> root@<original-server>/tmp/dbdumps/*.sql/tmp
(Note: compress/transfer/uncompress as needed from source to dest)
psql -d ambari -f /tmp/ambari.sql
psql -d ambarirca -f /tmp/ambarirca.sql
```

- 10 Start the new Server.

```
<exit to root>
ambari-server start
```

11. On each Agent host, start the Agent.

```
ambari-agent start
```

12. Open Ambari Web. Point your compatible browser to:

```
<new_Ambari_Server>:8080
```

13. Go to **Services** -> **MapReduce** and use the Management Header to **Stop** and **Start** the MapReduce service.

14. Start other services as necessary.

The new Server is ready to use.

20. Appendix: Using Non-Default Databases

Use the following instructions to prepare a non-default database for Hive/HCatalog, Oozie, or Ambari. You **must** complete these instructions before you setup the Ambari Server by running `ambari-server setup`.

20.1. Hive/HCatalog

1. On the Hive Metastore machine, install the appropriate JDBC .jar file.

- For **Oracle**:

a. Download the Oracle JDBC (OJDBC) driver from <http://www.oracle.com/technetwork/database/features/jdbc/index-091264.html>.

Select Oracle Database 11g Release 2 - ojdbc6.jar .

b. Copy the .jar file to the Java share directory.

```
cp ojdbc6.jar /usr/share/java
```

c. Make sure the .jar file has the appropriate permissions - 644.

- For **MySQL**:

a. Install the connector.

- RHEL/CentOS/Oracle Linux

```
yum install mysql-connector-java-5.0.8-1
```

- SLES

```
zypper install mysql-connector-java-5.0.8-1
```

b. Confirm that the MySQL .jar file is in the Java share directory

```
ls /usr/share/java/mysql-connector-java.jar
```

c. Make sure the .jar file has the appropriate permissions - 644.

2. On the Ambari Server host, install the appropriate JDBC .jar file.

- For **Oracle**:

a. Download the Oracle JDBC (OJDBC) driver from <http://www.oracle.com/technetwork/database/features/jdbc/index-091264.html>.

Select Oracle Database 11g Release 2 - ojdbc6.jar .

b. Copy the .jar file to the Java share directory.

```
cp ojdbc6.jar /var/lib/ambari-server/resources
```

c. Make sure the .jar file has the appropriate permissions - 644.

- For **MySQL**:

a. Download the mysql connector driver from the host on which you installed mysql-connector-java.

b. Copy the .jar file to the Java share directory.

```
cp mysql-connector-java.jar /var/lib/ambari-server/resources
```

c. Make sure the .jar file has the appropriate permissions - 644.

3. Create a user for Hive and grant it permissions.

- For **Oracle**, create the Hive user and grant it database permissions.

```
# sqlplus sys/root as sysdba
SQL> CREATE USER $HIVEUSER IDENTIFIED BY $HIVEPASSWORD;
SQL> GRANT SELECT_CATALOG_ROLE TO $HIVEUSER;
SQL> GRANT CONNECT, RESOURCE TO $HIVEUSER;
SQL> QUIT;
```

Where `$HIVEUSER` is the Hive user name and `$HIVEPASSWORD` is the Hive user password.

- For **MySQL**, create the Hive user and grant it database permissions.

```
# mysql -u root -p
mysql> CREATE USER '$HIVEUSER'@'%' IDENTIFIED BY '$HIVEPASSWORD';
mysql> GRANT ALL PRIVILEGES ON *.* TO '$HIVEUSER'@'%';
mysql> flush privileges;
```

Where `$HIVEUSER` is the Hive user name and `$HIVEPASSWORD` is the Hive user password.

4. For **Oracle only** Load the Hive Metastore Schema.

- The Hive Metastore database schema must be pre-loaded into your Oracle database using the schema script.

```
sqlplus $HIVEUSER/$HIVEPASSWORD < hive-schema-0.12.0.oracle.sql
```

The file `hive-schema-0.12.0.oracle.sql` is found in the `/var/lib/ambari-server/resources/` directory of the Ambari Server machine, once you have completed the [Set Up the Bits](#) step in the install process.

20.1.1. Troubleshooting Hive/HCatalog

Use these entries to help you troubleshoot any issues you might have installing Hive/HCatalog with non-default databases.

20.1.1.1. Problem: Hive Metastore Install Fails Using Oracle

Check the install log:

```
cp /usr/share/java/${jdbc_jar_name} ${target}] has failures: true
```

The Oracle JDBC .jar file cannot be found.

20.1.1.1.1. Soution

Make sure the file is in the appropriate directory on the Hive Metastore server and click **Retry**.

20.1.1.2. Problem: Install Warning when "Hive Check Execute" Fails Using Oracle

Check the install log:

```
java.sql.SQLException: ORA-01754:  
a table may contain only one column of type LONG
```

The Hive Metastore schema was not properly loaded into the database.

20.1.1.2.1. Soution

Complete the install with the warning. Check your database to confirm the Hive Metastore schema is loaded. Once in the Ambari Web GUI, browse to **Services > Hive/HCat**. Use the Management Header to re-run the smoke test (**Maintenance ->Run Smoke Test**) to check that the schema is correctly in place.

20.1.1.3. Problem: Hive Check Execute may fail after completing an Ambari upgrade to version 1.4.2

For secure and non-secure clusters, with Hive security authorization enabled, the Hive service check may fail. Hive security authorization may not be configured properly.

20.1.1.3.1. Solution

Two workarounds are possible. Using Ambari Web, in **Hive Configs Advanced**:

- Disable Hive security authorization, by setting the hive.security.authorization.enabled value to false.
- or
- Properly configure Hive security authorization. For example, set the following properties:

Table 20.1. Hive Security Authorization Settings

Property	Value
hive.security.authorization.manager	org.apache.hadoop.hive.ql.security.authorization.StorageBasedAuthorizationP

Property	Value
hive.security.metastore.authorization.manager	org.apache.hadoop.hive.ql.security.authorization.StorageBasedAuthorizationP
hive.security.authenticator.manager	org.apache.hadoop.hive.ql.security.ProxyUserAuthenticator

For more information about configuring Hive security, see [Metastore Server Security](#) in [Hive Authorization](#) and the HCatalog document [Storage Based Authorization](#).

20.2. Oozie

1. On the Oozie Server machine, install the appropriate JDBC .jar file.

- For **Oracle**:

a. Download the Oracle JDBC (OJDBC driver from <http://www.oracle.com/technetwork/database/features/jdbc/index-091264.html>.

Select Oracle Database 11g Release 2 - ojdbc6.jar .

b. Copy the .jar file to the Java share directory.

```
cp ojdbc6.jar /usr/share/java
```

c. Make sure the .jar file has the appropriate permissions - 644.

- For **MySQL**:

a. Install the connector.

- RHEL/CentOS/Oracle Linux

```
yum install mysql-connector-java-5.0.8-1
```

- SLES

```
zypper install mysql-connector-java-5.0.8-1
```

b. Confirm that the MySQL .jar file is in the Java share directory.

```
ls /usr/share/java/mysql-connector-java.jar
```

c. Make sure the .jar file has the appropriate permissions - 644.

2. On the Ambari Server host, install the appropriate JDBC .jar file.

- For **Oracle**:

a. Download the Oracle JDBC (OJDBC) driver from <http://www.oracle.com/technetwork/database/features/jdbc/index-091264.html>.

Select Oracle Database 11g Release 2 - ojdbc6.jar .

b. Copy the .jar file to the Java share directory.

```
cp ojdbc6.jar /var/lib/ambari-server/resources
```

c. Make sure the .jar file has the appropriate permissions - 644.

- For **MySQL**:
 - a. Download the mysql connector driver from the host on which you installed mysql-connector-java.
 - b. Copy the .jar file to the Java share directory.

```
cp mysql-connector-java.jar /var/lib/ambari-server/resources
```

- c. Make sure the .jar file has the appropriate permissions - 644.

3. Create a user for Oozie and grant it permissions.

- For **Oracle**, create the Oozie user and grant it database permissions:

```
# sqlplus sys/root as sysdba
SQL> CREATE USER $OOZIEUSER IDENTIFIED BY $OOZIEPASSWORD;
SQL> GRANT ALL PRIVILEGES TO $OOZIEUSER;
SQL> QUIT;
```

Where `$OOZIEUSER` is the Oozie user name and `$OOZIEPASSWORD` is the Oozie user password.

- For **MySQL**:

- a. Create the Oozie user and grant it database permissions.

