

Hortonworks Data Platform

Apache Ambari Upgrade

(May 19, 2017)

Hortonworks Data Platform: Apache Ambari Upgrade

Copyright © 2012-2017 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 4.0 License.
<http://creativecommons.org/licenses/by-sa/4.0/legalcode>

Table of Contents

1. Upgrading Ambari and HDP	1
2. Getting Ready to Upgrade Ambari and HDP	2
3. Upgrading Ambari	4
3.1. Preparing to Upgrade	4
3.2. Upgrade Ambari	5
3.3. Post-Upgrade Tasks	10
3.3.1. Upgrade Ambari Metrics	10
3.3.2. Adding Grafana to Ambari Metrics	12
3.3.3. Upgrading Configurations	13
4. Upgrading HDP	14
4.1. Prerequisites	16
4.1.1. Rolling Upgrade Prerequisites	17
4.2. Prepare to Upgrade	18
4.3. Register and Install Target Version	20
4.4. Perform the Upgrade	22
4.4.1. Perform Rolling Upgrade	23
4.4.2. Perform Express Upgrade	25
4.5. Post-upgrade Tasks	26
4.5.1. Post-upgrade Tasks for Ranger with Kerberos	26

List of Figures

3.1. installgrafanapending.png 12

List of Tables

4.1. HDP Upgrade Options 15

1. Upgrading Ambari and HDP

Ambari and the HDP Stack being managed by Ambari can be upgraded independently.

This guide provides information on:

- [Getting Ready to Upgrade Ambari and HDP](#)
- [Upgrading Ambari](#)
- [Upgrading HDP](#)



Important

Ambari 2.4 **does not support managing an HDP 2.1 or 2.0** cluster. If you are running HDP 2.1 or 2.0, in order to use Ambari 2.4 **you must first upgrade to HDP 2.2 or higher** using either Ambari 2.2.2/2.2.1/2.2/2.1/2.0 **before upgrading to Ambari 2.4**. Once completed, upgrade your current Ambari to Ambari 2.4.

2. Getting Ready to Upgrade Ambari and HDP

When preparing to upgrade Ambari and the HDP Cluster, we strongly recommend you review this checklist of items to confirm your cluster operation is healthy. Attempting to upgrade a cluster that is operating in an unhealthy state can produce unexpected results.



Important

If planning to upgrade Ambari and upgrade to a new HDP minor version (for example: moving from HDP 2.2 to HDP 2.3), upgrade Ambari to the latest version before upgrading the cluster. Refer to the [Stack Compatibility Matrix](#) to see which HDP versions each Ambari version supports.

- Ensure all services in the cluster are running.
- Run each Service Check (found under the Service Actions menu) and confirm they execute successfully.
- Clear all alerts, or understand why they are being generated. Remediate as necessary.
- Confirm start and stop for all services are executing successfully.
- Time service start and stops. The time to start and stop services is a big contributor to overall upgrade time so having this information handy is useful.
- Download the software packages prior to the upgrade. Place them in a local repository and/or consider using a storage proxy since multi-gigabyte downloads will be required on all nodes in the cluster. Refer to [Using a Local Repository](#) for more information.
- Ensure point-in-time backups are taken of all DBs supporting the clusters. This includes (among others) Ambari, Hive Metastore, Ranger and Oozie.



Note

If you are upgrading to Ambari 2.4, you must make sure that the Ranger db root password is NOT set to blank before performing any HDP upgrade. Ambari requires that the Ranger db root password has a value. If you upgrade Ambari to 2.2 and upgrade HDP without setting a value for the Ranger db root password, Ranger Admin will fail to start after the upgrade.

To prepare Ranger for upgrades, set the password for the Ranger DB root user to a non-blank value. Then, set the Ranger db root password field in Ambari Web to match this value. Finally, restart Ranger Admin using Ambari Web.

For Ambari Upgrades

- This (Ambari 2.4) Upgrade Guide will help you upgrade your existing Ambari install to version 2.4. If you are upgrading to another Ambari version, please be sure to use the Ambari Upgrade Guide for that version.

- Be sure to review the Known Issues and Behavioral Changes for this Ambari release in the [Ambari 2.4.3.0 Release Notes](#).

For HDP Cluster Upgrades

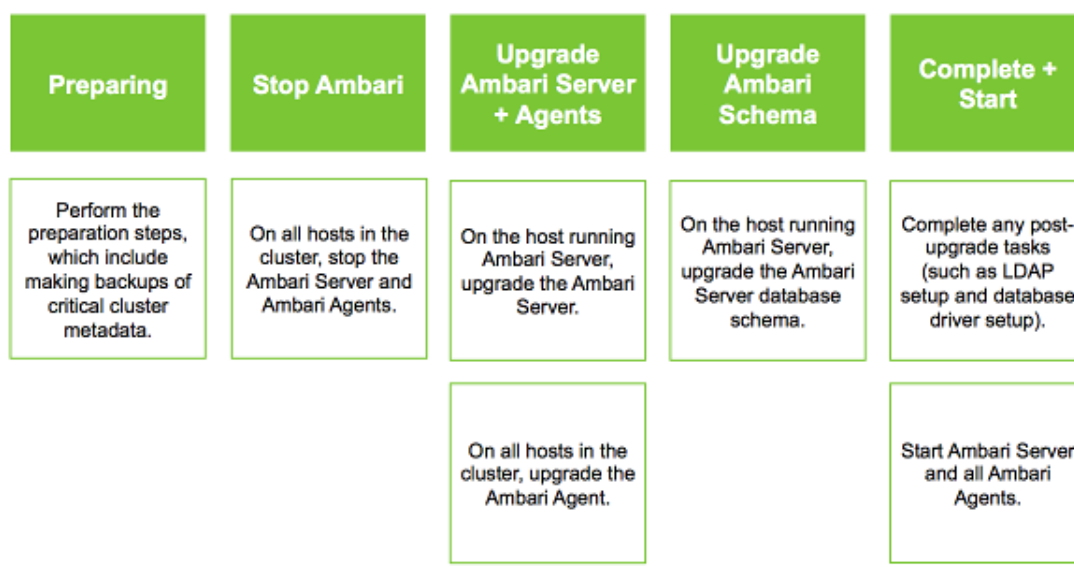
- Ensure sufficient disk space on `/usr/hdp/<version>` (roughly 3GB for each additional HDP release).
- If you plan to add new services available with HDP to your cluster, the new services might include new service accounts. Any operational procedures required to support these new service accounts should be performed prior to the upgrade. The accounts will typically be required on all nodes in the cluster.
- If your cluster includes Storm, document any running Storm topologies.

3. Upgrading Ambari

Ambari and the HDP cluster being managed by Ambari can be upgraded independently. This section describes the process to upgrade Ambari. You are **strongly encouraged** to read completely through this entire document before starting the upgrade process, to that you understand the interdependencies and order of the steps. It is **highly recommended** you validate these steps in a test environment to adjust + account for any special configurations for your cluster.

