Hortonworks DataFlow

Installing HDF Services on an Existing HDP Cluster fo

(December 22, 2017)

Hortonworks DataFlow: Installing HDF Services on an Existing HDP Cluster for IBM Power Systems

Copyright © 2012-2017 Hortonworks, Inc. Some rights reserved.

Table of Contents

1. Upgrading Ambari and the HDF Management Pack

Ambari and the 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 [1]

- Prepare Ambari for Upgrade [2]

- *Mandatory* Post-Upgrade Tasks [6]

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

⚠️ **Important**

Completing post-upgrade tasks is mandatory.
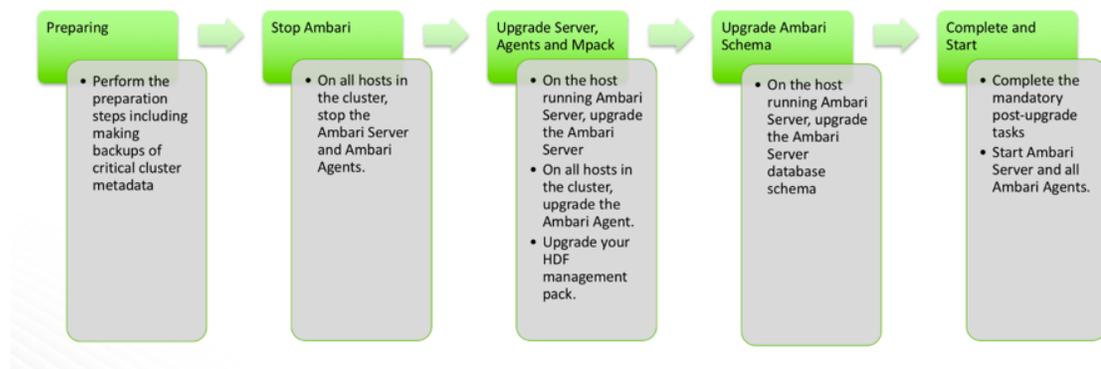
1.1. Preparing to Upgrade

- Be sure to review the Ambari 2.6.1.0 release notes 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 System and add the Grafana component.

- After upgrading Ambari, you must also upgrade SmartSense.

- Upgrade Ambari to version 2.5x or 2.6x, based on your current Ambari Server version.

The following table lists recommended (✓), and unsupported (X) upgrade paths.

From / To	2.5.x	2.6.x
2.4.0.x	✓	✓
2.4.2.x	N/A	✓
2.5.x	N/A	✓

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.

Next Steps

Prepare Ambari for Upgrade [2]

More Information

Ambari 2.6.2.0 Release Notes

1.2. Prepare Ambari for Upgrade

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
```

1.3. Get the Ambari Repository

1. 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:

   For RHEL/CENTOS/Oracle Linux 7 running on IBM Power Systems:

   ```
   wget -nv http://public-repo-1.hortonworks.com/ambari/centos7-ppc/2.x/
   updates/2.6.2.0/ambari.repo -O /etc/yum.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.

   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.

1.4. Upgrade Ambari Server

1. 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 "2.6"

   ```
   yum upgrade ambari-server
   ```

2. 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 0:2.6.1-143 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-2.6.*.jar to `/tmp` before proceeding with upgrade.

1.5. Upgrade the Ambari Agents

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

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

   ```
   yum upgrade ambari-agent
   ```

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

   For RHEL/CentOS/Oracle Linux 7:   `rpm -qa | grep ambari-agent`

1.6. Upgrade the HDF Management Pack

About This Task

A management pack bundles service definitions, stack definitions, and stack add-on service definitions so they do not need to be included with the Ambari core functionality and can be updated in between major releases. Upgrade the management pack to ensure that you have the latest versions of the available Apache components.

Before You Begin

Get the HDF Management Pack location and build number from the HDF Release Notes.

Steps

1. Back up your Ambari resources folder:

   ```
   cp -r /var/lib/ambari-server/resources /var/lib/ambari-server/resources.
   backup
   ```

2. Upgrade the HDF management pack:

```
ambari-server upgrade-mpack \
--mpack=http://public-repo-1.hortonworks.com/HDF/centos7/3.x/updates/
<version>/tars/hdf_ambari_mp/hdf-ambari-mpack-<version>-<build-number>.tar.
gz \
--verbose
```

1.7. Upgrade the Ambari Database Schema

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

```
ambari-server upgrade
```

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

```
ambari-server start
```

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

```
ambari-agent start
```

4. 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.

1.8. Restart Ambari

1. 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".

2. If you have configured Ambari to authenticate against an external LDAP or Active Directory, you **must** re-run

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

3. If you are running **Ambari Metrics** service in your cluster, you **must** upgrade Ambari Metrics System and add the Grafana component.

4. If your cluster includes the SmartSense service, you **must** upgrade SmartSense along with Ambari.

1.9. *Mandatory* Post-Upgrade Tasks

Depending on the configuration of your cluster and your current Ambari version, you must upgrade any of the following features in your cluster, as described in the following topics:

Upgrading Ambari Infra	If your cluster includes Ambari Infra service, you must upgrade it along with Ambari.
Upgrading Ambari Log Search	If your cluster includes Ambari Log Search service, you must upgrade it along with Ambari.
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.
Upgrading Configurations	Certain scenarios may require that you modify configurations that Ambari did not upgrade automatically.
Upgrading SmartSense	If your cluster includes the SmartSense service, you must upgrade it after upgrading Ambari.

1.9.1. Upgrading Ambari Infra

If you have Ambari Solr installed, you must upgrade Ambari Infra after upgrading Ambari.

Steps

1. Make sure Ambari Infra services are stopped. From **Ambari Web**, browse to **Services > Ambari Infra** and select **Stop** from the **Service Actions** menu.

2. On every host in your cluster with an Infra Solr Client installed, run the following commands:

   For RHEL/CentOS/Oracle Linux:

   ```
   yum clean all
   ```

   ```
   yum upgrade ambari-infra-solr-client
   ```

3. Execute the following command on all hosts running an Ambari Infra Solr Instance:

   For RHEL/CentOS/Oracle Linux:

   ```
   yum upgrade ambari-infra-solr
   ```

4. Start the Ambari Infra services.

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

1.9.2. Upgrading Ambari Log Search

If you have Ambari Log Search installed, you must upgrade Ambari Log Search after upgrading Ambari.

Prerequisites

Before starting this upgrade, ensure the Ambari Infra components have been upgraded.