```
# mysql -u root -p
mysql> CREATE USER '$OOZIEUSER'@'%' IDENTIFIED BY '$OOZIEPASSWORD';
mysql> GRANT ALL PRIVILEGES ON *.* TO '$OOZIEUSER'@'%' ;
mysql> flush privileges;
```

Where `$OOZIEUSER` is the Oozie user name and `$OOZIEPASSWORD` is the Oozie user password.

- b. Create the Oozie database.

```
# mysql -u root -p
mysql> CREATE DATABASE oozie;
```

20.2.1. Troubleshooting Oozie

Use these entries to help you troubleshoot any issues you might have installing Oozie with non-default databases.

20.2.1.1. Problem: Oozie Server Install Fails Using MySQL

Check the install log:

```
cp /usr/share/java/mysql-connector-java.jar
/usr/lib/oozie/libext/mysql-connector-java.jar]
has failures: true
```

The MySQL JDBC .jar file cannot be found.

20.2.1.1.1. Soution

Make sure the file is in the appropriate directory on the Oozie server and click **Retry**.

20.2.1.2. Problem: Oozie Server Install Fails Using Oracle or MySQL

Check the install log:

```
Exec[exec cd /var/tmp/oozie &&
/usr/lib/oozie/bin/ooziedb.sh create -sqlfile oozie.sql -run ]
has failures: true
```

Oozie was unable to connect to the database or was unable to successfully setup the schema for Oozie.

20.2.1.2.1. Soution

Check the database connection settings provided during the **Customize Services** step in the install wizard by browsing back to **Customize Services** -> **Oozie**. After confirming (and adjusting) your database settings, proceed forward with the install wizard.

If the Install Oozie Server continues to fail, get more information by connecting directly to the Oozie server and executing the following command as `$OOZIEUSER`:

```
su oozie
/usr/lib/oozie/bin/ooziedb.sh create -sqlfile oozie.sql -run
```

20.3. Ambari

1. On the Ambari Server machine, install the appropriate .jar file.

- For **Oracle**:

- a. Download the Oracle JDBC (OJDBC) driver from <http://www.oracle.com/technetwork/database/features/jdbc/index-091264.html>.
- b. Select Oracle Database 11g Release 2 - ojdbc6.jar
- c. Copy the .jar file to the Java share directory.

```
cp ojdbc6.jar /usr/share/java
```

- For **MySQL**:

Install the MySQL connector.

```
yum install mysql-connector-java*
```

2. Make sure the .jar file has the appropriate permissions.

- For **Oracle**:

The .jar file should have 644 permissions.

- For **MySQL**:

```
cd /usr/share/java ls mysql*
```

Should return present and 644 permissions.

3. Create the Ambari user, password, and tablespace, and grant the account database permissions.

- For **Oracle**:

```
# sqlplus sys/root as sysdba
SQL> create user $AMBARIUSER identified by $AMBARIPASSWORD default
tablespace "USERS" temporary tablespace "TEMP";
SQL> grant unlimited tablespace to $AMBARIUSER;
SQL> grant create session to $AMBARIUSER;
SQL> grant create table to $AMBARIUSER;
SQL> quit;
```

Where `$AMBARIUSER` is the Ambari user name and `$AMBARIPASSWORD` is the Ambari user password.

- For **MySQL**, on the MySQL server host:

```
mysql -u root
CREATE USER '$AMBARIUSER'@'%' IDENTIFIED BY '$AMBARIPASSWORD';
GRANT ALL PRIVILEGES ON *.* TO '$AMBARIUSER'@'%';
CREATE USER '$AMBARIUSER'@'localhost' IDENTIFIED BY '$AMBARIPASSWORD';
GRANT ALL PRIVILEGES ON *.* TO '$AMBARIUSER'@'localhost';
CREATE USER '$AMBARIUSER'@'$AMBARISERVERFQDN' IDENTIFIED BY
'$AMBARIPASSWORD';
GRANT ALL PRIVILEGES ON *.* TO '$AMBARIUSER'@'$AMBARISERVERFQDN';
FLUSH PRIVILEGES;
```

Where `$AMBARISERVERFQDN` is the Fully Qualified Domain Name of the Ambari server host.

4. Load the database schema on the Ambari Server host.

- For **Oracle**:

Create the Ambari Server schema by running a script.

```
sqlplus $AMBARIUSER/$AMBARIPASSWORD <
/var/lib/ambari-server/resources/Ambari-DDL-Oracle-CREATE.sql
```

The file `Ambari-DDL-Oracle-CREATE.sql` is found in the `/var/lib/ambari-server/resources/` directory of the Ambari Server machine, once you have completed the [Set Up the Bits](#) step in the install process.

- For **MySQL**:

Create the Ambari Server database and schema.

```
mysql -u $AMBARIUSER -p
CREATE DATABASE $AMBARIDATABASE;
USE $AMBARIDATABASE;
```

```
SOURCE Ambari-DDL-MySQL-CREATE.sql;
```

Where \$AMBARIDATABASE is the Ambari database name.

When setting up the Ambari server, select Advanced Database Configuration > option 3: MySQL.

On the Ambari server host, find the database schema in `/var/lib/ambari-server/resources/` after setting up the bits during the installation process.

20.3.1. Troubleshooting Ambari

Use these entries to help you troubleshoot any issues you might have installing Ambari with an existing Oracle database.

20.3.1.1. Problem: Ambari Server Fails to Start: No Driver

Check `/var/log/ambari-server/ambari-server.log` for :

```
ExceptionDescription:Configurationerror.  
Class[oracle.jdbc.driver.OracleDriver] not found.
```

The Oracle JDBC .jar file cannot be found.

20.3.1.1.1. Soution

Make sure the file is in the appropriate directory on the Ambari server and re-run `ambari-server setup`. See [Step 1](#) above.

20.3.1.2. Problem: Ambari Server Fails to Start: No Connection

Check `/var/log/ambari-server/ambari-server.log` for :

```
The Network Adapter could not establish the connection  
Error Code: 17002
```

Ambari Server cannot connect to the database.

20.3.1.2.1. Soution

Confirm the database host is reachable from the Ambari Server and is correctly configured by reading `/etc/ambari-server/conf/ambari.properties`

```
server.jdbc.url=jdbc:oracle:thin:  
    @oracle.database.hostname:1521/ambaridb  
server.jdbc.rca.url=jdbc:oracle:thin:  
    @oracle.database.hostname:1521/ambaridb
```

20.3.1.3. Problem: Ambari Server Fails to Start: Bad Username

Check `/var/log/ambari-server/ambari-server.log` for :

```
Internal Exception: java.sql.SQLException: ORA01017:  
invalid username/password; logon denied
```

You are using an invalid username/password.

20.3.1.3.1. Soution

Confirm the user account is set up in the database and has the correct privileges. See [Step 3](#) above.

20.3.1.4. Problem: Ambari Server Fails to Start: No Schema

Check `/var/log/ambari-server/ambari-server.log` for :