- [Preparing to Upgrade](#)
- [Upgrade Ambari](#)
- [Post-Upgrade Tasks](#)

The high-level process for upgrading Ambari is as follows:



3.1. Preparing to Upgrade

- Be sure to review the [Release Notes](#) for this Ambari release for Known Issues and Behavioral Changes.
- You **must** have root, administrative, or root-equivalent authorization on the Ambari server host and all servers in the cluster.
- You **must** backup the Ambari Server database.
- You **must** make a safe copy of the Ambari Server configuration file found at `/etc/ambari-server/conf/ambari.properties`.
- **Plan to upgrade the Ambari Metrics service:**
 - Record the location of the **Metrics Collector** component before you begin the upgrade process.

- You **must** stop the Ambari Metrics service from **Ambari Web**.
- After upgrading Ambari, you must also [Upgrade Ambari Metrics](#) service and [add the Grafana](#) component.
- Proceed to [Upgrade to Ambari 2.4](#).
- After upgrading Ambari, you must also [Upgrade Hortonworks SmartSense](#).



Important

If your current Ambari version is 1.7 or below, you must [upgrade the Ambari Server version to 2.2](#) before upgrading to version 2.4.



Note

During Ambari upgrade, the existing `/var/lib/ambari-server/ambari-env.sh` file is overwritten and a backup copy of `ambari-env.sh` (with extension `.rpmsave`) is created. If you have manually modified `ambari-env.sh` (for example, to change Ambari Server heap), you will need to re-apply your changes to the new file.

3.2. Upgrade Ambari

1. If you are running Ambari Metrics service in your cluster, stop the service. From **Ambari Web**, browse to **Services > Ambari Metrics** and select **Stop** from the **Service Actions** menu.
2. Stop the Ambari Server. On **the host** running Ambari Server:

```
ambari-server stop
```
3. Stop all Ambari Agents. On **each host** in your cluster running an Ambari Agent:

```
ambari-agent stop
```
4. Fetch the new Ambari repo and replace the old repository file with the new repository file **on all hosts** in your cluster.



Important

Check your current directory before you download the new repository file to make sure that there are no previous versions of the `ambari.repo` 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.

Select the repository appropriate for your environment from the following list:

- **For RHEL/CentOS/Oracle Linux 6:**

```
wget -nv http://public-repo-1.hortonworks.com/ambari/centos6/2.x/updates/2.4.3.0/ambari.repo -O /etc/yum.repos.d/ambari.repo
```

- **For RHEL/CentOS/Oracle Linux 7:**

```
wget -nv http://public-repo-1.hortonworks.com/ambari/centos7/2.x/updates/2.4.3.0/ambari.repo -O /etc/yum.repos.d/ambari.repo
```

- **For SLES 11:**

```
wget -nv http://public-repo-1.hortonworks.com/ambari/suse11/2.x/updates/2.4.3.0/ambari.repo -O /etc/zypp/repos.d/ambari.repo
```

- **For SLES 12:**

```
wget -nv http://public-repo-1.hortonworks.com/ambari/sles12/2.x/updates/2.4.3.0/ambari.repo -O /etc/zypp/repos.d/ambari.repo
```

- **For Ubuntu 12:**

```
wget -nv http://public-repo-1.hortonworks.com/ambari/ubuntu12/2.x/updates/2.4.3.0/ambari.list -O /etc/apt/sources.list.d/ambari.list
```

- **For Ubuntu 14:**

```
wget -nv http://public-repo-1.hortonworks.com/ambari/ubuntu14/2.x/updates/2.4.3.0/ambari.list -O /etc/apt/sources.list.d/ambari.list
```

- **For Debian 7:**

```
wget -nv http://public-repo-1.hortonworks.com/ambari/debian7/2.x/updates/2.4.3.0/ambari.list -O /etc/apt/sources.list.d/ambari.list
```



Note

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



Note

Ambari Server does not automatically turn off `iptables`. Check that your installation setup does not depend on `iptables` being disabled. After upgrading the server, you must either disable `iptables` manually or make sure that you have appropriate ports available on all cluster hosts. For more information about ports, see [Configuring Network Port Numbers](#).

5. Upgrade Ambari Server. On the host running Ambari Server:

- **For RHEL/CentOS/Oracle Linux:**

```
yum clean all  
yum info ambari-server
```

In the info output, visually validate that there is an available version containing "2.4"

```
yum upgrade ambari-server
```

- **For SLES:**

```
zypper clean
zypper info ambari-server
```

In the info output, visually validate that there is an available version containing "2.4"

```
zypper up ambari-server
```

- **For Ubuntu/Debian:**

```
apt-get clean all
apt-get update
apt-cache show ambari-server | grep Version
```

In the info output, visually validate that there is an available version containing "2.4"

```
apt-get install ambari-server
```



Important

When performing upgrade on SLES, you will see a message "There is an update candidate for 'ambari-server', but it is from different vendor. Use 'zypper install ambari-server-2.4-101.noarch' to install this candidate". You will need to use yast to update the package, as follows:

- a. From the command line run: > yast.

```
> yast
```

You will see command line UI for YaST program.

- b. Choose `Software > Software Management`, then click the `Enter` button.
- c. In the `Search Phrase` field, enter `ambari-server`, then click the `Enter` button.
- d. On the right side you will see the search result `ambari-server 2.4`. Click `Actions`, choose `Update`, then click the `Enter` button.
- e. Go to `Accept`, and click `enter`.

6. Check for upgrade success by noting progress during the Ambari Server installation process you started in Step 5.

- As the process runs, the console displays output similar, although not identical, to the following:

```
Setting up Upgrade Process
Resolving Dependencies
--> Running transaction check
```

- If the upgrade fails, the console displays output similar to the following:

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

- A successful upgrade displays output similar to the following:

```
Updated:
  ambari-server.noarch 0:2.4-111
Complete!
```



Note

Confirm there is only one `ambari-server*.jar` file in `/usr/lib/ambari-server`. If there is more than one JAR file with name `ambari-server*.jar`, move all JARs except `ambari-server-2.4.*.jar` to `/tmp` before proceeding with upgrade.

7. Upgrade all Ambari Agents. On **each host** in your cluster running an Ambari Agent:

- **For RHEL/CentOS/Oracle Linux:**

```
yum upgrade ambari-agent
```

- **For SLES:**

```
zypper up ambari-agent
```



Note

Ignore the warning that begins with "There are some running programs that use files deleted by recent upgrade".



Important

When performing upgrade on SLES, you will see a message "There is an update candidate for 'ambari-agent', but it is from different vendor. Use 'zypper install ambari-agent-2.4-101.noarch' to install this candidate". You will need to use `yast` to update the package, as follows:

- a. From the command line run: `> yast`

```
> yast
```

You will see command line UI for YaST program.

- b. Choose `Software > Software Management`, then click the `Enter` button.
- c. In the `Search Phrase` field, enter `ambari-agent`, then click the `Enter` button.
- d. On the right side you will see the search result `ambari-agent 2.4`. Click `Actions`, choose `Update`, then click the `Enter` button.
- e. Go to `Accept`, and click `enter`.