Steps

1. Make sure Ambari Log Search service is stopped. From **Ambari Web**, browse to **Services > Log Search** and select **Stop** from the **Service Actions** menu.

2. On every host in your cluster running a Log Feeder, run the following commands:

   For RHEL/CentOS/Oracle Linux:

   ```
   yum clean all
   ```

   ```
   yum upgrade ambari-logsearch-logfeeder
   ```

3. Execute the following command on all hosts running the Log Search Server:

   For RHEL/CentOS/Oracle Linux:

   ```
   yum upgrade ambari-logsearch-portal
   ```

4. Start Log Search Service.

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

1.9.3. Upgrading Ambari Metrics

Prerequisites

Ensure all services are up and healthy.

Steps

1. Make sure Ambari Metrics service is stopped. From **Ambari Web**, browse to **Services > Ambari Metrics** and select **Stop** from the **Service Actions** menu.

2. 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
   ```

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

   For RHEL/CentOS/Oracle Linux:

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

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

   For RHEL/CentOS/Oracle Linux:

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

5. 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 and you must restart each service to pick up the latest sink implementations.

   Please wait to restart all services until after you have completed all applicable post-upgrade tasks, for example: HDFS, YARN, Kafka, HBase, Flume, Storm.

Next Steps

• Restart services, <u>only</u> after you complete all applicable, post-upgrade tasks.

> **Note**
>
> New Ambari Metrics Sinks will not be activated until all services are restarted.

1.9.4. Upgrading Configurations

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

Upgrading Kerberos krb5.conf [8]

Upgrading Log Rotation Configuration [9]

1.9.4.1. Upgrading 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 %}

Original Template Entry	Updated Template Entry
	kdc = {{kdc_host|trim()}}
	{%- endfor -%}

1.9.4.2. Upgrading Log Rotation Configuration

Ambari 2.6.x provides s a simplified log rotation configuration. These changes will be made automatically during your next stack upgrade, but are not automatically made during the Ambari upgrade. After upgrading Ambari from version 2.x to 2.6.0, if you want to utilize the simplified log rotation configuration, you must update configurations for all services in your cluster, using the following steps:

Steps

1. **ZooKeeper**

   a. In **Ambari Web**, browse to **ZooKeeper > Configs**.

   b. Scroll down to **Custom zookeeper-log4j**.

   c. In **Custom zookeeper-log4j**, click **Add Property**.

   d. In **Add Property**, type the following properties and values:

   zookeeper_log_max_backup_size=10

   zookeeper_log_number_of_backup_files=10

   For example:

   e. Click **Add**.

   f. Browse to **Advanced zookeeper-log4j**.

   g. In **Advanced zookeeper-log4j** *content section*, find and replace the following properties and values:

   Find: log4j.appender.ROLLINGFILE.MaxFileSize=<value>

Replace:
log4j.appender.ROLLINGFILE.MaxFileSize={{zookeeper_log_number_of_backup_files}}MB

Find: #log4j.appender.ROLLINGFILE.MaxBackupIndex=<value>MB

Replace:
#log4j.appender.ROLLINGFILE.MaxBackupIndex={{zookeeper_log_number_of_backup_files}}

For example:

h. In **Configs**, click **Save**.

For example:

i. Restart **ZooKeeper**, as prompted.

2. **Kafka**

a. In **Ambari Web**, browse to **Kafka > Configs**.

b. Scroll down to **Custom Kafka-log4j**.

c. In **Custom Kafka-log4j**, click **Add Property**.

d. In **Add Property**, type the following properties and values:

kafka_log_maxfilesize=256

kafka_log_maxbackupindex=20

controller_log_maxfilesize=256

controller_log_maxbackupindex=20

e. Click **Add**.

f. Browse to **Advanced kafka-log4j**.

g. In **Advanced kafka-log4j** *content section*, find and replace the following properties and values:

Find: log4j.appender.kafkaAppender=org.apache.log4j.DailyRollingFileAppender

Add: log4j.appender.kafkaAppender.MaxFileSize = {{kafka_log_maxfilesize}}MB

Add: log4j.appender.kafkaAppender.MaxBackupIndex = {{kafka_log_maxbackupindex}}MB

Find: log4j.appender.controllerAppender=org.apache.log4j.DailyRollingFileAppender

Add: log4j.appender.controllerAppender.MaxFileSize = {{controller_log_maxfilesize}}MB

Add: log4j.appender.controllerAppender.MaxBackupIndex = {{controller_log_maxbackupindex}}

h. In **Configs**, click **Save**.

i. Restart **Kafka**, as prompted.

3. **Ranger**

a. In **Ambari Web**, browse to **Ranger > Configs > Advanced**.

b. Scroll down to **Custom admin-log4j**.

c. In **Custom admin-log4j,** click **Add Property**.

d. In **Add Property**, type the following properties and values:

ranger_xa_log_maxfilesize=256

ranger_xa_log_maxbackupindex=20

e. Click **Add**.

f. Browse to **Advanced admin-log4j**.

g. In **Advanced admin-log4j** *content section*, find and replace the following properties and values:

Find: log4j.appender.xa_log_appender=org.apache.log4j.DailyRollingFileAppender

Add:
log4j.appender.xa_log_appender.MaxFileSize={{ranger_xa_log_maxfilesize}}MB

Add:
log4j.appender.xa_log_appender.MaxBackupIndex={{ranger_xa_log_maxbackupindex}}

h. Scroll down to **Custom usersync-log4j**.

i. In **Custom usersync-log4j**, click **Add Property**.

j. In **Add Property**, type the following properties and values:

ranger_usersync_log_maxfilesize=256

ranger_usersync_log_number_of_backup_files=20

k. Click **Add**.

l. Browse to **Advanced usersync-log4j**.

m. In **Advanced usersync-log4j** *content section*, find and replace the following properties and values:

Find: log4j.appender.logFile=org.apache.log4j.DailyRollingFileAppender

Add: log4j.appender.logFile.MaxFileSize = {{ranger_usersync_log_maxfilesize}}MB

Add: log4j.appender.logFile.MaxBackupIndex = {{ranger_usersync_log_number_of_backup_files}}

n. Scroll down to **Custom tagsync-log4j**.

o. In **Custom tagsync-log4j**, click **Add Property**.

p. In **Add Property**, type the following properties and values:

ranger_tagsync_log_maxfilesize=256

ranger_tagsync_log_number_of_backup_files=20

q. Click **Add**.

r. Browse to **Advanced tagsync-log4j**.

s. In **Advanced tagsync-log4j** *content section*, find and replace the following properties and values:

Find: log4j.appender.logFile=org.apache.log4j.DailyRollingFileAppender

Add: log4j.appender.logFile.MaxFileSize = {{ranger_tagsync_log_maxfilesize}}MB

Add: log4j.appender.logFile.MaxBackupIndex = {{ranger_tagsync_log_number_of_backup_files}}

t. In **Configs**, click **Save**.

u. Restart **Ranger**, as prompted.

4. **Ranger-KMS**

   a. In **Ambari Web**, browse to **Ranger-KMS > Configs > Advanced**.

   b. Scroll down to **Custom kms-log4j**.

   c. In **Custom kms-log4j**, click **Add Property**.

   d. In **Add Property**, type the following properties and values:

      ranger_kms_log_maxfilesize=256

      ranger_kms_log_maxbackupindex=20

      ranger_kms_audit_log_maxfilesize=256

      ranger_kms_audit_log_maxbackupindex=20

   e. Click **Add**.

   f. Browse to **Advanced kms-log4j** .

   g. In **Advanced kms-log4j** *content section*, find and replace the following properties and values:

      Find: log4j.appender.kms=org.apache.log4j.DailyRollingFileAppender

      Add: log4j.appender.kms.MaxFileSize = {{ranger_kms_log_maxfilesize}}MB

      Add: log4j.appender.kms.MaxBackupIndex = {{ranger_kms_log_maxbackupindex}}

      Find: log4j.appender.kms-audit=org.apache.log4j.DailyRollingFileAppender

      Add: log4j.appender.kms-audit.MaxFileSize={{ranger_kms_audit_log_maxfilesize}}MB

      Add: log4j.appender.kms-audit.MaxBackupIndex = {{ranger_kms_audit_log_maxbackupindex}}

   h. In **Configs**, click **Save**.

   i. Restart **Ranger-KMS**.

5. **Storm**

   a. In **Ambari Web**, browse to **Storm > Configs**.

   b. Scroll down to **Custom cluster-log4j property**.

   c. In **Custom cluster-log4j property**, click **Add Property**.

   d. In **Add Property**, type the following properties and values:

      storm_a1_maxfilesize=100

      storm_a1_maxbackupindex=9

   e. Click **Add**.

   f. Browse to **Advanced storm-cluster-log4j** .

   g. In **Advanced storm-cluster-log4j** *content section*, find and replace the following properties and values:

      Find: In RollingFile="A1"<SizeBasedTriggeringPolicy size="<value>MB"/>

      Replace: <SizeBasedTriggeringPolicy size="{{storm_a1_maxfilesize}}MB"/>

      Find: In RollingFile="A1"<DefaultRolloverStrategy max="<value>"/>

      Replace: <DefaultRolloverStrategy max="{{storm_a1_maxbackupindex}}"/>

   h. Scroll down to **Custom worker-log4j property**.

   i. In **Custom worker-log4j property**, click **Add Property**.

   j. In **Add Property**, type the following properties and values:

      storm_wrkr_a1_maxfilesize=100

      storm_wrkr_a1_maxbackupindex=9

      storm_wrkr_out_maxfilesize=100

      storm_wrkr_out_maxbackupindex=4

      storm_wrkr_err_maxfilesize=100

      storm_wrkr_err_maxbackupindex=4

   k. Click **Add**.

   l. Browse to **Advanced storm-worker-log4j** .

   m. In **Advanced storm-worker-log4j** *content section*, find and replace the following properties and values:

      Find: In RollingFile="A1"<SizeBasedTriggeringPolicy size="<value> MB"/>

Replace: <SizeBasedTriggeringPolicy size="{{storm_wrkr_a1_maxfilesize}} MB"/>

Find: In RollingFile="A1"<DefaultRolloverStrategy max="<value>"/>

Replace: <DefaultRolloverStrategy max="{{storm_wrkr_a1_maxbackupindex}}"/>

Find: In RollingFile="STDOUT"<SizeBasedTriggeringPolicy size="<value>" MB/>

Replace: <SizeBasedTriggeringPolicy size="{{storm_wrkr_out_maxfilesize}} MB"/>

Find: In RollingFile="STDOUT"<DefaultRolloverStrategy max="<value>"/>

Replace: <DefaultRolloverStrategy max="{{storm_wrkr_out_maxbackupindex}}"/>

Find: In RollingFile="STDERR"<SizeBasedTriggeringPolicy size="<value>" MB/>

Replace: <SizeBasedTriggeringPolicy size="{{storm_wrkr_err_maxfilesize}} MB"/>

Find: In RollingFile="STDOUT"<DefaultRolloverStrategy max="<value>"/>

Replace: <DefaultRolloverStrategy max="{{storm_wrkr_err_maxbackupindex}}"/>

n. In **Configs**, click **Save**.

o. Restart **Storm**, as prompted.

1.9.5. Upgrading SmartSense

If your cluster includes the SmartSense service, you must upgrade it after upgrading Ambari.

More Information

Upgrading SmartSense

Next Steps

Restart services.

2. Upgrading to HDP 2.6.5

If you already have HDP 2.6.0 installed, upgrading your cluster to HDP 2.6.5 means:

• Keeping the same configuration files you used for HDP 2.6.0.

• Keeping the same data and metadata in the same location you used for HDP 2.6.0

• Installing any new components (added for the first time in HDP 2.6.0) side-by-side with existing components

2.1. Before you begin

• Ensure that you know which HDP components you need to upgrade at your installation.

• Decide whether you are going to upgrade using a local repository or a remote repository.

• If you are using the Falcon service, install the Berkeley DB prior to performing an upgrade.

  See the Prerequisite to Installing or Upgrading Falcon in the Data Movement and Integration guide.

2.2. Upgrade options

• If you are upgrading your cluster manually, use the Non-Ambari Upgrade Guide.

• If you are upgrading your cluster through Ambari, use the Ambari Upgrade Guide

More information:

• Upgrading HDP

• Register and Install HDP Version

• Obtain the HDP repos

3. Installing Databases

Schema Registry, SAM, Druid, and Superset require a relational data store to store metadata. You can use either MySQL, Postgres, or Oracle for this. This chapter describes how to install either MySQL, Postgres, and Oracle and how create a databases for SAM and Registry. If you are installing on an existing HDP cluster by using Superset, you can skip the installation instructions, because MySQL was installed with Druid. In this case, configure the databases.

Note

You should install either Postgres, Oracle or MySQL; both are not necessary. It is recommended that you use MySQL.

Warning

If you are installing Postgres, you must install Postgres 9.5 or later for SAM and Schema Registry. Ambari does not install Postgres 9.5, so you must perform a manual Postgres installation.

Installing and Configuring MySQL

- Installing MySQL [17]

- Configuring SAM and Schema Registry Metadata Stores in MySQL [18]

- Configuring Druid and Superset Metadata Stores in MySQL [19]

Installing and Configuring Postgres

- Install Postgres [19]

- Configure Postgres to Allow Remote Connections [20]

- Configure SAM and Schema Registry Metadata Stores in Postgres [21]

- Configure Druid and Superset Metadata Stores in Postgres [21]

Using an Oracle Database

- Section 3.8, "Specifying an Oracle Database to Use with SAM and Schema Registry" [22]

- Section 3.9, "Switching to an Oracle Database After Installation" [22]

3.1. Installing MySQL

About This Task

You can install MySQL 5.5 or later.

Before You Begin

On the Ambari host, install the JDBC driver for MySQL, and then add it to Ambari:

```
yum install mysql-connector-java* \
sudo ambari-server setup --jdbc-db=mysql \
--jdbc-driver=/usr/share/java/mysql-connector-java.jar
```

Steps

1. Log in to the node on which you want to install the MySQL metastore to use for SAM, Schema Registry, and Druid.

2. Install MySQL and the MySQL community server, and start the MySQL service:

```
yum localinstall \
https://dev.mysql.com/get/mysql57-community-release-el7-8.noarch.rpm

yum install mysql-community-server

systemctl start mysqld.service
```

3. Obtain a randomly generated MySQL root password:

```
grep 'A temporary password is generated for root@localhost' \
/var/log/mysqld.log |tail -1
```

4. Reset the MySQL root password. Enter the following command. You are prompted for the password you obtained in the previous step. MySQL then asks you to change the password.

```
/usr/bin/mysql_secure_installation
```

3.2. Configuring SAM and Schema Registry Metadata Stores in MySQL

Steps

1. Launch the MySQL monitor:

```
mysql -u root -p
```

2. Create the database for the Registry and SAM metastore:

```
create database registry;
create database streamline;
```

3. Create Schema Registry and SAM user accounts, replacing the last string with your password:

```
CREATE USER 'registry'@'%' IDENTIFIED BY 'R12$%34qw';
CREATE USER 'streamline'@'%' IDENTIFIED BY 'R12$%34qw';
```

4. Assign privileges to the user account:

```
GRANT ALL PRIVILEGES ON registry.* TO 'registry'@'%' WITH GRANT OPTION ;
GRANT ALL PRIVILEGES ON streamline.* TO 'streamline'@'%' WITH GRANT OPTION ;
```

5. Commit the operation:

```
commit;
```

3.3. Configuring Druid and Superset Metadata Stores in MySQL

About This Task

Druid and Superset require a relational data store to store metadata. To use MySQL for this, install MySQL and create a database for the Druid metastore.

Steps

1. Launch the MySQL monitor:

   ```
   mysql -u root -p
   ```

2. Create the database for the Druid and Superset metastore:

   ```
   CREATE DATABASE druid DEFAULT CHARACTER SET utf8;
   CREATE DATABASE superset DEFAULT CHARACTER SET utf8;
   ```

3. Create `druid` and `superset` user accounts, replacing the last string with your password:

   ```
   CREATE USER 'druid'@'%' IDENTIFIED BY '9oNio)ex1ndL';
   CREATE USER 'superset'@'%' IDENTIFIED BY '9oNio)ex1ndL';
   ```

4. Assign privileges to the `druid` account:

   ```
   GRANT ALL PRIVILEGES ON *.* TO 'druid'@'%' WITH GRANT OPTION;
   GRANT ALL PRIVILEGES ON *.* TO 'superset'@'%' WITH GRANT OPTION;
   ```

5. Commit the operation:

   ```
   commit;
   ```

3.4. Install Postgres

Before You Begin

If you have already installed a MySQL database, you may skip these steps.

> **Warning**
>
> You must install Postgres 9.5 or later for SAM and Schema Registry. Ambari does not install Postgres 9.5, so you must perform a manual Postgres installation.

Steps

1. Install RPM according to the requirements of your operating system:

   ```
   yum install https://yum.postgresql.org/9.6/redhat/rhel-7-x86_64/pgdg-redhat96-9.6-3.noarch.rpm
   ```

2. Install 9.5+ Postgres database:

```
yum install postgresql96-server postgresql96-contrib postgresql96
```

3. Initialize the database:

- For CentOS 7, use the following syntax:

  ```
  /usr/pgsql-9.6/bin/postgresql96-setup initdb
  ```

- For CentOS 6, use the following syntax:

  ```
  sudo service postgresql initdb
  ```

4. Start Postgres.

   For example, if you are using CentOS 7, use the following syntax:

   ```
   systemctl enable postgresql-9.6.service
   systemctl start postgresql-9.6.service
   ```

5. Verify that you can log in:

   ```
   sudo su postgres
   psql
   ```

3.5. Configure Postgres to Allow Remote Connections

About This Task

It is critical that you configure Postgres to allow remote connections before you deploy a cluster. If you do not perform these steps in advance of installing your cluster, the installation fails.

Steps

1. Open /var/lib/pgsql/9.6/data/pg_hba.conf and update to the following

   ```
   # "local" is for Unix domain socket connections only
   local all all trust


   # IPv4 local connections:
   host all all 0.0.0.0/0 trust


   # IPv6 local connections:
   host all all ::/0 trust
   ```

2. Open /var/lib//pgsql/9.6/data/postgresql.conf and update to the following:

   ```
   listen_addresses = '*'
   ```

3. Restart Postgres:

4.
```
systemctl stop postgresql-9.6.service
systemctl start postgresql-9.6.service
```

3.6. Configure SAM and Schema Registry Metadata Stores in Postgres

About This Task

If you have already installed MySQL and configured SAM and Schema Registry metadata stores using MySQL, you do not need to configure additional metadata stores in Postgres.

Steps

1. Log in to Postgres:

```
sudo su postgres
psql
```

2. Create a database called `registry` with the password `registry`:

```
create database registry;
CREATE USER registry WITH PASSWORD 'registry';
GRANT ALL PRIVILEGES ON DATABASE "registry" to registry;
```

3. Create a database called `streamline` with the password `streamline`:

```
create database streamline;
CREATE USER streamline WITH PASSWORD 'streamline';
GRANT ALL PRIVILEGES ON DATABASE "streamline" to streamline;
```

3.7. Configure Druid and Superset Metadata Stores in Postgres

About This Task

Druid and Superset require a relational data store to store metadata. To use Postgres for this, install Postgres and create a database for the Druid metastore. If you have already created a data store using MySQL, you do not need to configure additional metadata stores in Postgres.

Steps

1. Log in to Postgres:

```
sudo su postgres
psql
```

2. Create a database, user, and password, each called `druid`, and assign database privileges to the user `druid`:

```
create database druid;
CREATE USER druid WITH PASSWORD 'druid';
GRANT ALL PRIVILEGES ON DATABASE  "druid" to druid;
```

3. Create a database, user, and password, each called `superset`, and assign database privileges to the user `superset`:

```
create database superset;
CREATE USER superset WITH PASSWORD 'superset';
GRANT ALL PRIVILEGES ON DATABASE "superset" to superset;
```

3.8. Specifying an Oracle Database to Use with SAM and Schema Registry

About This Task

You may use an Oracle database with SAM and Schema Registry. Oracle databases 12c and 11g Release 2 are supported

Prerequisites

You have an Oracle database installed and configured.

Steps

1. Register the Oracle JDBC driver jar.

```
sudo ambari-server setup --jdbc-db=oracle --jdbc-driver=/usr/share/java/
ojdbc.jar
```

2. From the SAM an Schema Registry configuration screen, select Oracle as the database type and provide the necessary Oracle Server JDBC credentials and connection string.

3.9. Switching to an Oracle Database After Installation

About This Task

If you want to use an Oracle database with SAM or Schema Registry after you have performed your initial HDF installation or upgrade, you can switch to an Oracle database. Oracle databases 12c and 11g Release 2 are supported

Prerequisites

You have an Oracle database installed and configured.

Steps

1. Log into Ambari Server and shut down SAM or Schema Registry.

2. From the configuration screen, select Oracle as the database type and provide Oracle credentials, the JDBC connection string and click **Save**.

3. From the command line where Ambari Server is running, register the Oracle JDBC driver jar:

```
sudo ambari-server setup --jdbc-db=oracle --jdbc-driver=/usr/share/java/
ojdbc.jar
```

4. From the host where SAM or Schema Registry are installed, copy the JDBC jar to the following location, depending on which component you are updating.

```
cp ojdbc6.jar /usr/hdf/current/registry/bootstrap/lib/.
cp ojdbc6.jar /usr/hdf/current/streamline/bootstrap/lib/.
```

5. From the host where SAM or Schema Registry are installed, run the following command to create the required schemas for SAM or Schema Registry.

```
export JAVA_HOME=/usr/jdk64/jdk1.8.0_112 ; source /usr/hdf/current/
streamline/conf/streamline-env.sh ; /usr/hdf/current/streamline/bootstrap/
bootstrap-storage.sh create

export JAVA_HOME=/usr/jdk64/jdk1.8.0_112 ; source /usr/hdf/current/registry/
conf/registry-env.sh ; /usr/hdf/current/registry/bootstrap/bootstrap-
storage.sh create
```

> **Note**
>
> You only this command run once, from a single host, to prepare the database.

6. Confirm that new tables are created in the Oracle database.

7. From Ambari, restart SAM or Schema Registry.

8. If you are specifying an Oracle database for SAM, run the following command after you have restarted SAM.

```
export JAVA_HOME=/usr/jdk64/jdk1.8.0_112 ; source /usr/hdf/current/
streamline/conf/streamline-env.sh ; /usr/hdf/current/streamline/bootstrap/
bootstrap.sh
```

9. Confirm that Sam or Schema Registry are available and turn off maintenance mode.

4. Installing the HDF Management Pack

About This Task

A management pack (mpack) bundles service definitions, stack definitions, and stack add-on service definitions so they do not need to be included with the Ambari core functionality and can be updated in between major releases.

Steps

1. Download the Hortonworks HDF management pack. You can find the download location for your operating system in the *HDF Release Notes*.

2. Copy the bundle to `/tmp` on the node where you installed Ambari.

3. Install the management pack:

```
ambari-server install-mpack \
--mpack=/tmp/hdf-ambari-mpack-<version>.tar.gz \
--verbose
```

4. Restart the Ambari server:

```
ambari-server restart
```

More Information

HDF Release Notes

5. Update the HDF Base URL

About This Task

Adding the base URL tells Ambari where to look for the HDF repository. This step is necessary when you are using an existing Ambari instance, already managing an HDP cluster, to install and manage an HDF cluster.

Steps

1. From the Ambari menu, click the **admin** drop-down in the top right of your Ambari Dashboard view. Then select **Manage Ambari**.

2. From the **Clusters** view on the left, click **Versions**, and then click the **HDP version** link.

3. Configure the HDF Base URL to the base URL appropriate for your operating system. Find the HDF Base URLs in the HDF Release Notes.

4. Click **Save**.

6. Add HDF Services to an HDP Cluster

About This Task

You can use the HDF management pack and Ambari to add HDF services to an HDP cluster.

⚠️ **Important**

You cannot install SAM and Schema Registry for HDF 3.1.x on an HDP 2.6.4 or 2.6.5 cluster, and you cannot upgrade these services from a previous HDP cluster.
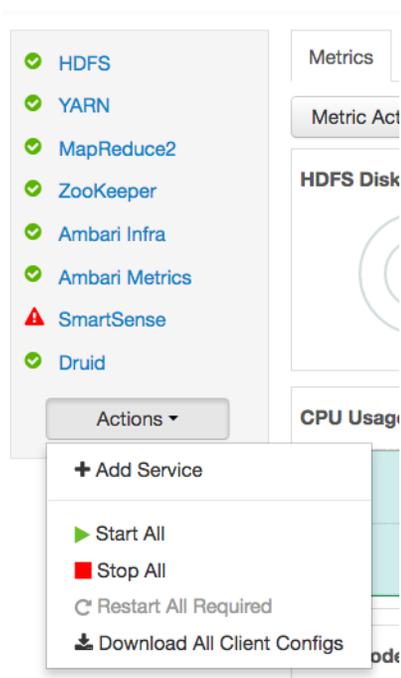
⚠️ **Important**

You cannot upgrade your HDF Storm and Kafka versions if they exist on an HDP cluster.

Steps

1. If you are installing HDF services on an existing HDP Cluster, on the Ambari home page, click the button **Actions** and select **+ Add Service**.

2. Chose select the HDF Services (NiFi and NiFi Registry) you want to install.

3. On the Assign Masters screen, distribute master services using the preceding deployment diagram of the Stream Processing cluster.

4. On the Assign Slaves and Clients screen, distribute slave services using the deployment diagram of the Stream Processing cluster.

7. Configure HDF Components

You can customize your Hortonworks DataFlow (HDF) component configurations either during or after installation. During installation, you customize HDF component configurations in the **Customize Services** page of the installation wizard. After installation, you can navigate to Services > Configs in the Ambari dashboard.

- Configure NiFi [27]

- Configure NiFi for Atlas Integration [27]

- Configure Kafka [29]

- Configure Storm [29]

- Configure Log Search [29]

- Deploy the Cluster Services [30]

- Access the UI for Deployed Services [30]

7.1. Configure NiFi

About This Task

You use the **NiFi** tab in the **Customize Services** step to configure Apache NiFi. Generally, you can accept the defaults during initial installation. However, there are some settings that you must set before proceeding.

Steps

1. From **Advanced-nifi-ambari-config**, specify the **Encrypt Configuration Master Key Passwords**.

   This password is used when you generate the master key for sensitive properties encryption in the NiFi properties file when it is written to disk. It must contain at least 12 characters.

2. From **Advanced-nifi-ambari-config**, provide the **Sensitive property values encryption password**.

   This is the password used when you encrypt any sensitive property values that are configured in processors. For enhanced security, it should contain at least 10 characters.

7.2. Configure NiFi for Atlas Integration

About This Task

You can integrate NiFi with Apache Atlas to take advantage of robust dataset and application lineage support. You do this by configuring the NiFi ReportLineageToAtlas Reporting Task once you have NiFi configured and running.

Before You Begin

If NiFi is installed on an HDP cluster, you must be running HDP 2.6.4. If NiFi is installed on an HDF cluster managed by a separate Ambari instance, you must be running HDP 2.6.1 or later, and Apache Atlas 0.8.0 or later.

Steps

1. From the Global Menu located in NiFi's upper right corner, select **Controller Services** and click the **Reporting Tasks** tab.

2. Click the **Add** (+) icon to launch the **Add Reporting Task** dialog.

3. Select ReportLineageToAtlas and click **Add**.

4. Click the **Edit** icon to launch the **Configure Reporting Task** dialog. The following Properties are required:

   - Atlas URLs – a comma-separated list of Atlas Server URLs. Once you have started reporting, you cannot modify an existing Reporting Task to add a new Atlas Server. When you need to add a new Atlas Server, you must create a new reporting task.

   - Atlas Authentication Method – Specifies how to authenticate the Reporting Task to the Atlas Server. Basic authentication is the default.

   - NiFi URL for Atlas – Specifies the NiFi cluster URL

   - NiFi Lineage Strategy – Specifies the level of granularity for your NiFi dataflow reporting to Atlas. Once you have started reporting, you should not switch between simple and complete lineage reporting strategies.

   - Provenance Record Start Position – Specifies where in the Provenance Events stream the Reporting Task should start.

   - Provenance Record Batch Size – Specifies how many records you want to send in a single batch

   - Create Atlas Configuration File – If enabled, the `atlas-application-properties` file and the `Atlas Configuration Directory` are automatically created when the Reporting Task starts.

   - Kafka Security Protocol – Specifies the protocol used to communicate with Kafka brokers to send Atlas hook notification messages. This value should match Kafka's security.protocol property value.

Result

Once you have ReportLineageToAtlas up and running, you may view dataset level lineage graphs in the Atlas UI.

> **Note**
>
> The default time interval for the Reporting Task to start sending data to an Atlas Server is 5 minutes so do not expect to see immediate lineage graphs.

You can change the default time interval in the Reporting Task property configuration.

More Information

For complete information, see the help included with the Reporting Task.

7.3. Configure Kafka

About This Task

You can configure Apache Kafka from the **Kafka** tab in the **Customize Services** step.

Steps

1. For your initial installation, accept the default values set by Apache Ambari.

2. If Ambari prompts you with Some configurations need your attention before you can proceed, review the list of properties and provide the required information.

3. Review the *Apache Kafka Component Guide* for information about configuring Apache Storm to meet your operational objectives.

More Information

Configuring Kafka for Production Environments

7.4. Configure Storm

About This Task

You can configure Storm from the **Storm** tab in the **Customize Services** step.

Steps

1. For your initial installation, accept the default values set by Ambari.

2. If Ambari prompts you with:

   Some configurations need your attention before you can proceed.

   Review the list of properties and provide the required information.

3. Review the *Apache Storm Component Guide* for information about configuring storm to meet your operational objectives.

More Information

Configuring Storm for Production Environments

7.5. Configure Log Search

About This Task

To ensure that you can view logs in the new SAM Log Search, you can manually review and adjust Log Search Settings for storm_worker and storm_worker_event.

Steps

1. From the left-hand navigation pane, select **Log Search | Configs**.

2. Manually set the Log Feeder Log Levels Filter for storm_worker and storm_worker_event to include **Info**, **Debug**, and **Trace**.

7.6. Deploy the Cluster Services