```
Internal Exception: java.sql.SQLException: ORA00942:  
table or view does not exist
```

The schema has not been loaded.

20.3.1.4.1. Soution

Confirm you have loaded the database schema. See [Step 4](#) above.

21. Setting Up Kerberos for Use with Ambari

This section provides information on setting up Kerberos for an Ambari-installed version of Hadoop.

- [Setting Up Kerberos for Hadoop 2.x](#)
- [Setting Up Kerberos for Hadoop 1.x](#)

21.1. Setting Up Kerberos for Hadoop 2.x

Use the following instructions to prepare to deploy a secure Hadoop cluster version 2.x:

1. [Preparing Kerberos](#)
2. [Setting Up Hadoop Users](#)
3. [Enabling Kerberos Security](#)

21.1.1. Preparing Kerberos

This section provides information on setting up Kerberos.

1. [Kerberos Overview](#)
2. [Installing and Configuring the KDC](#)
3. [Creating the Database](#)
4. [Starting the KDC](#)
5. [Installing and Configuring the Kerberos Clients](#)
6. [Creating Service Principals and Keytab Files for Hadoop](#)
7. [Setup Kerberos JAAS Configuration for Ambari](#)

21.1.1.1. Kerberos Overview

Establishing identity with strong authentication is the basis for secure access in Hadoop. Users need to be able to reliably “identify” themselves and then have that identity propagated throughout the Hadoop cluster. Once this is done those users can access resources (such as files or directories) or interact with the cluster (like running MapReduce jobs). As well, Hadoop cluster resources themselves (such as Hosts and Services) need to authenticate with each other to avoid potential malicious systems “posing as” part of the cluster to gain access to data.

To create that secure communication among its various components, Hadoop uses Kerberos. Kerberos is a third party authentication mechanism, in which users and services that users want to access rely on a third party - the Kerberos server - to authenticate each to the other. The Kerberos server itself is known as the *Key Distribution Center*, or KDC. At a high level, it has three parts:

- A database of the users and services (known as *principals*) that it knows about and their respective Kerberos passwords
- An *authentication server* (AS) which performs the initial authentication and issues a *Ticket Granting Ticket* (TGT)
- A *Ticket Granting Server* (TGS) that issues subsequent service tickets based on the initial TGT

A user principal requests authentication from the AS. The AS returns a TGT that is encrypted using the user principal's Kerberos password, which is known only to the user principal and the AS. The user principal decrypts the TGT locally using its Kerberos password, and from that point forward, until the ticket expires, the user principal can use the TGT to get service tickets from the TGS. Service tickets are what allow a principal to access various services.

Because cluster resources (hosts or services) cannot provide a password each time to decrypt the TGT, they use a special file, called a *keytab*, which contains the resource principal's authentication credentials.

The set of hosts, users, and services over which the Kerberos server has control is called a *realm*.

Table 21.1. Kerberos terminology

Term	Description
Key Distribution Center, or KDC	The trusted source for authentication in a Kerberos-enabled environment.
Kerberos KDC Server	The machine, or server, that serves as the Key Distribution Center.
Kerberos Client	Any machine in the cluster that authenticates against the KDC.
Principal	The unique name of a user or service that authenticates against the KDC.
Keytab	A file that includes one or more principals and their keys.
Realm	The Kerberos network that includes a KDC and a number of Clients.

21.1.1.2. Installing and Configuring the KDC

To use Kerberos with Hadoop you can either use an existing KDC or install a new one just for Hadoop's use. The following gives a very high level description of the installation process. To get more information see [RHEL documentation](#), [CentOS documentation](#), or [SLES documentation](#).



Note

Because Kerberos is a time-sensitive protocol, all hosts in the realm must be time-synchronized, for example, by using the Network Time Protocol (NTP). If the local system time of a client differs from that of the KDC by as little as 5 minutes (the default), the client will not be able to authenticate.

1. To install a new version of the server:

```
[On RHEL, CentOS, or Oracle Linux]
yum install krb5-server krb5-libs krb5-auth-dialog krb5-workstation
```

OR

```
[On SLES]
zypper install krb5 krb5-server krb5-client
```



Note

The host on which you install the KDC must itself be secure.

2. When the server is installed use a text editor to edit the configuration file, located by default here:

```
/etc/krb5.conf
```

Change the `[realms]` section of this file by replacing the default “kerberos.example.com” setting for the `kdc` and `admin_server` properties with the Fully Qualified Domain Name of the KDC server. In this example below, “kerberos.example.com” has been replaced with “my.kdc.server”.

```
[realms]
EXAMPLE.COM = {
    kdc = my.kdc.server
    admin_server = my.kdc.server
}
```

21.1.1.3. Creating the Database

Use the utility `kdb5_util` to create the Kerberos database.

```
[on RHEL, CentOS, or Oracle Linux]
/usr/sbin/kdb5_util create -s
```

OR

```
[on SLES]
kdb5_util create -s
```

21.1.1.4. Starting the KDC

Start the KDC.

```
[on RHEL, CentOS, or Oracle Linux]
/etc/rc.d/init.d/krb5kdc start
/etc/rc.d/init.d/kadmin start
```

OR

```
[on SLES]
rckrb5kdc start
rckadmind start
```

21.1.1.5. Installing and Configuring the Kerberos Clients

1. To install the Kerberos clients, on every server in the cluster:

```
[on RHEL, CentOS, or Oracle Linux]
yum install krb5-workstation
```

OR

```
[on SLES]
zypper install krb5-client
```

2. Copy the `krb5.conf` file you modified in [Installing and Configuring the KDC](#) to all the servers in the cluster.

21.1.1.6. Creating Service Principals and Keytab Files for Hadoop 2.x

Each service and sub-service in Hadoop must have its own principal. A principal name in a given realm consists of a primary name and an instance name, which in this case is the FQDN of the host that runs that service. As services do not login with a password to acquire their tickets, their principal's authentication credentials are stored in a keytab file, which is extracted from the Kerberos database and stored locally with the service principal on the service component host.

First you must create the principal, using mandatory naming conventions.

Then you must create the keytab file with that principal's information and copy the file to the keytab directory on the appropriate service host.



Note

Principals can be created either on the KDC machine itself or through the network, using an "admin" principal. The following instructions assume you are using the KDC machine and using the `kadmin.local` command line administration utility. Using `kadmin.local` on the KDC machine allows you to create principals without needing to create a separate "admin" principal before you start.

1. Open the `kadmin.local` utility on the KDC machine

```
/usr/sbin/kadmin.local
```

2. Create the service principals:

```
$kadmin.local
addprinc -randkey $primary_name/$fully.qualified.domain.name@EXAMPLE.COM
```

The `-randkey` option is used to generate the password.

Note that in the example each service principal's primary name has appended to it the instance name, the FQDN of the host on which it runs. This provides a unique principal name for services that run on multiple hosts, like `DataNodes` and `NodeManagers`. The addition of the hostname serves to distinguish, for example, a request from `DataNode A` from a request from `DataNode B`. This is important for two reasons:

- If the Kerberos credentials for one `DataNode` are compromised, it does not automatically lead to all `DataNodes` being compromised

- If multiple DataNodes have exactly the same principal and are simultaneously connecting to the NameNode, and if the Kerberos authenticator being sent happens to have same timestamp, then the authentication would be rejected as a replay request.

The principal name must match the values in the table below:

Table 21.2. Service Principals

Service	Component	Mandatory Principal Name
HDFS	NameNode	nn/\$FQDN
HDFS	NameNode HTTP	HTTP/\$FQDN
HDFS	SecondaryNameNode	nn/\$FQDN
HDFS	SecondaryNameNode HTTP	HTTP/\$FQDN
HDFS	DataNode	dn/\$FQDN
MR2	History Server	jhs/\$FQDN
MR2	History Server HTTP	HTTP/\$FQDN
YARN	ResourceManager	rm/\$FQDN
YARN	NodeManager	nm/\$FQDN
Oozie	Oozie Server	oozie/\$FQDN
Oozie	Oozie HTTP	HTTP/\$FQDN
Hive	Hive Metastore	hive/\$FQDN
	HiveServer2	
Hive	WebHCat	HTTP/\$FQDN
HBase	MasterServer	hbase/\$FQDN
HBase	RegionServer	hbase/\$FQDN
ZooKeeper	ZooKeeper	zookeeper/\$FQDN
Nagios Server	Nagios	nagios/\$FQDN
JournalNode Server ^a	JournalNode	jn/\$FQDN

^aOnly required if you are setting up NameNode HA.

For example: To create the principal for a DataNode service, issue this command:

```
$kadmin.local
addprinc -randkey dn/$DataNode-Host@EXAMPLE.COM
```

3. In addition you must create four special principals for Ambari's own use.



Note

The names in table below can be customized in the Customize Services step of the Ambari Install Wizard. If this is the case in your installation, the principal names should match the customized names. For example, if the HDFS Service User has been set to `hdfs1`, the respective principal for the Ambari HDFS User should also be created as `hdfs1`.

These principals do not have the FQDN appended to the primary name:

Table 21.3. Ambari Principals

User	Mandatory Principal Name
Ambari User ^a	ambari
Ambari Smoke Test User	ambari-qa
Ambari HDFS User	hdfs
Ambari HBase User	hbase

^aThis principal is used with the JAAS configuration. See [Setup Kerberos JAAS Configuration for Ambari](#) for more information.

- Once the principals are created in the database, you can extract the related keytab files for transfer to the appropriate host:

```
$kadmin.local
xst -norandkey -k $keytab_file_name $primary_name/fully.qualified.domain.name@EXAMPLE.COM
```

You must use the mandatory names for the `$keytab_file_name` variable shown in this table.



Note

Some older versions of Kerberos do not support the `xst -norandkey` option. You can use the command without the `-norandkey` flag, except in cases where you need to copy a principal from one keytab file to another keytab file on a host. This might be a requirement if the Hadoop configurations on a host have keytab path properties that point to different keytab locations but have corresponding principal name properties that have the same values.

In situations like this, you can use the two step `kadmin/kutil` procedure. This description assumes MIT Kerberos. If you are using another version, please check your documentation.

- Extract the keytab file information:

```
$kadmin
xst -k $keytab_file_name-temp1 $primary_name/fully.qualified.domain.name@EXAMPLE.COM
xst -k $keytab_file_name-temp2 $primary_name/fully.qualified.domain.name@EXAMPLE.COM
```

- Write the keytab to a file. :

```
$kutil
kutil: rkt $keytab_file_name-temp1
kutil: rkt $keytab_file_name-temp2
kutil: wkt $keytab_file_name
kutil: clear
```

Table 21.4. Service Keytab File Names

Component	Principal Name	Mandatory Keytab File Name
NameNode	nn/\$FQDN	nn.service.keytab

Component	Principal Name	Mandatory Keytab File Name
NameNode HTTP	HTTP/\$FQDN	spnego.service.keytab
SecondaryNameNode	nn/\$FQDN	nn.service.keytab
SecondaryNameNode HTTP	HTTP/\$FQDN	spnego.service.keytab
DataNode	dn/\$FQDN	dn.service.keytab
MR2 History Server	jhs/\$FQDN	jhs.service.keytab
MR2 History Server HTTP	HTTP/\$FQDN	spnego.service.keytab
YARN	rm/\$FQDN	rm.service.keytab
YARN	nm/\$FQDN	nm.service.keytab
Oozie Server	oozie/\$FQDN	oozie.service.keytab
Oozie HTTP	HTTP/\$FQDN	spnego.service.keytab
Hive Metastore	hive/\$FQDN	hive.service.keytab
HiveServer2		
WebHCat	HTTP/\$FQDN	spnego.service.keytab
HBase Master Server	hbase/\$FQDN	hbase.service.keytab
HBase RegionServer	hbase/\$FQDN	hbase.service.keytab
ZooKeeper	zookeeper/\$FQDN	zk.service.keytab
Nagios Server	nagios/\$FQDN	nagios.service.keytab
Journal Server ^a	jn/\$FQDN	jn.service.keytab
Ambari User ^b	ambari	ambari.keytab
Ambari Smoke Test User	ambari-qa	smokeuser.headless.keytab
Ambari HDFS User	hdfs	hdfs.headless.keytab
Ambari HBase User	hbase	hbase.headless.keytab

^aOnly required if you are setting up NameNode HA.

^bThis principal is used with the JAAS configuration. See [Setup Kerberos JAAS Configuration for Ambari](#) for more information.

For example: To create the keytab files for NameNode HTTP, issue this command:

```
$kadmin.local
xst -norandkey -k spnego.service.keytab HTTP/<namenode-host>
```



Note

If you have a large cluster, you may want to create a script to automate creating your principals and keytabs. To help with that, you can download a CSV-formatted file of all the required principal names and keytab files from the Ambari Web GUI. Select **Admin** view->**Security**->**Enable Security**-> and run the **Enable Security Wizard**, using the default values. At the bottom of the third page, **Create Principals and Keytabs**, click **Download CSV**.

- When the keytab files have been created, on each host create a directory for them and set appropriate permissions.

```
mkdir -p /etc/security/keytabs/
chown root:hadoop /etc/security/keytabs
chmod 750 /etc/security/keytabs
```

- Copy the appropriate keytab file to each host. If a host runs more than one component (for example, both NodeManager and DataNode), copy keytabs for both components.

The Ambari Smoke Test User, the Ambari HDFS User, and the Ambari HBase User keytabs should be copied to the all hosts on the cluster.



Important

If you have customized service user names, replace the default values below with your appropriate service user, group, and keytab names.

7. Set appropriate permissions for the keytabs.

- a. Optionally, if you have [Setup Kerberos JAAS Configuration for Ambari](#) on the Ambari server host:

```
chown ambari:ambari /etc/security/keytabs/ambari.keytab
chmod 400 /etc/security/keytabs/ambari.keytab
```

- b. On the HDFS NameNode and SecondaryNameNode hosts:

```
chown hdfs:hadoop /etc/security/keytabs/nn.service.keytab
chmod 400 /etc/security/keytabs/nn.service.keytab
chown root:hadoop /etc/security/keytabs/spnego.service.keytab
chmod 440 /etc/security/keytabs/spnego.service.keytab
```

- c. On the HDFS NameNode host, for the Ambari Test Users:

```
chown ambari-qa:hadoop /etc/security/keytabs/smokeuser.headless.keytab
chmod 440 /etc/security/keytabs/smokeuser.headless.keytab
chown hdfs:hadoop /etc/security/keytabs/hdfs.headless.keytab
chmod 440 /etc/security/keytabs/hdfs.headless.keytab
chown hbase:hadoop /etc/security/keytabs/hbase.headless.keytab
chmod 440 /etc/security/keytabs/hbase.headless.keytab
```

- d. On each host that runs an HDFS DataNode:

```
chown hdfs:hadoop /etc/security/keytabs/dn.service.keytab
chmod 400 /etc/security/keytabs/dn.service.keytab
```

- e. On the host that runs the MR2 History Server:

```
chown mapred:hadoop /etc/security/keytabs/jhs.service.keytab
chmod 400 /etc/security/keytabs/jhs.service.keytab
chown root:hadoop /etc/security/keytabs/spnego.service.keytab
chmod 440 /etc/security/keytabs/spnego.service.keytab
```

- f. On the host that runs the YARN ResourceManager:

```
chown yarn:hadoop /etc/security/keytabs/rm.service.keytab
chmod 400 /etc/security/keytabs/rm.service.keytab
```

- g. On each host that runs a YARN NodeManager:

```
chown yarn:hadoop /etc/security/keytabs/nm.service.keytab
chmod 400 /etc/security/keytabs/nm.service.keytab
```

- h. On the host that runs the Oozie Server:

```
chown oozie:hadoop /etc/security/keytabs/oozie.service.keytab
chmod 400 /etc/security/keytabs/oozie.service.keytab
```

```
chown root:hadoop /etc/security/keytabs/spnego.service.keytab
chmod 440 /etc/security/keytabs/spnego.service.keytab
```

- i. On the host that runs the Hive Metastore, HiveServer2 and WebHCat:

```
chown hive:hadoop /etc/security/keytabs/hive.service.keytab
chmod 400 /etc/security/keytabs/hive.service.keytab
chown root:hadoop /etc/security/keytabs/spnego.service.keytab
chmod 440 /etc/security/keytabs/spnego.service.keytab
```

- j. On hosts that run the HBase MasterServer, RegionServer and ZooKeeper:

```
chown hbase:hadoop /etc/security/keytabs/hbase.service.keytab
chmod 400 /etc/security/keytabs/hbase.service.keytab
chown zookeeper:hadoop /etc/security/keytabs/zk.service.keytab
chmod 400 /etc/security/keytabs/zk.service.keytab
```

- k. On the host that runs the Nagios server:

```
chown nagios:nagios /etc/security/keytabs/nagios.service.keytab
chmod 400 /etc/security/keytabs/nagios.service.keytab
```

- l. On each host that runs a JournalNode, if you are setting up NameNode HA:

```
chown hdfs:hadoop /etc/security/keytabs/jn.service.keytab
chmod 400 /etc/security/keytabs/jn.service.keytab
```

8. Verify that the correct keytab files and principals are associated with the correct service using the `klist` command. For example, on the NameNode:

```
klist -k -t /etc/security/keytabs/nn.service.keytab
```

Do this on each respective service in your cluster.

21.1.1.7. Setup Kerberos JAAS Configuration for Ambari

If you want to set up JAAS configuration for Ambari so that independent access to native Hadoop GUIs like the NameNode UI are secure, use the [Apache community docs](#) to setup your configurations, and then do the following:

1. Log into the Ambari server host.



Important

Ambari Server should not be running when you do this: either make the edits before you start Ambari Server the first time or bring the server down to make the edits.

2. Run the special setup command and answer the prompts:

```
ambari-server setup-security
```

- a. Select 5 for **Setup Ambari kerberos JAAS configuration**.
- b. Enter the Kerberos principal name for the Ambari server you set up [here](#).
- c. Enter the path to the keytab for Ambari's principal.

d. Restart Ambari Server:

```
ambari-server restart
```

21.1.2. Setting Up Hadoop Users

This section provides information on setting up Hadoop users for Kerberos.

- [Overview](#)
- [Creating Mappings Between Principals and UNIX Usernames](#)

21.1.2.1. Overview

Hadoop uses users' group memberships at various places for things like determining group ownership for files or for access control. To configure Hadoop for use with Kerberos and Ambari you must create a mapping between service principals and these UNIX usernames.

A user is mapped to the groups it belongs to using an implementation of the `GroupMappingServiceProvider` interface. The implementation is pluggable and is configured in `core-site.xml`.

By default Hadoop uses `ShellBasedUnixGroupsMapping`, which is an implementation of `GroupMappingServiceProvider`. It fetches the group membership for a username by executing a UNIX shell command. In secure clusters, since the usernames are actually Kerberos principals, `ShellBasedUnixGroupsMapping` will work only if the Kerberos principals map to valid UNIX usernames. Hadoop provides a feature that lets administrators specify mapping rules to map a Kerberos principal to a local UNIX username .

21.1.2.2. Creating Mappings Between Principals and UNIX Usernames

Hadoop uses a rule-based system to create mappings between service principals and their related UNIX usernames. The rules are specified in the `core-site.xml` configuration file as the value to the optional key `hadoop.security.auth_to_local`.

The default rule is simply named `DEFAULT`. It translates all principals in your default domain to their first component. For example, `myusername@APACHE.ORG` and `myusername/admin@APACHE.ORG` both become `myusername`, assuming your default domain is `APACHE.ORG`.

Use the following instructions to configure the mappings between principals and UNIX usernames:

1. [Creating Rules](#)
2. [Examples](#)

21.1.2.2.1. Creating Rules

- Simple Rules

To make a simple map between principal names and UNIX users, you create a straightforward substitution rule. For example, to map the `ResourceManager(rm)` and `NodeManager(nm)` principals in the `EXAMPLE.COM` realm to the UNIX `$YARN_USER` user and the `NameNode(nn)` and `DataNode(dn)` principals to the UNIX `$HDFS_USER`

user, you would make this the value for the `hadoop.security.auth_to_local` key in `core-site.xml`.

```
RULE:[2:$1@$0]([jt]t@.*EXAMPLE.COM)s/.*/$YARN_USER /
RULE:[2:$1@$0]([nd]n@.*EXAMPLE.COM)s/.*/$HDFS_USER/
DEFAULT
```

- **Complex Rules**

To accommodate more complex translations, you create a hierarchical set of rules to add to the default. Each rule is divided into three parts: base, filter, and substitution.

- **The Base:**

The base begins with the number of components in the principal name (excluding the realm), followed by a colon, and the pattern for building the username from the sections of the principal name. In the pattern section `$0` translates to the realm, `$1` translates to the first component and `$2` to the second component.

For example:

`[1:$1@$0]` translates `myusername@APACHE.ORG` to `myusername@APACHE.ORG`

`[2:$1]` translates `myusername/admin@APACHE.ORG` to `myusername`

`[2:$1%$2]` translates `myusername/admin@APACHE.ORG` to `myusername%admin`

- **The Filter:**

The filter consists of a regex in a parentheses that must match the generated string for the rule to apply.

For example:

`(.*%admin)` matches any string that ends in `%admin`

`(.*@SOME.DOMAIN)` matches any string that ends in `@SOME.DOMAIN`

- **The Substitution:**

The substitution is a sed rule that translates a regex into a fixed string.

For example:

`s/@ACME\.COM//` removes the first instance of `@SOME.DOMAIN`.

`s/@[A-Z]*\.COM//` removes the first instance of `@` followed by a name followed by `COM`.

`s/X/Y/g` replaces all of the `X` in the name with `Y`

21.1.2.2.2. Examples

- If your default realm was `APACHE.ORG`, but you also wanted to take all principals from `ACME.COM` that had a single component `joe@ACME.COM`, you would create this rule:

```
RULE:[1:$1@$0](.*@ACME\.COM)s/@.*//
DEFAULT
```

- To also translate names with a second component, you would use these rules:

```
RULE:[1:$1@$0](.*@ACME\.COM)s/@.*//
RULE:[2:$1@$0](.*@ACME\.COM)s/@.*//
DEFAULT
```

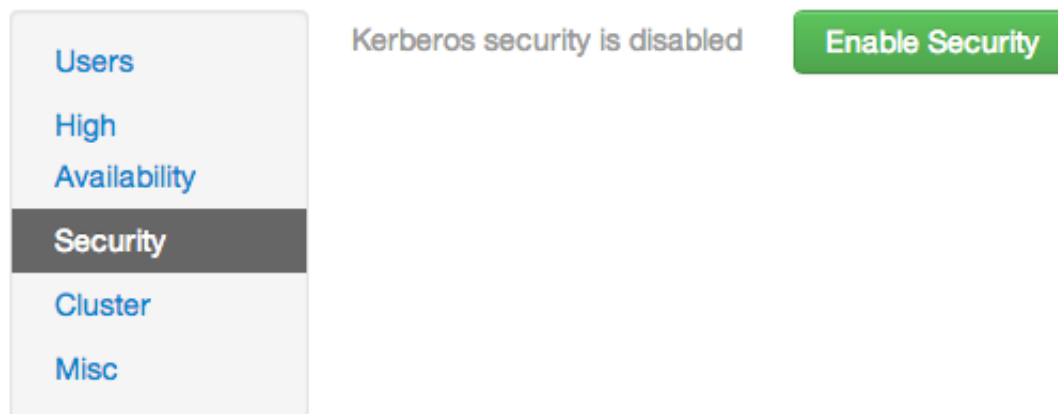
- To treat all principals from APACHE.ORG with the extension /admin as admin, your rules would look like this:

```
RULE[2:$1%$2@$0](.*%admin@APACHE\.ORG)s/.*admin/
DEFAULT
```

21.1.3. Enabling Kerberos Security

To turn on Kerberos-based security in the Ambari Web GUI you must:

1. Have already set up Kerberos for your cluster. For more information, see [Setting Up Kerberos for Hadoop 2.x](#).
2. Go to the **Admin** tab.
3. Select **Security**.
4. Click **Enable Security** and follow the **Enable Security Wizard**.



- a. **Get Started:** This step just reminds you that you need to set up Kerberos before you start.
- b. **Configure Services:** This step asks you for your Kerberos information: principals and paths to keytabs, path to your Kerberos tools, realm names and so on. For more information about a specific field, hover over it, and a popup with a definition appears.

- c. **Create Principals and Keytabs:** Use this step to check that all your information is correct. Click **Back** to make any changes. Click **Apply** when you are satisfied with the assignments.



Note

If you have a large cluster, you may want to go to the **Create Principals and Keytabs** step first. Step through the wizard accepting the defaults to get to the appropriate page. On the page, use the **Download CSV** button to get a list of all the necessary principals and keytabs in CSV form, which can be used to set up a script. The list includes hostname, principal description, principal name, keytab user, keytab group, keytab permissions, absolute keytab path, and keytab filename.

SECURITY WIZARD

Get Started

Configure Services

Create Principals and Keytabs

Save and Apply Configuration

Create Principals and Keytabs

You need to create the following principals and keytabs on the hosts shown. You can download the list as a CSV file and use it to create a script to generate the principals and keytabs. Once the principals and keytabs have been created, click on Proceed to continue. If you need to make configuration changes, click Back.

Host	Component	Principal	Keytab
FQDN	Ambari Smoke Test User	ambari-qa@EXAMPLE.COM	/etc/security/keytabs/smokeuser.headless.keytab
FQDN	HDFS User	hdfs@EXAMPLE.COM	/etc/security/keytabs/hdfs.headless.keytab
FQDN	HBase User	hbase@EXAMPLE.COM	/etc/security/keytabs/hbase.headless.keytab
FQDN	SPNEGO User	HTTP/p-10-191-37-233.ec2.internal@EXAMPLE.COM	/etc/security/keytabs/spnego.service.keytab
FQDN	DataNode	dn/p-10-191-37-233.ec2.internal@EXAMPLE.COM	/etc/security/keytabs/dn.service.keytab
FQDN	HBase Master	hbase/p-10-191-37-233.ec2.internal@EXAMPLE.COM	/etc/security/keytabs/hbase.service.keytab
FQDN	HiveServer2	hive/p-10-191-37-233.ec2.internal@EXAMPLE.COM	/etc/security/keytabs/hive.service.keytab
FQDN		jt/p-10-191-37-233.ec2.internal@EXAMPLE.COM	/etc/security/keytabs/jt.service.keytab

← Back
Download CSV
Apply →

- d. **Save and Apply Configuration:** This step displays a bar showing the progress of integrating the Kerberos information into your Ambari Server.

21.2. Setting Up Kerberos for Hadoop 1.x

- [Preparing Kerberos](#)
- [Setting Up Hadoop Users](#)
- [Enabling Kerberos Security](#)

21.2.1. Preparing Kerberos

This section provides information on setting up Kerberos.

1. [Kerberos Overview](#)
2. [Installing and Configuring the KDC](#)
3. [Creating the Database](#)
4. [Starting the KDC](#)
5. [Installing and Configuring the Kerberos Clients](#)
6. [Creating Service Principals and Keytab Files for Hadoop](#)
7. [Setup Kerberos JAAS Configuration for Ambari](#)

21.2.1.1. Kerberos Overview

Establishing identity with strong authentication is the basis for secure access in Hadoop. Users need to be able to reliably “identify” themselves and then have that identity propagated throughout the Hadoop cluster. Once this is done those users can access resources (such as files or directories) or interact with the cluster (like running MapReduce jobs). As well, Hadoop cluster resources themselves (such as Hosts and Services) need to authenticate with each other to avoid potential malicious systems “posing as” part of the cluster to gain access to data.

To create that secure communication among its various components, Hadoop uses Kerberos. Kerberos is a third party authentication mechanism, in which users and services that users want to access rely on a third party - the Kerberos server - to authenticate each to the other. The Kerberos server itself is known as the *Key Distribution Center*, or KDC. At a high level, it has three parts:

- A database of the users and services (known as *principals*) that it knows about and their respective Kerberos passwords
- An *authentication server* (AS) which performs the initial authentication and issues a *Ticket Granting Ticket* (TGT)
- A *Ticket Granting Server* (TGS) that issues subsequent service tickets based on the initial TGT

A user principal requests authentication from the AS. The AS returns a TGT that is encrypted using the user principal's Kerberos password, which is known only to the user principal and the AS. The user principal decrypts the TGT locally using its Kerberos password, and from that point forward, until the ticket expires, the user principal can use the TGT to get service tickets from the TGS. Service tickets are what allow a principal to access various services.

Because cluster resources (hosts or services) cannot provide a password each time to decrypt the TGT, it uses a special file, called a *keytab*, which contains the resource principal's authentication credentials.

The set of hosts, users, and services over which the Kerberos server has control is called a *realm*.

Table 21.5. Kerberos terminology

Term	Description
Key Distribution Center, or KDC	The trusted source for authentication in a Kerberos-enabled environment.
Kerberos KDC Server	The machine, or server, that serves as the Key Distribution Center.
Kerberos Client	Any machine in the cluster that authenticates against the KDC.
Principal	The unique name of a user or service that authenticates against the KDC.
Keytab	A file that includes one or more principals and their keys.
Realm	The Kerberos network that includes a KDC and a number of Clients.

21.2.1.2. Installing and Configuring the KDC

To use Kerberos with Hadoop you can either use an existing KDC or install a new one just for Hadoop's use. The following gives a very high level description of the installation process. To get more information see [RHEL documentation](#) or [CentOS documentation](#) or [SLES documentation](#).



Note

Because Kerberos is a time-sensitive protocol, all hosts in the realm must be time-synchronized, for example, by using the Network Time Protocol (NTP). If the local system time of a client differs from that of the KDC by as little as 5 minutes (the default), the client will not be able to authenticate.

1. To install a new version of the server:

```
[On RHEL, CentOS, or Oracle Linux]
yum install krb5-server krb5-libs krb5-auth-dialog krb5-workstation
```

OR

```
[On SLES]
zypper install krb5 krb5-server krb5-client
```



Note

The host on which you install the KDC must itself be secure.

2. When the server is installed use a text editor to edit the configuration file, located by default here:

```
/etc/krb5.conf
```

Change the `[realms]` section of this file by replacing the default "kerberos.example.com" setting for the `kdc` and `admin_server` properties with the Fully Qualified Domain Name of the KDC server. In this example below, "kerberos.example.com" has been replaced with "my.kdc.server".

```
[realms]
```

```
EXAMPLE.COM = {  
  kdc = my.kdc.server  
  admin_server = my.kdc.server  
}
```

21.2.1.3. Creating the Database

Use the utility `kdb5_util` to create the Kerberos database.

```
[on RHEL,CentOS, or Oracle Linux]  
/usr/sbin/kdb5_util create -s
```

OR

```
[on SLES]  
kdb5_util create -s
```

The `-s` option stores the master server key for the database in a *stash* file. If the stash file is not present, you must log into the KDC with the master password (specified during installation) each time it starts.

21.2.1.4. Starting the KDC

Start the KDC.

```
[on RHEL, CentOS, or Oracle Linux]  
/etc/rc.d/init.d/krb5kdc start  
/etc/rc.d/init.d/kadmin start
```

OR

```
[on SLES]  
rckrb5kdc start  
rckadmind start
```

21.2.1.5. Installing and Configuring the Kerberos Clients

1. To install the Kerberos clients, on every server in the cluster:

```
[on RHEL, CentOS, or Oracle Linux]  
yum install krb5-workstation
```

OR

```
[on SLES]  
zypper install krb5-client
```

2. Copy the `krb5.conf` file you modified in [Installing and Configuring the KDC](#) to all the servers in the cluster.

21.2.1.6. Creating Service Principals and Keytab Files for Hadoop

Each service and sub-service in Hadoop must have its own principal. A principal name in a given realm consists of a primary name and an instance name, which in this case is the FQDN of the host that runs that service. As services do not login with a password to acquire their tickets, their principal's authentication credentials are stored in a keytab file, which is

extracted from the Kerberos database and stored locally with the service principal on the service component host.

First you must create the principal, using mandatory naming conventions.

Then you must create the keytab file with that principal's information and copy the file to the keytab directory on the appropriate service host.



Note

Principals can be created either on the KDC machine itself or through the network, using an "admin" principal. The following instructions assume you are using the KDC machine and using the `kadmin.local` command line administration utility. Using `kadmin.local` on the KDC machine allows you to create principals without needing to create a separate "admin" principal before you start.

1. Open the `kadmin.local` utility on the KDC machine

```
/usr/sbin/kadmin.local
```

2. Create the service principals:

```
$kadmin.local
addprinc -randkey $primary_name/$fully.qualified.domain.name@EXAMPLE.COM
```

The `-randkey` option is used to generate the password.

Note that in the example each service principal's primary name has appended to it the instance name, the FQDN of the host on which it runs. This provides a unique principal name for services that run on multiple hosts, like DataNodes and TaskTrackers. The addition of the hostname serves to distinguish, for example, a request from DataNode A from a request from DataNode B. This is important for two reasons:

- If the Kerberos credentials for one DataNode are compromised, it does not automatically lead to all DataNodes being compromised
- If multiple DataNodes have exactly the same principal and are simultaneously connecting to the NameNode, and if the Kerberos authenticator being sent happens to have same timestamp, then the authentication would be rejected as a replay request.

The principal name must match the values in the table below:

Table 21.6. Service Principals

Service	Component	Mandatory Principal Name
HDFS	NameNode	<code>nn/\$FQDN</code>
HDFS	NameNode HTTP	<code>HTTP/\$FQDN</code>
HDFS	SecondaryNameNode	<code>nn/\$FQDN</code>
HDFS	SecondaryNameNode HTTP	<code>HTTP/\$FQDN</code>
HDFS	DataNode	<code>dn/\$FQDN</code>
MapReduce	JobTracker	<code>jt/\$FQDN</code>
MapReduce	TaskTracker	<code>tt/\$FQDN</code>

Service	Component	Mandatory Principal Name
Oozie	Oozie Server	oozie/\$FQDN
Oozie	Oozie HTTP	HTTP/\$FQDN
Hive	Hive Metastore	hive/\$FQDN
	HiveServer2	
Hive	WebHCat	HTTP/\$FQDN
HBase	MasterServer	hbase/\$FQDN
HBase	RegionServer	hbase/\$FQDN
ZooKeeper	ZooKeeper	zookeeper/\$FQDN
Nagios Server	Nagios	nagios/\$FQDN

For example: To create the principal for a DataNode service, issue this command:

```
$kadmin.local
addprinc -randkey dn/$DataNode-Host@EXAMPLE.COM
```

- In addition you must create four special principals for Ambari's own use.



Note

The names in table below can be customized in the Customize Services step of the Ambari Install Wizard. If this is the case in your installation, the principal names should match the customized names. For example, if the HDFS Service User has been set to `hdfs1`, the respective principal for the Ambari HDFS User should also be created as `hdfs1`.

These principals do not need the FQDN appended to the primary name:

Table 21.7. Ambari Principals

User	Default Principal Name
Ambari User ^a	ambari
Ambari Smoke Test User	ambari-ga
Ambari HDFS User	hdfs
Ambari HBase User	hbase

^aThis principal is used with the JAAS configuration. See [Setup Kerberos JAAS Configuration for Ambari](#) for more information.

- Once the principals are created in the database, you can extract the related keytab files for transfer to the appropriate host:

```
$kadmin.local
xst -norandkey -k $keytab_file_name $primary_name/fully.qualified.domain.name@EXAMPLE.COM
```



Note

Some older versions of Kerberos do not support the `xst -norandkey` option. You can use the command without the `-norandkey` flag, except in cases where you need to merge two principals with the same name into a single keytab file for a single host. In this case, you can use the two step

kadmin/kutil procedure. This description assumes MIT Kerberos. If you are using another version, please check your documentation.

- a. Extract the keytab file information:

```
$kadmin
xst -k $keytab_file_name-temp1 $primary_name/fully.qualified.
domain.name@EXAMPLE.COM
xst -k $keytab_file_name-temp2 $primary_name/fully.qualified.
domain.name@EXAMPLE.COM
```

- b. Merge the keytabs into a single file. :

```
$kutil
rkt $keytab_file_name-temp1
rkt $keytab_file_name-temp2
wkt $keytab_file_name
clear
```

You must use the mandatory names for the `keytab_file_name` variable shown in this table. Adjust the principal names if necessary.

Table 21.8. Service Keytab File Names

Component	Principal Name	Mandatory Keytab File Name
NameNode	nn/\$FQDN	nn.service.keytab
NameNode HTTP	HTTP/\$FQDN	spnego.service.keytab
SecondaryNameNode	nn/\$FQDN	nn.service.keytab
SecondaryNameNode HTTP	HTTP/\$FQDN	spnego.service.keytab
DataNode	dn/\$FQDN	dn.service.keytab
JobTracker	jt/\$FQDN	jt.service.keytab
TaskTracker	tt/\$FQDN	tt.service.keytab
Oozie Server	oozie/\$FQDN	oozie.service.keytab
Oozie HTTP	HTTP/\$FQDN	spnego.service.keytab
Hive Metastore	hive/\$FQDN	hive.service.keytab
HiveServer2		
WebHCat	HTTP/\$FQDN	spnego.service.keytab
HBase Master Server	hbase/\$FQDN	hbase.service.keytab
HBase RegionServer	hbase/\$FQDN	hbase.service.keytab
ZooKeeper	zookeeper/\$FQDN	zk.service.keytab
Nagios Server	nagios/\$FQDN	nagios.service.keytab
Ambari User ^a	ambari	ambari.keytab
Ambari Smoke Test User	ambari-qa	smokeuser.headless.keytab
Ambari HDFS User	hdfs	hdfs.headless.keytab
Ambari HBase User	hbase	hbase.headless.keytab

^aThis principal is used with the JAAS configuration. See [Setup Kerberos JAAS Configuration for Ambari](#) for more information.

For example: To create the keytab files for NameNode HTTP, issue this command:

```
xst -norandkey -k spnego.service.keytab HTTP/<namenode-host>
```



Note

If you have a large cluster, you may want to create a script to automate creating your principals and keytabs. The Ambari Web GUI can help. See [Create Principals and Keytabs \[168\]](#) for more information.

5. When the keytab files have been created, on each host create a directory for them and set appropriate permissions.

```
mkdir -p /etc/security/keytabs/  
chown root:hadoop /etc/security/keytabs  
chmod 750 /etc/security/keytabs
```

6. Copy the appropriate keytab file to each host. If a host runs more than one component (for example, both TaskTracker and DataNode), copy keytabs for both components. The Ambari Smoke Test User, the Ambari HDFS User, and the Ambari HBase User keytabs should be copied to all hosts.
7. Set appropriate permissions for the keytabs.



Important

If you have customized service user names, replace the default values below with your appropriate service user, group, and keytab names.

- a. Optionally, if you have [Setup Kerberos JAAS Configuration for Ambari](#) on the Ambari server host:

```
chown ambari:ambari /etc/security/keytabs/ambari.keytab  
chmod 400 /etc/security/keytabs/ambari.keytab
```

- b. On the HDFS NameNode and SecondaryNameNode hosts:

```
chown hdfs:hadoop /etc/security/keytabs/nn.service.keytab  
chmod 400 /etc/security/keytabs/nn.service.keytab  
chown root:hadoop /etc/security/keytabs/spnego.service.keytab  
chmod 440 /etc/security/keytabs/spnego.service.keytab
```

- c. On the HDFS NameNode host, for the Ambari Test Users:

```
chown ambari-qa:hadoop /etc/security/keytabs/smokeuser.headless.keytab  
chmod 440 /etc/security/keytabs/smokeuser.headless.keytab  
chown hdfs:hadoop /etc/security/keytabs/hdfs.headless.keytab  
chmod 440 /etc/security/keytabs/hdfs.headless.keytab  
chown hbase:hadoop /etc/security/keytabs/hbase.headless.keytab  
chmod 440 /etc/security/keytabs/hbase.headless.keytab
```

- d. On each host that runs an HDFS DataNode:

```
chown hdfs:hadoop /etc/security/keytabs/dn.service.keytab  
chmod 400 /etc/security/keytabs/dn.service.keytab
```

- e. On the host that runs the MapReduce JobTracker:

```
chown mapred:hadoop /etc/security/keytabs/jt.service.keytab  
chmod 400 /etc/security/keytabs/jt.service.keytab
```

- f. On each host that runs a MapReduce TaskTracker:

```
chown mapred:hadoop /etc/security/keytabs/tt.service.keytab
chmod 400 /etc/security/keytabs/tt.service.keytab
```

- g. On the host that runs the Oozie Server:

```
chown oozie:hadoop /etc/security/keytabs/oozie.service.keytab
chmod 400 /etc/security/keytabs/oozie.service.keytab
chown root:hadoop /etc/security/keytabs/spnego.service.keytab
chmod 440 /etc/security/keytabs/spnego.service.keytab
```

- h. On the host that runs the Hive Metastore, HiveServer2 and WebHCat:

```
chown hive:hadoop /etc/security/keytabs/hive.service.keytab
chmod 400 /etc/security/keytabs/hive.service.keytab
chown root:hadoop /etc/security/keytabs/spnego.service.keytab
chmod 440 /etc/security/keytabs/spnego.service.keytab
```

- i. On hosts that run the HBase MasterServer, RegionServer and ZooKeeper:

```
chown hbase:hadoop /etc/security/keytabs/hbase.service.keytab
chmod 400 /etc/security/keytabs/hbase.service.keytab
chown zookeeper:hadoop /etc/security/keytabs/zk.service.keytab
chmod 400 /etc/security/keytabs/zk.service.keytab
```

- j. On the host that runs the Nagios server:

```
chown nagios:nagios /etc/security/keytabs/nagios.service.keytab
chmod 400 /etc/security/keytabs/nagios.service.keytab
```

8. Verify that the correct keytab files and principals are associated with the correct service using the `klist` command. For example, on the NameNode:

```
klist -k -t /etc/security/keytabs/nn.service.keytab
```

Do this on each respective service in your cluster.

21.2.1.7. Setup Kerberos JAAS Configuration for Ambari

If you want to set up JAAS configuration for Ambari so that independent access to native Hadoop GUIs like the NameNode UI are secure, use the [Apache community docs](#) to setup your configurations, and then do the following:



Important

Ambari Server should not be running when you do this: either make the edits before you start Ambari Server the first time or bring the server down to make the edits.

Run the special setup command and answer the prompts

```
ambari-server setup-security
```

1. Select 5 for **Setup Ambari kerberos JAAS configuration**.
2. Enter the Kerberos principal name for the Ambari server you set up [here](#).

3. Enter the path to the keytab for Ambari's principal.
4. Restart Ambari Server:

```
ambari-server restart
```

For more information

21.2.2. Setting Up Hadoop Users

This section provides information on setting up Hadoop users for Kerberos.

- [Overview](#)
- [Creating Mappings Between Principals and UNIX Usernames](#)

21.2.2.1. Overview

Hadoop uses users' group memberships at various places for things like determining group ownership for files or for access control. To configure Hadoop for use with Kerberos and Ambari you must create a mapping between service principals and these UNIX usernames..

A user is mapped to the groups it belongs to using an implementation of the `GroupMappingServiceProvider` interface. The implementation is pluggable and is configured in `core-site.xml`.

By default Hadoop uses `ShellBasedUnixGroupsMapping`, which is an implementation of `GroupMappingServiceProvider`. It fetches the group membership for a username by executing a UNIX shell command. In secure clusters, since the usernames are actually Kerberos principals, `ShellBasedUnixGroupsMapping` will work only if the Kerberos principals map to valid UNIX usernames. Hadoop provides a feature that lets administrators specify mapping rules to map a Kerberos principal to a local UNIX username .

21.2.2.2. Creating Mappings Between Principals and UNIX Usernames

Hadoop uses a rule-based system to create mappings between service principals and their related UNIX usernames. The rules are specified in the `core-site.xml` configuration file as the value to the optional key `hadoop.security.auth_to_local`.

The default rule is simply named `DEFAULT`. It translates all principals in your default domain to their first component. For example, `myusername@APACHE.ORG` and `myusername/admin@APACHE.ORG` both become `myusername`, assuming your default domain is `APACHE.ORG`.

Use the following instructions to configure the mappings between principals and UNIX usernames:

1. [Creating Rules](#)
2. [Examples](#)

21.2.2.2.1. Creating Rules

- [Simple Rules](#)

To make a simple map between principal names and UNIX users, you create a straightforward substitution rule. For example, to map the JobTracker(jt) and TaskTracker(tt) principals in the EXAMPLE.COM realm to the UNIX `$MAPREDUCE_USER` user and the NameNode(nn) and DataNode(dn) principals to the UNIX `$HDFS_USER` user, you would make this the value for the `hadoop.security.auth_to_local` key in `core-site.xml`.

```
RULE:[2:$1@$0]([jt]t@.*EXAMPLE.COM)s/./$MAPREDUCE_USER /
RULE:[2:$1@$0]([nd]n@.*EXAMPLE.COM)s/./ $HDFS_USER/
DEFAULT
```

- **Complex Rules**

To accommodate more complex translations, you create a hierarchical set of rules to add to the default. Each rule is divided into three parts: base, filter, and substitution.

- **The Base:**

The base begins with the number of components in the principal name (excluding the realm), followed by a colon, and the pattern for building the username from the sections of the principal name. In the pattern section `$0` translates to the realm, `$1` translates to the first component and `$2` to the second component.

For example:

`[1:$1@$0]` translates `myusername@APACHE.ORG` to `myusername@APACHE.ORG`

`[2:$1]` translates `myusername/admin@APACHE.ORG` to `myusername`

`[2:$1%$2]` translates `myusername/admin@APACHE.ORG` to `myusername%admin`

- **The Filter:**

The filter consists of a regex in a parentheses that must match the generated string for the rule to apply.

For example:

`(.*%admin)` matches any string that ends in `%admin`

`(.*@SOME.DOMAIN)` matches any string that ends in `@SOME.DOMAIN`

- **The Substitution:**

The substitution is a sed rule that translates a regex into a fixed string.

For example:

`s/@ACME\.COM//` removes the first instance of `@SOME.DOMAIN`.

`s/[A-Z]*\.` removes the first instance of `@` followed by a name followed by `COM`.

`s/X/Y/g` replaces all of the `X` in the name with `Y`

21.2.2.2. Examples

- If your default realm was `APACHE.ORG`, but you also wanted to take all principals from `ACME.COM` that had a single component `joe@ACME.COM`, you would create this rule:

```
RULE:[1:$1@$0](.*@ACME\.COM)s/@.*//
DEFAULT
```

- To also translate names with a second component, you would use these rules:

```
RULE:[1:$1@$0](.*@ACME\.COM)s/@.*//
RULE:[2:$1@$0](.*@ACME\.COM)s/@.*//
DEFAULT
```

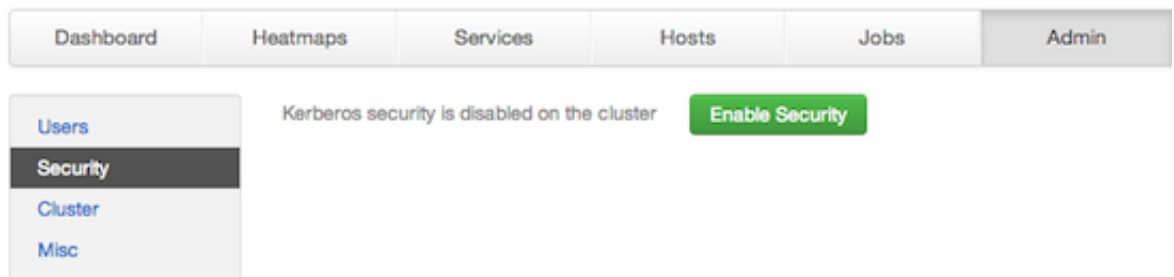
- To treat all principals from `APACHE.ORG` with the extension `/admin` as `admin`, your rules would look like this:

```
RULE[2:$1%$2@$0](.*%admin@APACHE\.ORG)s/.*%admin/
DEFAULT
```

21.2.3. Enabling Kerberos Security

To turn on Kerberos-based security in the Ambari Web GUI you must:

1. Have already set up Kerberos for your cluster. For more information, see [Setting Up Kerberos for Hadoop 1.x](#).
2. Go to the **Admin** tab.
3. Click **Security**.
4. Click **Enable Security** and follow the **Add security wizard**.



- a. **Get Started:** This step just reminds you that you need to set up Kerberos before you start.
- b. **Configure Services:** This step asks you for your Kerberos information: principals and paths to keytabs, path to your Kerberos tools, realm names and so on. For more information about a specific field, hover over it, and a popup with a definition appears.

- c. **Create Principals and Keytabs:** Use this step to check that all your information is correct. Click **Back** to make any changes. Click **Apply** when you are satisfied with the assignments.



Note

If you have a large cluster, you may want to go to the **Create Principals and Keytabs** step first. Step through the wizard accepting the defaults to get to the appropriate page. On the page, use the **Download CSV** button to get a list of all the necessary principals and keytabs in CSV form, which can be used to set up a script. The list includes hostname, principal description, principal name, keytab user, keytab group, keytab permissions, absolute keytab path, and keytab filename.

SECURITY WIZARD

Get Started

Configure Services

Create Principals and Keytabs

Save and Apply Configuration

Create Principals and Keytabs

You need to create the following principals and keytabs on the hosts shown. You can download the list as a CSV file and use it to create a script to generate the principals and keytabs. Once the principals and keytabs have been created, click on Proceed to continue. If you need to make configuration changes, click Back.

Host	Component	Principal	Keytab
FQDN	Ambari Smoke Test User	ambari-qa@EXAMPLE.COM	/etc/security/keytabs/smokeuser.headless.keytab
FQDN	HDFS User	hdfs@EXAMPLE.COM	/etc/security/keytabs/hdfs.headless.keytab
FQDN	HBase User	hbase@EXAMPLE.COM	/etc/security/keytabs/hbase.headless.keytab
FQDN	SPNEGO User	HTTP/p-10-191-37-233.ec2.internal@EXAMPLE.COM	/etc/security/keytabs/spnego.service.keytab
FQDN	DataNode	dn/p-10-191-37-233.ec2.internal@EXAMPLE.COM	/etc/security/keytabs/dn.service.keytab
FQDN	HBase Master	hbase/p-10-191-37-233.ec2.internal@EXAMPLE.COM	/etc/security/keytabs/hbase.service.keytab
FQDN	HiveServer2	hive/p-10-191-37-233.ec2.internal@EXAMPLE.COM	/etc/security/keytabs/hive.service.keytab
FQDN		j/p-10-191-37-233.ec2.internal@EXAMPLE.COM	/etc/security/keytabs/j.service.keytab

← Back
Download CSV
Apply →

- d. **Save and Apply Configuration:** This step displays a bar showing the progress of integrating the Kerberos information into your Ambari Server.