- **For Ubuntu/Debian:**

```
apt-get update
apt-get install ambari-agent
```

8. After the upgrade process completes, check each host to make sure the new files have been installed:

- **For RHEL/CentOS/Oracle Linux 6:**

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

- **For RHEL/CentOS/Oracle Linux 7:**

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

- **For SLES 11:**

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

- **For SLES 12:**

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

- **For Ubuntu 12:**

```
dpkg -l ambari-agent
```

- **For Ubuntu 14:**

```
dpkg -l ambari-agent
```

- **For Debian 7:**

```
dpkg -l ambari-agent
```

9. Upgrade Ambari Server database schema. On **the host** running Ambari Server:

```
ambari-server upgrade
```

10 Start the Ambari Server. On **the host** running Ambari Server:

```
ambari-server start
```

11 Start all Ambari Agents. On **each host** in your cluster running an Ambari Agent:

```
ambari-agent start
```

12 Open Ambari Web.

Point your browser to `http://<your.ambari.server>:8080`

where `<your.ambari.server>` is the name of your ambari server host. For example, `c6401.ambari.apache.org`.



Important

Refresh your browser so that it loads the new version of the Ambari Web code. If you have problems, clear your browser cache manually, then restart Ambari Server.

13. Log in, using the Ambari administrator credentials that you have set up.

For example, the default name/password is `admin/admin`.

You will see a Restart indicator next to each service after upgrading. Ambari upgrade has added to/adjusted the configuration properties of your cluster based on new configuration types and properties being made available for each service with this release of Ambari. Review these changes by comparing the previous configuration with the latest version created by "ambari-upgrade".

14. If you have configured Ambari to authenticate against an external LDAP or Active Directory, you **must** re-run "ambari-server setup-ldap". For more information, see [Set Up LDAP or Active Directory Authentication](#).

15. If you have configured your cluster for Hive or Oozie with an external database (Oracle, MySQL or PostgreSQL), you **must** re-run "ambari-server setup -jdbc-db and -jdbc-driver" to get the JDBC driver JAR file in place. For more information, see [Using Non-Default Databases - Hive](#) and [Using Non-Default Databases - Oozie](#).

16. If you are running **Ambari Metrics** service in your cluster, [Upgrade Ambari Metrics](#) service and [add the Grafana](#) component.

17. If your cluster includes the SmartSense service, you must upgrade it along with Ambari, [Upgrade SmartSense Service](#).

18. Perform any other [Post-Upgrade Tasks](#) as required.

3.3. Post-Upgrade Tasks

Depending on the configuration of your cluster and the Ambari version you are starting with, review the following post-upgrade tasks.

Task	Description
Upgrading Ambari Metrics	If your cluster includes the Ambari Metrics System ("AMS") service, you must upgrade the system along with Ambari. This will include adding the Grafana component to the system.
Adding Grafana to Ambari Metrics	Grafana is now included as a component of Ambari Metrics. If you are upgrading from Ambari 2.2.1 or earlier, and your Ambari Metrics service does not contain Grafana, proceed to add Grafana to Ambari Metrics.
Upgrading SmartSense	If your cluster includes the SmartSense service, you must upgrade it along with Ambari.
Upgrading Configurations	Certain scenarios may require that you modify configurations that Ambari did not upgrade automatically. Review scenarios described in this section for information about upgrading configurations in your cluster.

3.3.1. Upgrade Ambari Metrics

1. Upgrade to Ambari 2.4 and perform needed post-upgrade checks. Make sure all services are up and healthy.
2. Make sure Ambari Metrics service is stopped. From Ambari Web, browse to Services > Ambari Metrics and select Stop from the Service Actions menu.

3. On every host in your cluster running a Metrics Monitor, run the following commands:

For RHEL/CentOS/Oracle Linux:

```
yum clean all
yum upgrade ambari-metrics-monitor ambari-metrics-hadoop-sink
```

For SLES:

```
zypper clean
zypper up ambari-metrics-monitor ambari-metrics-hadoop-sink
```

For Ubuntu/Debian:

```
apt-get clean all
apt-get update
apt-get install ambari-metrics-assembly
```

4. Execute the following command on all hosts running the Metrics Collector:

For RHEL/CentOS/Oracle Linux:

```
yum upgrade ambari-metrics-collector
```

For SLES:

```
zypper up ambari-metrics-collector
```

For Ubuntu/Debian:

```
apt-get clean all
apt-get update
apt-get install ambari-metrics-collector
```

5. Execute the following command on the host running the Grafana component:

For RHEL/CentOS/Oracle Linux:

```
yum upgrade ambari-metrics-grafana
```

For SLES:

```
zypper up ambari-metrics-grafana
```

For Ubuntu/Debian:

```
apt-get clean all
apt-get update
apt-get install ambari-metrics-grafana
```

6. Start Ambari Metrics Service.

From **Ambari Web**, browse to **Services > Ambari Metrics** select **Service Actions** then choose **Start**.

Updated Ambari Metrics Sink jars will be installed on all hosts. You must restart each service to pick up the latest sink implementations.

(For example: HDFS, YARN, Kafka, HBase, Flume, Storm)

- If you are upgrading from Ambari 2.2.1 or earlier, and your Ambari Metrics service does not contain Grafana, proceed to [Add Grafana to Ambari Metrics](#).

3.3.2. Adding Grafana to Ambari Metrics

As of Ambari 2.4, Grafana is now included as a component of Ambari Metrics. You must add Grafana to the system and install Grafana on a host in the cluster.



Note

When using the API commands below, be sure to replace the **ambari.server** with the Ambari Server hostname, **cluster.name** with your cluster name and **host.name** with the host where you will run Grafana. This can be the same host that is running the Metrics Collector.

- [Upgrade to Ambari 2.4](#) and perform needed post-upgrade checks. Make sure all services are up and healthy.
- Add the METRICS_GRAFANA component to Ambari:

```
curl -u admin:admin -H "X-Requested-By:ambari" -i -X POST
http://ambari.server:8080/api/v1/clusters/cluster.name/services/
AMBARI_METRICS/components/METRICS_GRAFANA
```

- Add METRICS_GRAFANA to a host in the cluster.

```
curl -u admin:admin -H "X-Requested-By:ambari" -i -X POST -d
'{"host_components":[{"HostRoles":{"component_name":"METRICS_GRAFANA"}}]}'
http://ambari.server:8080/api/v1/clusters/cluster.name/hosts?Hosts/
host_name=host.name
```

- From **Ambari Web**, browse to **Services > Ambari Metrics** and you will see Grafana is in the **Install Pending...** state. You need to complete the configuration of Grafana before installing and starting.
- To complete the configuration, click on **Services > Ambari Metrics > Configs** and you will need to enter the default Grafana Admin Password in the **General** section. Click **Save**.
- Browse to **Hosts > host.name** (i.e. the **host.name** used in the API call where you added Grafana). You will see the Grafana component is in an **Install Pending...** state. Use the **Install Pending...** action button and select **Re-install**.