After you finish the wizard and deploy the cluster, some services might fail to start. If this is the case, you can start those services individually by launching them from the Ambari dashboard Services pane.

Steps

1. From Ambari's left-hand **Services** pane, click the service you want.

2. From the **Quick Links** drop-down, select the UI option.

3. Find links for the SAM UI under **Streaming Analytics Manager** and for the Schema Registry UI under **Registry.**

Result

The UI for your HDF service opens in a new window.

7.7. Access the UI for Deployed Services

About This Task

Once you have deployed your Ambari-managed cluster, you can launch the UI for any of the services from Ambari.

Steps

1. From Ambari's left-hand **Services** pane, click the service you want.

2. From the **Quick Links** drop-down, select the UI option.

3. Find links for the SAM UI under **Streaming Analytics Manager** and for the Schema Registry UI under **Registry.**

Result

The UI for your HDF service opens in a new window.

8. Configuring Schema Registry and SAM for High Availability

About This Task

You can configure Schema Registry and SAM for high availability.

Steps for Configuring SAM for HA

1. Install two or more instances of SAM on unique nodes.

2. From the **Services** pane, select **Streaming Analytics Manager** and click the **Configs** tab.

3. In the **Jar Storage Type** drop down, select **HDFS** or **Database**.

> **Note**
>
> If you are using a MySQL database, ensure that you make adjustments to the database configuration as well. `max_allowed_packet` must be greater than the maximum file size of any custom processor or user defined function that will be uploaded.

Steps for Configuring Schema Registry for HA

1. Install two or more instances of Schema Registry on unique nodes.

2. From the **Services** pane, select **Schema Registry** and click the **Configs** tab.

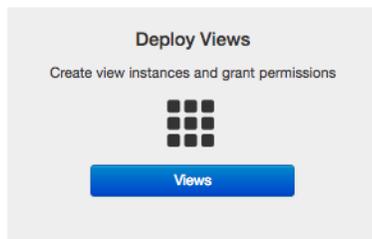3. In the **Jar Storage Type** drop down, select **HDFS**.

9. Install the Storm Ambari View

About This Task

The Storm Ambari view provides you a number of different troubleshooting and debugging tools.

Steps

1. From the **admin** drop-down, select **Manage Ambari**.

2. Click the **Views** button.

3. From the list of available Views, expand **Storm_Monitoring** and click **+ Create Instance**.

4. Configure the Storm Ambari View.

a. **Instance Name** and **Display Name** may not have an spaces.

b. The **Storm Hostname** refers to the host where the Storm UI Server is deployed.

c. The Storm port is the Storm UI port server (keep it as default 8744 if you have not changed it).

d. Click **Save**.

Result

After saving it, you should see a menu item for the Storm Ambari View.

10. Using a Local Repository

Local repositories are frequently used in enterprise clusters that have limited outbound internet access. In these scenarios, having packages available locally provides more governance, and better installation performance. These repositories are used heavily during installation for package distribution, as well as post-install for routine cluster operations such as service start/restart operations. The following sections describe the steps necessary to set up and use a local repository:

* Obtaining the Repositories [34]

* Set up a local repository having:

    * Setting Up a Local Repository with No Internet Access [36]

    * Setting up a Local Repository With Temporary Internet Access [38]

* Preparing The Ambari Repository Configuration File [39]

⚠️ **Important**

Setting up and using a local repository is **optional**. After obtaining

10.1. Obtaining the Repositories

This section describes how to obtain:

* Ambari Repositories [34]

* HDP Stack Repositories [34]

10.1.1. Ambari Repositories

If you do not have Internet access, use the link appropriate for your OS family to **download a tarball** that contains the software for setting up Ambari.

If you have temporary Internet access, use the link appropriate for your OS family to **download a repository file** that contains the software for setting up Ambari.

Ambari 2.6.2 Repositories

OS	Format	URL
RedHat 7	Base URL	http://public-repo-1.hortonworks.com/ambari/centos7-ppc/2.x/updates/2.6.2.0
	Repo File	https://public-repo-1.hortonworks.com/ambari/centos7-ppc/2.x/updates/2.6.2.0/ambari.repo
	Tarball md5 | asc	http://public-repo-1.hortonworks.com/ambari/centos7-ppc/2.x/updates/2.6.2.0/ambari-2.6.2.0-centos7-ppc.tar.gz

10.1.2. HDP Stack Repositories

If you do not have Internet access, use the link appropriate for your OS family to download a tarball that contains the software for setting up the Stack.

If you have temporary Internet access, use the link appropriate for your OS family to download a repository file that contains the software for setting up the Stack.

- HDP 2.6 Repositories [35]

10.1.2.1. HDP 2.6 Repositories

OS	Version Number	Repository Name	Format	URL
RedHat 7	HDP-2.6.5.0	HDP	Version Definition File (VDF)	http://public-repo-1.hortonworks.com/HDP/centos7-ppc/2.x/updates/2.6.5.0-292.xml
			Base URL	http://public-repo-1.hortonworks.com/HDP/centos7-ppc/2.x/updates/2.6.5.0
			Repo File	http://public-repo-1.hortonworks.com/HDP/centos7-ppc/2.x/updates/2.6.5.0/hdp.repo
			Tarball md5 | asc	http://public-repo-1.hortonworks.com/HDP/centos7-ppc/2.x/updates/2.6.5.0/HDP-2.6.5.0-centos7-ppc-rpm.tar.gz
		HDP-UTILS	Base URL	http://public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.22/repos/centos7-ppc
			Repo File	http://public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.22/repos/centos7-ppc/hdp-utils.repo
			Tarball md5 | asc	http://public-repo-1.hortonworks.com/HDP-UTILS-1.1.0.22/repos/centos7-ppc/HDP-UTILS-1.1.0.22-centos7-ppc.tar.gz
		HDP-GPL	URL	http://public-repo-1.hortonworks.com/HDP-GPL/centos7-ppc/2.x/updates/2.6.5.0/hdp.gpl.repo
			Tarball	http://public-repo-1.hortonworks.com/HDP-GPL/centos7-ppc/2.x/updates/2.6.5.0/HDP-GPL-2.6.5.0-centos7-ppc-gpl.tar.gz

10.1.3. Setting Up a Local Repository

Based on your Internet access, choose one of the following options:

- No Internet Access

  This option involves downloading the repository tarball, moving the tarball to the selected mirror server in your cluster, and extracting to create the repository.

- Temporary Internet Access

  This option involves using your temporary Internet access to sync (using reposync) the software packages to your selected mirror server and creating the repository.

  Both options proceed in a similar, straightforward way. Setting up for each option presents some key differences, as described in the following sections:

  - Getting Started Setting Up a Local Repository [35]

  - Setting Up a Local Repository with No Internet Access [36]

  - Setting up a Local Repository With Temporary Internet Access [38]

10.1.4. Getting Started Setting Up a Local Repository

To get started setting up your local repository, complete the following prerequisites:

- Select an existing server in, or accessible to the cluster, that runs a supported operating system.

- Enable network access from all hosts in your cluster to the mirror server.

- Ensure the mirror server has a package manager installed such as yum (RHEL7).

- **Optional:** If your repository has temporary Internet access, and you are using RHEL/CentOS/Oracle Linux as your OS, install yum utilities:

```
yum install yum-utils createrepo
```

1. Create an HTTP server.

   a. On the mirror server, install an HTTP server (such as Apache httpd) using the Apache instructions.

   b. Activate this web server.

   c. Ensure that any firewall settings allow inbound HTTP access from your cluster nodes to your mirror server.

   > **Note**
   >
   > If you are using Amazon EC2, make sure that SELinux is disabled.

2. On your mirror server, create a directory for your web server.

   - For example, from a shell window, type:

   ```
   mkdir –p /var/www/html/
   ```

   - If you are using a symlink, enable the

   ```
   followsymlinks
   ```

   on your web server.

Next Steps

After you have completed the steps in this section, move on to specific set up for your repository internet access type.

More Information

httpd://apache.org/download.cgi

10.1.4.1. Setting Up a Local Repository with No Internet Access

Prerequisites

Complete the Getting Started Setting up a Local Repository procedure.

--

To finish setting up your repository, complete the following:

Steps

1. Obtain the tarball for the repository you would like to create.

2. Copy the repository tarballs to the web server directory and untar.

   a. Browse to the web server directory you created.

   ```
   cd /var/www/html/
   ```

   b. Untar the repository tarballs to the following locations: where <web.server>, <web.server.directory>, <OS>, <version>, and <latest.version> represent the name, home directory, operating system type, version, and most recent release version, respectively.

   Untar Locations for a Local Repository - No Internet Access

Repository Content	Repository Location
Ambari Repository	Untar under <web.server.directory>
HDP Stack Repositories	Create directory and untar under <web.server.directory>/hdp

3. Confirm you can browse to the newly created local repositories.

   URLs for a Local Repository - No Internet Access

Repository	Base URL
Ambari Base URL	http://<web.server>/Ambari-2.6.1.0/centos7
HDP Base URL	http://<web.server>/hdp/HDP/centos7/2.x/updates/<latest.version>
HDP-UTILS Base URL	http://<web.server>/hdp/HDP-UTILS-<version>/repos/centos7

   where <web.server> = FQDN of the web server host.

   ⚠️ **Important**

   Be sure to record these Base URLs. You will need them when installing Ambari and the cluster.

4. Optional: If you have multiple repositories configured in your environment, deploy the following plug-in on all the nodes in your cluster.

   • Install the plug-in.

     • ```yum install yum-plugin-priorities```

   • Edit the `/etc/yum/pluginconf.d/priorities.conf` file to add the following:

   ```
   [main]
   ```

   ```
   enabled=1
   ```

   ```
   gpgcheck=0
   ```

More Information

Obtaining the Repositories [34]

10.1.4.2. Setting up a Local Repository With Temporary Internet Access

Prerequisites

Complete the Getting Started Setting up a Local Repository procedure.

--

To finish setting up your repository, complete the following:

Steps

1. Put the repository configuration files for Ambari and the Stack in place on the host.

2. Confirm availability of the repositories.

   ```
   yum repolist
   ```

3. Synchronize the repository contents to your mirror server.

   • Browse to the web server directory:

   • For Ambari, create `ambari` directory and reposync.

     ```
     mkdir -p ambari/centos7

     cd ambari/centos7

     reposync -a ppc64le -r ambari-2.6.2.0
     ```

   • For HDP Stack Repositories, create `hdp` directory and reposync.

     ```
     mkdir -p hdp/centos7

     cd hdp/centos7

     reposync -a ppc64le -r HDP-2.6.5.0

     reposync -a ppc64le -r HDP-UTILS-1.1.0.22
     ```

4. Generate the repository metadata.

   • For Ambari:

     ```
     createrepo <web.server.directory>/ambari/centos7/Updates-Ambari-2.6.2.0
     ```

   • For HDP Stack Repositories:

     ```
     createrepo <web.server.directory>/hdp/centos7/HDP-<latest.version>

     createrepo <web.server.directory>/hdp/centos7/HDP-UTILS-<version>
     ```

5. Confirm that you can browse to the newly created repository.

   URLs for the New Repository

Repository	Base URL
Ambari Base URL	http://<web.server>/ambari/centos7/Updates-Ambari-2.6.2.0

Repository	Base URL
HDP Base URL	http://<web.server>/hdp/centos7/HDP-<latest.version>
HDP-UTILS Base URL	http://<web.server>/hdp/centos7/HDP-UTILS-<version>

where <web.server> = FQDN of the web server host.

> ⚠️ **Important**
>
> Be sure to record these Base URLs. You will need them when installing
> Ambari and the Cluster.

6. Optional. If you have multiple repositories configured in your environment, deploy the
   following plug-in on all the nodes in your cluster.

   • Install the plug-in.

     • `yum install yum-plugin-priorities`

   • Edit the `/etc/yum/pluginconf.d/priorities.conf` file to add the following:

```
[main]
```
```
enabled=1
```
```
gpgcheck=0
```

More Information

Obtaining the Repositories [34]

10.1.5. Preparing The Ambari Repository Configuration File

1. Download the `ambari.repo` file from the public repository.

```
http://public-repo-1.hortonworks.com/ambari/centos7-ppc/2.x/updates/2.6.2.0/
ambari.repo
```

2. In the ambari.repo file, replace the base URL `baseurl` with the local repository URL.

> 📝 **Note**
>
> You can disable the GPG check by setting gpgcheck =0. Alternatively, you
> can keep the check enabled but replace the gpgkey with the URL to the
> GPG-KEY in your local repository.

```
[Updates-Ambari-2.6.2.0]
```
```
name=Ambari-2.6.2.0-Updates
```
```
baseurl=INSERT-BASE-URL
```
```
gpgcheck=1
```
```
gpgkey=http://public-repo-1.hortonworks.com/ambari/centos7-ppc/RPM-GPG-KEY/
RPM-GPG-KEY-Jenkins
```
```
enabled=1
```

```
priority=1
```

Base URL for a Local Repository

Local Repository	Base URL
Built with Repository Tarball (No Internet Access)	http://<web.server>/Ambari-2.6.2.0/centos7
Built with Repository File (Temporary Internet Access)	http://<web.server>/ambari/centos7/Updates-Ambari-2.6.2.0

where <web.server> = FQDN of the web server host.

3. Place the ambari.repo file on the machine you plan to use for the Ambari Server.

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

- Edit the `/etc/yum/pluginconf.d/priorities.conf` file to add the following:

  ```
  [main]
  ```

  ```
  enabled=1
  ```

  ```
  gpgcheck=0
  ```

Next Steps

Proceed to *Installing Ambari Server* to install and setup Ambari Server.

More Information

Setting Up a Local Repository with No Internet Access [36]

Setting up a Local Repository With Temporary Internet Access [38]

11. Navigating the HDF Library

To navigate the Hortonworks DataFlow (HDF) documentation library, begin by deciding your current goal.

If you want to...	See this document...
Install or upgrade an HDF cluster using Apache Ambari	• Release Notes • Support Matrix • Planning Your Deployment • Ambari Upgrade • MiNiFi Java Agent Quick Start
Get started with HDF	• Getting Started with Apache NiFi • Getting Started with Stream Analytics
Use and administer HDF Flow Management capabilities	• Apache NiFi User Guide • Apache NiFi Administration Guide • Apache NiFi Developer Guide • Apache NiFi Expression Language Guide • MiNiFi Java Agent Administration Guide
Use and administer HDF Stream Analytics capabilities	• Streaming Analytics Manager User Guide • Schema Registry User Guide • Apache Storm Component Guide • Apache Kafka Component Guide