Figure 3.1. installgrafanapending.png

 Grafana / Ambari Metrics

Install Pending... ▾

- Once the install operation completes, select **Start** to start Grafana.
- To access Grafana, browse to **Services > Ambari Metrics**, select **Quick Links** and then choose **Grafana**.
- Refer to Hortonworks Data Platform Apache Ambari User Guide for more information on [Using Grafana](#).

3.3.3. Upgrading Configurations

This section describes potential cluster configuration updates that may be required.

Kerberos krb5.conf

Ambari has added support for handling more than one KDC host. Only one kadmin host is supported by the Kerberos infrastructure. This required modifications for the **krb5.conf** template. In order for Ambari to properly construct the krb5.conf configuration file, make the following configuration change if your cluster meets all of these criteria:

- Kerberos is enabled and [Ambari is configured for Automated Setup](#), and
- Ambari is managing the krb5.conf, and
- You **have modified** the krb5.conf template content from the default content. If you have not modified the default content, Ambari will automatically update the template content as part of upgrade and these configuration updates do not need to be applied manually.

If you meet all of the above criteria, you must update the **krb5.conf** template content found in **Services > Kerberos > Advanced**:

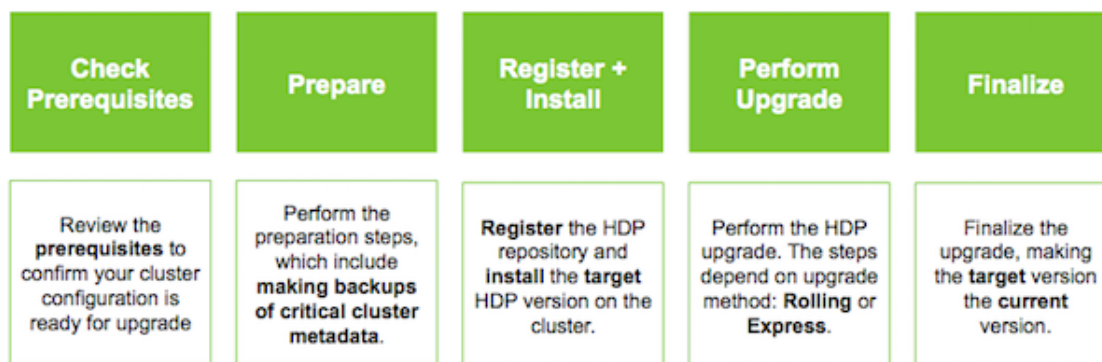
Original Template Entry	Updated Template Entry
admin_server = {{ admin_server_host default(kdc_host, True)}}	admin_server = {{ admin_server_host default(kdc_host_list[0] trim(), True)}}
kdc = {{ kdc_host}}	{% for kdc_host in kdc_host_list %} kdc = {{ kdc_host trim()}} {%- endfor -%}

4. Upgrading HDP

- [Prerequisites](#)
- [Prepare to Upgrade](#)
- [Register and Install Target Version](#)
- [Perform the Upgrade](#)

There are different HDP upgrade options based on your current HDP and the target HDP. This section describes the different upgrade options, their prerequisites, and the overall process. You are **strongly encouraged** to read completely through this entire document before starting the upgrade process, to that you understand the interdependencies and order of the steps. It is **highly recommended** you validate these steps in a test environment to adjust and account for any special configurations for your cluster.

The high-level process for performing an HDP upgrade is as follows:



Ambari will guide you through the steps required to upgrade HDP. Make sure Ambari and the cluster are healthy, operating normally and all service checks are passing. See [Preparing to Upgrade Ambari and HDP](#) for more information on performing the Ambari upgrade.



Note















Be sure to review the available HDP upgrade scenarios below. It is **strongly recommended** that you **first upgrade to Ambari 2.4** before upgrading HDP unless otherwise noted. After upgrading Ambari, be sure the cluster is operating normally and service checks are passing prior to attempting an HDP upgrade.

There are two methods for upgrading HDP with Ambari: **Rolling Upgrade** and **Express Upgrade**.

- A **Rolling Upgrade** orchestrates the HDP upgrade in an order that is meant to preserve cluster operation and minimize service impact during upgrade. This process has more [stringent prerequisites](#) (particularly regarding cluster high availability configuration) and can take longer to complete than an Express Upgrade.
- An **Express Upgrade** orchestrates the HDP upgrade in an order that will incur cluster downtime but with less stringent prerequisites.

The following table describes which upgrade option is available based on the current and target HDP version combinations. It is important to note that the HDP Stack version is based on the following format: **major.minor.maintenance.patch-build#**. A "minor" release is a release that increments the second-digit. A "maintenance" release is one that increments the third-digit. This terminology is used in the following table to describe the different upgrade paths. For example: if you want to upgrade from HDP 2.2 to HDP 2.3, that is a "Minor Upgrade". If you want to upgrade from HDP 2.2.x to HDP 2.2.y, that is a "Maintenance" Upgrade.

Table 4.1. HDP Upgrade Options

Upgrade Path	Current HDP*	Target HDP*	Rolling Upgrade	Express Upgrade
Maintenance	HDP 2.5.x.y	HDP 2.5.x.y		
Maintenance	HDP 2.4.x.y	HDP 2.4.x.y		
Maintenance	HDP 2.3.x.y	HDP 2.3.x.y		
Maintenance	HDP 2.2.x.y	HDP 2.2.x.y		
Minor **	HDP 2.4	HDP 2.5		
Minor **	HDP 2.3	HDP 2.4 or 2.5		
Minor ***	HDP 2.2	HDP 2.3 or 2.4		
Minor	HDP 2.0 or 2.1	HDP 2.2, 2.3, or 2.4	Ambari 2.2.2 does not support managing an HDP 2.1 or 2.0 cluster. You must first upgrade to HDP 2.2 or higher using either Ambari 2.2.1, 2.2, 2.1 or 2.0 prior to upgrading to Ambari 2.2.2.	

* The instructions in this document sometimes refer to an HDP "version" using the HDP #.#.x.y convention. For example, you will see the HDP 2.3.x.y convention to refer to any HDP 2.3 "maintenance" release. Therefore, to use a specific HDP 2.3 maintenance release, be sure to replace 2.3.x.y in the following instructions with the appropriate maintenance version, such as 2.3.0.0 for the HDP 2.3 GA release, or 2.3.2.0 for an HDP 2.3 maintenance release.

** A Minor upgrade can be performed from any current maintenance release to any target maintenance release in that minor release. For example, you can go from HDP 2.2.x.y to HDP 2.3.x.y as part of a Minor upgrade.

*** A direct upgrade from HDP 2.2 to HDP 2.5 is not available. You must first go from HDP 2.2 to HDP 2.4, then HDP 2.4 to HDP 2.5.



Important

Because HDP 2.3.6 contains features and fixes that are not applicable to HDP 2.4.0 or HDP 2.4.2, do not upgrade from HDP 2.3.6 to HDP 2.4.0 or HDP 2.4.2.

4.1. Prerequisites

To perform an HDP upgrade using Ambari, your cluster must meet the following prerequisites. These prerequisites are required regardless of Rolling or Express upgrade because these checks are important for Ambari to know the cluster is in a healthy operating mode and can be successfully managed from Ambari. However, there are additional [prerequisites for a Rolling Upgrade](#) that require your cluster to be configured a certain way to support the rolling upgrade orchestration process.

Check	Description
HDP Version	All hosts must have the target HDP version installed. See the Register and Install Target Version section for more information.
Disk Space	Be sure to have adequate space on <code>/usr/hdp</code> for the target HDP version. Each complete install of an HDP version will occupy about 2.5 GB of disk space.
Ambari Agent Heartbeats	All Ambari Agents must be communicating and heartbeating to Ambari Server. Any hosts that are not heartbeating must be in Maintenance Mode.
Host Maintenance Mode	The following two scenarios are checked: <ul style="list-style-type: none"> Any hosts in Maintenance Mode must not be hosting any Service Master Components. Any host in Maintenance Mode that is not hosting Master Components is allowed but you will receive a warning. You can proceed with your upgrade but these hosts will not be upgraded and before you can finalize the upgrade, you must delete the hosts from the cluster.
Service Maintenance Mode	No Services can be in Maintenance Mode.
Services Started	All Services must be started.
Service Checks	All Service Checks must pass. Be sure to run Service Actions > Run Service Check on all services (and remediate if necessary) prior to attempting an HDP upgrade.
Atlas Check	If Atlas is installed in your HDP 2.3 or 2.4 cluster, we strongly advise upgrading your cluster, as follows: <ol style="list-style-type: none"> Upgrade Ambari to version 2.4.x. Use the Remove Service capability to completely remove Atlas from your HDP 2.3 or 2.4 cluster. Do NOT remove Atlas service via REST API. Upgrade your cluster to HDP-2.5.x. Add Atlas (version 0.7x) to your cluster. <p>If you are using Atlas on HDP-2.3 or 2.4; Do Not upgrade to Ambari-2.4.x until you are ready to upgrade to HDP-2.5x (which includes Atlas 0.7x).</p>



Important

If your cluster has Kerberos enabled, includes Kafka and you are running HDP 2.2, since Kafka in HDP 2.2 did not support running Kerberized, after upgrade to HDP 2.3 or later (regardless if you perform an Express or Rolling Upgrade), Kafka will then be Kerberos enabled. Plan to adjust your use of Kafka accordingly. Ambari will perform a check to see if your cluster would encounter this issue and warn you prior to performing the HDP upgrade.

4.1.1. Rolling Upgrade Prerequisites

To perform a **Rolling Upgrade**, your cluster must meet the following prerequisites. If you do not meet these upgrade prerequisites, you can consider an **Express Upgrade**.

HDFS

Requirement	Description
NameNode HA	NameNode HA must be enabled and working properly. From Ambari Web, the NameNodes, ZKFailoverControllers and JournalNodes must be "green" and showing no alerts. See the Ambari User's Guide for more information, Configuring NameNode High Availability .
NameNode Truncate	In HDP 2.2, truncate must be disabled and this prerequisite is REQUIRED . There is an hdfs-site configuration <code>dfs.allow.truncate</code> , This configuration property is set to true if truncate is enabled. If the property is not set or is false, truncate is disabled. In 2.3 onwards, truncate is always enabled, which is acceptable. Therefore, this check is not applicable and will not be performed with HDP 2.3 maintenance upgrades.

YARN

Requirement	Description
Timeline Service State Recovery	YARN state recovery should be enabled. Check the Services > YARN > Configs > Advanced property <code>yarn.timeline-service.recovery.enabled</code> is set to true .
ResourceManager Work Preserving Recovery	YARN work preserving recovery should be enabled. Check the Services > YARN > Configs > Advanced property <code>yarn.resourcemanager.work-preserving-recovery.enabled</code> is set to true.
ResourceManager HA	ResourceManager HA should be enabled to prevent a disruption in service during the upgrade. See the Ambari User's Guide for more information on Configuring ResourceManager High Availability . This prerequisite is OPTIONAL .

MapReduce2

Requirement	Description
MapReduce Distributed Cache	MapReduce should reference Hadoop libraries from the distributed cache in HDFS. Refer to the YARN Resource Management guide for more information.
JobHistory State Recovery	JobHistory state recovery should be enabled. Check the following: Services > MapReduce > Configs > Advanced property <code>mapreduce.jobhistory.recovery.enable</code> is set to true . Once enabled, dependent properties for <code>mapreduce.jobhistory.recovery.store.class</code> and <code>mapreduce.jobhistory.recovery.store.leveldb.path</code> should also be set.
Encrypted Shuffle	If encrypted shuffle has been enabled, the <code>ssl-client.xml</code> file must be copied to <code>/etc/hadoop/conf/secure</code> directory on each node in the cluster.

Tez

Requirement	Description
Tez Distributed Cache	Tez should reference Hadoop libraries from the distributed cache in HDFS. Check the Services > Tez > Configs configuration and confirm <code>tez.lib.uris</code> is set with the path in HDFS for the <code>tez.tar.gz</code> file.

Hive

Requirement	Description
Hive Dynamic Service Discovery	HiveServer2 dynamic service discovery should be configured for Rolling Upgrade.
HiveServer2 Port Change	During the upgrade, Ambari will switch the HiveServer2 port from 10000 to 10010 (or 10011 if using HTTP transport mode). In order to support rolling upgrades, the Hive server is required to change its port. Applications and users which use a URL that includes the port will no longer be able to connect after Hive server has upgraded. If this behavior is not desired, then the port can be restored to its original value after the upgrade has been finalized.
Hive Client Retry	Hive client retry properties must be configured. Review the Services > Hive > Configs configuration and confirm <code>hive.metastore.failure.retries</code> and <code>hive.metastore.client.connect.retry.delay</code> are specified.
Multiple Hive Metastore	Multiple Hive Metastore instances are recommended for Rolling Upgrade. This ensures that there is at least one Hive Metastore running during the upgrade process. Additional Hive Metastore instances can be added through Ambari. This prerequisite is OPTIONAL .

Oozie

Requirement	Description
Oozie Client Retry	Oozie client retry properties must be configured. Review the Services > Oozie > Configs > oozie-env configuration and confirm <code>"export OOZIE_CLIENT_OPTS=\${OOZIE_CLIENT_OPTS} - Doozie.connection.retry.count=<number of retries>"</code> is specified.

4.2. Prepare to Upgrade

Review [Preparing to Upgrade Ambari and HDP](#).

It is also **strongly** recommended that you perform backups of your databases before beginning upgrade.

- Ambari database
- Hive Metastore database
- Oozie Server database
- Ranger Admin database
- Ranger Audit database

Checkpoint HDFS

1. Perform the following steps on the NameNode host. If you are configured for NameNode HA, perform the following steps on the Active NameNode. You can locate the Active NameNode from **Ambari Web > Services > HDFS** in the **Summary** area.
2. Check the NameNode directory to ensure that there is no snapshot of any prior HDFS upgrade. Specifically, using Ambari Web, browse to **Services > HDFS > Configs**, and examine the `dfs.namenode.name.dir` in the **NameNode Directories** property. Make sure that only a `"/current"` directory and no `"/previous"` directory exists on the NameNode host.
3. Create the following log and other files. Creating these logs allows you to check the integrity of the file system after the Stack upgrade.

As the HDFS user, `"su -l <HDFS_USER>"` run the following (where `<HDFS_USER>` is the HDFS Service user, for example, `hdfs`):

- 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.

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

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

```
hdfs dfsadmin -report > dfs-old-report-1.log
```

- **Optional:** Capture the complete namespace of the file system. The following command does a recursive listing of the root file system:

```
hdfs dfs -ls -R / > dfs-old-lsr-1.log
```

- **Optional:** Copy all unrecoverable data stored in HDFS to a local file system or to a backup instance of HDFS.

4. Save the namespace. As the HDFS user, `"su -l <HDFS_USER>"`, you must put the cluster in Safe Mode.

```
hdfs dfsadmin -safemode enter
```

```
hdfs dfsadmin -saveNamespace
```



Note

In a highly-available NameNode configuration, the command `hdfs dfsadmin -saveNamespace` sets a checkpoint in the first NameNode specified in the configuration, in `dfs.ha.namenodes.[nameserviceID]`.

You can also use the `dfsadmin -fs` option to specify which NameNode to connect. For example, to force a checkpoint in NameNode2:

```
hdfs dfsadmin -fs hdfs://namenode2-hostname:namenode2-port -saveNamespace
```


- Copy the checkpoint files located in `${dfs.namenode.name.dir}/current` into a backup directory.



Note

In a highly-available NameNode configuration, the location of the checkpoint depends on where the `saveNamespace` command is sent, as defined in the preceding step.

- Store the `layoutVersion` for the NameNode located at

`${dfs.namenode.name.dir}/current/VERSION`, into a backup directory where

`${dfs.namenode.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.

- As the HDFS user, "su -l <HDFS_USER>", take the NameNode out of Safe Mode.

```
hdfs dfsadmin -safemode leave
```

- Finalize any prior HDFS upgrade, if you have not done so already. As the HDFS user, "su -l <HDFS_USER>", run the following:

```
hdfs dfsadmin -finalizeUpgrade
```



Important

In HDP-2.5, Audit to DB is no longer supported in Apache Ranger. If you are upgrading from HDP-2.3 or HDP-2.4 to HDP-2.5, you should [use Apache Solr for Ranger audits](#). You should also [migrate your audit logs from DB to Solr](#).

Falcon

If your cluster includes Falcon, you must provide the Berkeley DB JAR file, under the Oracle license, as part of a Falcon installation or upgrade.

- [Falcon Prerequisites for HDP 2.5.0](#)
- [Falcon Prerequisites for HDP 2.5.3](#)

Proceed to [Register and Install Target Version](#).

4.3. Register and Install Target Version

This section describes the steps to register the software repositories with Ambari for the new target version (i.e. the version you will be upgrading to) and how to install the software on all hosts prior to performing the upgrade.

- [Register Target Version \[21\]](#)
- [Install Target Version \[22\]](#)



Note

Refer to the Ambari Admin Guide for more information on [registering and managing versions](#).

Register Target Version

1. Log in to Ambari.
2. Browse to **Admin > Stack and Versions**.
3. Click on the **Versions** tab. You will see the version currently running, marked as "Current".



Note

The full version is dependent on the HDP version you are actually running. For example, if you are currently running the HDP 2.3.0.0 release, you would see something like **HDP-2.3.0.0-2557** as the full version number.

4. Click **Manage Versions**.
5. Proceed to register a new version by clicking + **Register Version**.
6. Select the software version and method of delivery for your cluster.
 - **Choose HDP Stack.** The available HDP versions are shown in TABs. When you select a TAB, Ambari attempts to discover what specific version of that HDP Stack is available. That list is shown in a DROPDOWN. For that specific version, the available Services are displayed, with their Versions shown in the TABLE.
 - **Choose HDP Version.** If Ambari has access to the Internet, the specific Versions will be listed as options in the DROPDOWN. If you have a Version Definition File for a version that is not listed, you can click **Add Version...** and upload the VDF file. In addition, a **Default Version Definition** is also included in the list if you do not have Internet access or are not sure which specific version to install. If you choose the **Default Version Definition**, you must enter a "two-digit Version Number" in the **Name** input field.
 - **Choose Repository Delivery Method.** Using a Public Repository requires Internet connectivity. Using a Local Repository requires you have configured the software in a repository available in your network. To use the public software repositories, see the list of available [HDP Repositories](#) for each OS. Or, if you are using a local repository, enter the Base URLs for the local repository you have created. Refer to [Using a Local Repository](#) for more information in setting-up a local repository in your environment.
7. Click Save.
8. Click **Go To Dashboard**, and browse back to to **Admin > Stack and Versions > Versions**.
9. You will see the version currently running, marked "Current", and the version you just registered **HDP-2.3.2.0** which is showing an **Install** button.
10. Proceed to [Install Target Version \[22\]](#).

Install Target Version

1. Log in to Ambari.
2. Browse to **Admin > Stack and Versions**.
3. Click on the **Versions** tab.
4. Click **Install** and click **OK** to confirm.
5. The Install version operation will start. This will install the target version on all hosts in the cluster. You can monitor the progress of the install by clicking the "Installing" link.



Note

If you have Apache Ranger on your cluster, confirm that the `ambari-infra-solr-client` package is installed on the Ranger Admin host. The [Ambari Infra](#) service contains the default Solr instance used by Apache Ranger.

6. Once the install is complete, the **Install** button will be replaced with the option to **Upgrade**.
7. Proceed to [Perform the Upgrade](#).

4.4. Perform the Upgrade

1. Log in to Ambari.
2. Browse to **Admin > Stack and Versions**.
3. Click on the **Versions** tab. The [registered and installed](#) target HDP version should be visible.
4. Click **Perform Upgrade** on the target version.
5. Based on your current HDP version and the target HDP version, Ambari will perform a set of [prerequisite checks](#) to determine if you can perform a **Rolling Upgrade** or an **Express Upgrade**. A dialog will be presented with the options available.
6. Select your upgrade method: **Rolling Upgrade** or **Express Upgrade**. There are also advanced options available.

Option	Description
Skip all Service Check failures	Ambari will automatically skip any Service Check failures and complete the task without requiring user intervention to continue. After all the Service Checks have run in a task, you will be presented with summary of the failures and an option to continue the upgrade or pause.
Skip all Slave Component failures	Ambari will automatically skip any Slave Component failures and complete the task of upgrading Slave components without requiring user intervention to continue. After all Slave Components have been upgraded, you will be presented with a summary of the failures and an option to continue the upgrade or pause.

7. Click **Proceed**. Based on your option choice, proceed with the upgrade:

- If you selected Rolling Upgrade , proceed to [Perform Rolling Upgrade](#).
- If you selected Express Upgrade , proceed to [Perform Express Upgrade](#).

4.4.1. Perform Rolling Upgrade

1. Ambari will check that your cluster meets the [Rolling Upgrade Prerequisites](#). A dialog will be presented with the results:
 - a. If any prerequisites are not met but are required, the result will be shown with an error. You will not be allowed to proceed with the upgrade. Make the appropriate corrections and return to Perform Upgrade again.
 - b. If any prerequisites are not met but are optional, the result will be shown with a warning. You will be allowed to Proceed with the upgrade.
 - c. A list of configuration changes (if any) will be displayed.
2. Once the prerequisite checks are complete, the upgrade will start. The time it takes to perform the upgrade depends on many factors. As part of the upgrade process, each component in the cluster is restarted in a serial fashion as such, the stop/start time is a big contributor to the overall time.
3. The upgrade process involves a set of stages. This table lists the high-level stages and if the process requires any action by you during normal operation.



Note

If any stage fails, the upgrade will halt and prompt for action.

Stage	Description	Action Required
Prepare Backups	This is a confirmation step used to confirm that you have taken proper backups before proceeding.	You must acknowledge the prompt for database backups.
ZooKeeper	All ZooKeeper servers are upgraded and restarted.	None
Ranger	Ranger Admin and UserSync servers are upgraded and restarted.	None. If Ranger Admin does not function after the upgrade completes, run the upgrade scripts manually, then Retry Upgrading Ranger, as described here .
Core Masters	This stage upgrades the master components of core services. This includes JournalNodes & NameNodes (HDFS), HistoryServer (MapReduce2), ResourceManager & ATS (YARN) and HBase Masters (HBase).	None
All Service Checks	All Service Checks are performed against the cluster.	If any service check fails, you will be prompted to Ignore and Continue, Downgrade or Retry. If you selected the "Skip all Service Check failures" option, you will only be prompted at the completion of all of the Service Checks.
Core Slaves	This stage upgrades the slave components of core services. This includes DataNodes (HDFS), RegionServers (HBase) and NodeManagers (YARN). This is down	After the first 20% batch is complete, you will be prompted to Verify the cluster is operating correctly.

Stage	Description	Action Required
	in two batches: 20% of the slaves first, then the remaining slaves.	
All Service Checks	All Service Checks are performed against the cluster.	If any service check fails, you will be prompted to Ignore and Continue, Downgrade or Retry. If you selected the "Skip all Service Check failures" option, you will only be prompted at the completion of all of the Service Checks.
Hive	This stage upgrades the Hive Metastore, HiveServer2 and WebHCat components and the Hive clients.	Ambari will switch the HiveServer2 port from 10000 to 10010 (or, 10011 if using HTTP transport mode). You will be prompted to confirm prior to the switch.
Spark	The Spark Job History Server and clients are upgraded.	None
Oozie	The Oozie Server and clients are upgraded.	None
Falcon	The Falcon Server and clients are upgraded.	None
Client Components	All remaining clients are upgraded.	None
All Service Checks	All Service Checks are performed against the cluster.	If any service check fails, you will be prompted to Ignore and Continue, Downgrade or Retry. If you selected the "Skip all Service Check failures" option, you will only be prompted at the completion of all of the Service Checks.
Kafka	The Kafka Brokers are upgraded.	None
Knox	The Knox Gateways are upgraded.	None
Storm	Storm does not support rolling upgrade from a previous HDP version to HDP-2.5.	The rolling upgrade process will ask you to stop Storm topologies, perform the upgrade, and redeploy your topologies.
Slider	The Slider clients are upgraded.	None
Flume	The Flume agents are upgrade.	None
Finalize Upgrade Pre-Check	Checks if any hosts were not upgraded, either because the host was in Maintenance Mode, or one or more components on the host failed to upgrade (and were skipped).	Click on the list that displays "# hosts" for details on the hosts (and their components) that are not upgraded. You can Pause Upgrade , delete the hosts and return to finalize.
Finalize	The component upgrades are complete. You are presented the option to Finalize, which when selected, completes the upgrade process + saves the cluster state.	Prompted to Finalize, Finalize Later or Downgrade.

- Once the upgrade is complete, you have an option to **Finalize** the upgrade, to **Finalize Later** or to **Downgrade**. Finalizing later gives you a chance to perform more validation on the cluster. Downgrade moves the cluster version back to the previous version (basically: the reverse of the upgrade process stages). **Once finalized, you cannot downgrade back to the previous version.**



Note

If you choose to finalize later, both versions will be listed on the Stack and Versions tab with the starting version displaying as Current. It is not until you finalize that Ambari makes the target version the current version. Also, until you finalize, you will not be able to perform operational changes to the cluster (such as move components, change configurations, etc).

- Click **Finalize** and the upgrade process is complete.

4.4.2. Perform Express Upgrade

1. Ambari will check that your cluster meets the [Prerequisites](#). A dialog will be presented with the results:
 - a. If any prerequisites are not met but are required, the result will be shown with an error. You will not be allowed to proceed with the upgrade. Make the appropriate corrections and return to Perform Upgrade again.
 - b. If any prerequisites are not met but are optional, the result will be shown with a warning. You will be allowed to Proceed with the upgrade.
 - c. A list of configuration changes (if any) will be displayed.
2. Once the prerequisite checks are complete, the upgrade will start. The time it takes to perform the upgrade depends on many factors. As part of the upgrade process, each component in the cluster is restarted in a serial fashion so the stop/start time is a big contributor to the overall time.
3. The upgrade process involves a set of stages. This table lists the high-level stages and if the process requires any action by you during normal operation.



Note

If any stage fails, the upgrade will halt and prompt for action.

Stage	Description	Action Required
Prepare Upgrade	You should stop all YARN queues, all long-running applications on Slider, and deactivate & kill all running Storm topologies.	Perform the actions to prepare for the upgrade.
Stop Components for High-Level Services	This will stop all components for High-Level Services. This includes all master components except those of HDFS, HBase, ZooKeeper and Ranger.	None
Perform Backups	This is a confirmation step used to confirm that you have taken proper backups before proceeding.	You must acknowledge the prompt for database backups.
Stop Components for Core Service	Stops all components with HDFS, HBase, ZooKeeper and Ranger.	None
Update Target Stack	Updates the stack version in Ambari to the target version. There is no downgrade past this point.	None
Update Service Configs	Updates (i.e. transfers or replaces) any configurations that are necessary for the upgrade.	None
Restart Components	Restarts all core components such as ZooKeeper, Ranger, HDFS, YARN, MapReduce2 and various Clients (Tez, Pig, Sqoop).	None
All Service Checks	All Service Checks are performed against the cluster.	If any service check fails, you will be prompted to Ignore and Continue, Downgrade or Retry. If you selected the "Skip all Service Check failures" option, you will only be prompted at the completion of all of the Service Checks.

Stage	Description	Action Required
Restart Components	Restarts the remaining components such as Oozie, Falcon, Hive, Spark and others.	None
Set Version on All Hosts	Sets the HDP version on all hosts to the target HDP version.	None
Finalize Upgrade Pre-Check	Checks if any hosts were not upgraded, either because the host was in Maintenance Mode, or one or more components on the host failed to upgrade (and were skipped).	Click on the list that displays "# hosts" for details on the hosts (and their components) that are not upgraded. You can Pause Upgrade , delete the hosts and return to finalize.
Finalize Upgrade	The component upgrades are complete. You are presented the option to Finalize, which when selected, completes the upgrade process + saves the cluster state.	Prompted to Finalize or Finalize Later or Downgrade.

- Once the upgrade is complete, you have an option to **Finalize** the upgrade, to **Finalize Later** or to **Downgrade**. Finalizing later gives you a chance to perform more validation on the cluster. Downgrade moves the cluster version back to the previous version (basically: the reverse of the upgrade process stages). **Once finalized, you cannot downgrade back to the previous version.**



Note

If you choose to finalize later, both versions will be listed on the Stack and Versions tab with the starting version displaying as Current. It is not until you finalize that Ambari makes the target version the current version. Also, until you finalize, you will not be able to perform operational changes to the cluster (such as move components, change configurations, etc).

- Click **Finalize** and the upgrade process is complete.

4.5. Post-upgrade Tasks

4.5.1. Post-upgrade Tasks for Ranger with Kerberos

- If you have not already done so, you should [migrate your audit logs from DB to Solr](#).
- After successfully upgrading, use the procedure in "[How to Regenerate Keytabs](#)" to regenerate keytabs. Select the "Only regenerate keytabs for missing hosts and components" check box, confirm the keytab regeneration, then restart all required services.
- Log in to the Ranger Admin UI and select **Settings > Users/Groups**. Select the Users tab, then click **Add New User**. On the User Detail page, enter the short-name of the principal "ranger.lookup.kerberos.principal" (if that user does not already exist) in the **User Name** box. For example, if the lookup principal is "rangerlookup/_HOST@\${realm}", enter "rangelookup" as the user name. Enter the other settings shown below, then click **Save** to create the user. Add the new "rangerlookup" user in Ranger policies and services. You can use **Test Connection** on the Services pages to check the connection.

The screenshot shows the Ranger web interface for creating a user. The breadcrumb navigation is 'Users/Groups > User Create'. The 'User Detail' form includes the following fields:

- User Name *: rangerlookup
- New Password *: [masked]
- Password Confirm *: [masked]
- First Name *: rangerlookup
- Last Name: [empty]
- Email Address: [empty]
- Select Role *: User
- Group: Please select (with a plus icon)

 At the bottom of the form are 'Save' and 'Cancel' buttons.

Additional Steps for Ranger HA with Kerberos

If Ranger HA (High Availability) is set up along with Kerberos, perform the following additional steps:

1. Use SSH to connect to the KDC server host. Use the `kadmin.local` command to access the Kerberos CLI, then check the list of principals for each domain where Ranger Admin and the load-balancer are installed.

```
kadmin.local
kadmin.local: list_principals
```

For example, if Ranger Admin is installed on <host1> and <host2>, and the load-balancer is installed on <host3>, the list returned should include the following entries:

```
HTTP/ <host3>@EXAMPLE.COM
HTTP/ <host2>@EXAMPLE.COM
HTTP/ <host1>@EXAMPLE.COM
```

If the HTTP principal for any of these hosts is not listed, use the following command to add the principal:

```
kadmin.local: addprinc -randkey HTTP/<host3>@EXAMPLE.COM
```



Note

This step will need to be performed each time the Spnego keytab is regenerated.

2. Use the following `kadmin.local` commands to add the HTTP Principal of each of the Ranger Admin and load-balancer nodes to the Spnego keytab file:


```
kadmin.local: ktadd -norandkey -kt /etc/security/keytabs/spnego.service.
keytab HTTP/ <host3>@EXAMPLE.COM
kadmin.local: ktadd -norandkey -kt /etc/security/keytabs/spnego.service.
keytab HTTP/ <host2>@EXAMPLE.COM
kadmin.local: ktadd -norandkey -kt /etc/security/keytabs/spnego.service.
keytab HTTP/ <host1>@EXAMPLE.COM
```

Use the `exit` command to exit `ktadmin.local`.

3. Run the following command to check the Spnego keytab file:

```
klint -kt /etc/security/keytabs/spnego.service.keytab
```

The output should include the principals of all of the nodes on which Ranger Admin and the load-balancer are installed. For example:

```
Keytab name: FILE:/etc/security/keytabs/spnego.service.keytab
KVNO Timestamp      Principal
-----
1 07/22/16 06:27:31 HTTP/ <host3>@EXAMPLE.COM
1 07/22/16 06:27:31 HTTP/ <host3>@EXAMPLE.COM
1 07/22/16 06:27:31 HTTP/ <host3>@EXAMPLE.COM
1 07/22/16 06:27:31 HTTP/ <host3>@EXAMPLE.COM
1 07/22/16 06:27:31 HTTP/ <host3>@EXAMPLE.COM
1 07/22/16 08:37:23 HTTP/ <host2>@EXAMPLE.COM
1 07/22/16 08:37:23 HTTP/ <host2>@EXAMPLE.COM
1 07/22/16 08:37:23 HTTP/ <host2>@EXAMPLE.COM
1 07/22/16 08:37:23 HTTP/ <host2>@EXAMPLE.COM
1 07/22/16 08:37:23 HTTP/ <host2>@EXAMPLE.COM
1 07/22/16 08:37:23 HTTP/ <host2>@EXAMPLE.COM
1 07/22/16 08:37:35 HTTP/ <host1>@EXAMPLE.COM
1 07/22/16 08:37:36 HTTP/ <host1>@EXAMPLE.COM
1 07/22/16 08:37:36 HTTP/ <host1>@EXAMPLE.COM
1 07/22/16 08:37:36 HTTP/ <host1>@EXAMPLE.COM
1 07/22/16 08:37:36 HTTP/ <host1>@EXAMPLE.COM
1 07/22/16 08:37:36 HTTP/ <host1>@EXAMPLE.COM
1 07/22/16 08:37:36 HTTP/ <host1>@EXAMPLE.COM
```

4. Use `scp` to copy the Spnego keytab file to every node in the cluster on which Ranger Admin and the load-balancer are installed. Verify that the `/etc/security/keytabs/spnego.service.keytab` file is present on all Ranger Admin and load-balancer hosts.
5. On the Ambari dashboard, select **Ranger > Configs > Advanced**, then select **Advanced ranger-admin-site**. Set the value of the `ranger.spnego.kerberos.principal` property to `*`.

The screenshot shows a configuration field for the property `ranger.spnego.kerberos.principal`. The value entered in the text box is an asterisk (`*`). To the right of the text box are three small icons: a lock, a green circle, and a blue circle.

6. Click **Save** to save the configuration, then restart Ranger Admin.