



HDP 2.4

Release Notes

March 2016

Md5 **VMware** Virtual Appliance - e5b59ead89e2314ed5cbb0ad966422d9

Md5 **Virtualbox** Virtual Appliance- 6991bc1276053c27f82688fd42c5352d

HDP Stack and Ambari

The Sandbox uses the following versions of Ambari and HDP stack. Please use the following release note links provided to view Ambari and HDP stack specific information.

HDP 2.4 Build 169 Release Notes

http://docs.hortonworks.com/HDPDocuments/HDP2/HDP-2.4.0/bk_HDP_RelNotes/content/ch_relnotes_v240.html

Ambari 2.2.1.0 Build 161 Release Notes

http://docs.hortonworks.com/HDPDocuments/Ambari-2.2.1.0/bk_releasenotes_ambari_2.2.1.0/content/ch_relnotes-ambari-2.2.1.0.html

Behavior Changes

RMP-5904- Added maria_dev as default user login

RMP-5426 – Admin for Ambari no longer uses default password

Workaround: Resetting Ambari Password via SSH

1. Start your sandbox, ssh using a terminal or the web browser link:
2. Run the following commands:

```
ambari-admin-password-reset
```

The following text will appear

Please set the password for admin:

Please retype the password for admin:

After setting and retyping your new password, type the command:

```
ambari-agent restart
```

3. Ambari Admin password should be reset
4. Open Ambari login page. Verify your new password allows you to login as admin user.

RMP-5326 – Turn off alerts in ambari

BUG-52540 – Removed the line from console “You can access ssh [root@127.0.0.2 ...](#)”

New Features – Serviceability

RMP-5079 – Command to output the current blueprint being used (sandbox-blueprint)

RMP-5895- Add java version output to sandbox-version command.

Known Issues

BUG-54706 - HDFS replication set to 3

Workaround : Set configuration property “Block replication “ HDFS via Ambari to 1

Fixed Issues

BUG-54859 - No Tez view visible for Maria_dev user

BUG-48978 – Regression of Bug-40732 blueprint values are missing for Falcon

BUG-48765- Receiving message " [Dr. Who] is not authorized to view the log" when viewing logs in the sandbox

Limitations

This is a list of common limitations along with their workarounds.

RMP-4362 - Solr is not administered by Ambari

Workaround: You can use the following commands to start and stop solr

```
start service as root user /opt/lucidworks-hdpsearch/solr/bin/solr
start -c -z sandbox.hortonworks.com
stop service as root user /opt/lucidworks-hdpsearch/solr/bin/solr stop
-c -z sandbox.hortonworks.com
```

RMP-3586 - Due to dependency of the underlying OS and Virtual machine application, the following may occur when suspending the virtual machine:

-Region Server service for HBase may be stopped when returning back from suspended state. It will need to be restarted.

-Ambari Metrics may be stopped when returning back from suspended state since it now uses an embedded HBase.

Workaround: Avoid having to suspend your virtual machine.

System Information

Operating System and Java versions that the Sandbox has installed.

OS Version

CentOS release 6.7 (Final)

LSB_VERSION=base-4.0-amd64:base-4.0-noarch:core-4.0-amd64:core-4.0-noarch:graphics-4.0-amd64:graphics-4.0-noarch:printing-4.0-amd64:printing-4.0-noarch

*Java Version

java version "1.7.0_95"

OpenJDK Runtime Environment (rhel-2.6.4.0.el6_7-x86_64 u95-b00)

OpenJDK 64-Bit Server VM (build 24.95-b01, mixed mode)

**Updated from previous version*

Image File Sizes

VMware – 9 GB

Virtualbox – 9 GB

Tech Preview Packages

These packages are included to allow an early look as to what we are working on.

- Zeppelin 0.6.0

Users / Password

Hive

Database User: hive

Database Pass: hive

Databases Used

These are a list of databases used within Sandbox along with the corresponding HDP components that use them.

- Ambari: postgres
- Hive Metastore : Mysql
- Ranger: Mysql
- Oozie: derby (embedded)

HDP Supported Components Not Installed

These components are offered by the Hortonworks distribution, but not included in the Sandbox.

- Apache Accumulo
- Apache Mahout

Newly Added HDP Supported Packages

These are packages that have recently been included into the Sandbox for this release.

Apache Ambari

ambari-metrics-monitor-2.2.1.0-161.x86_64

ambari-server-2.2.1.0-161.x86_64

ambari-metrics-collector-2.2.1.0-161.x86_64

ambari-metrics-hadoop-sink-2.2.1.0-161.x86_64

ambari-agent-2.2.1.0-161.x86_64

Apache Ambari Views

ambari-admin-2.2.1.0.161.jar

capacity-scheduler-2.2.1.0.161.jar

files-2.2.1.0.161.jar
hive-2.2.1.0.161.jar
pig-2.2.1.0.161.jar
slider-2.2.1.0.161.jar
tez-view-2.2.1.0.161.jar

Apache Hadoop (HDFS, YARN, Mapreduce)

hadoop_2_4_0_0_169-2.7.1.2.4.0.0-169.el6.x86_64
hadoop_2_4_0_0_169-hdfs-2.7.1.2.4.0.0-169.el6.x86_64
hadoop_2_4_0_0_169-libhdfs-2.7.1.2.4.0.0-169.el6.x86_64
hadoop_2_4_0_0_169-yarn-2.7.1.2.4.0.0-169.el6.x86_64
hadoop_2_4_0_0_169-client-2.7.1.2.4.0.0-169.el6.x86_64
hadoop_2_4_0_0_169-mapreduce-2.7.1.2.4.0.0-169.el6.x86_64

Apache Falcon

falcon_2_4_0_0_169-0.6.1.2.4.0.0-169.el6.noarch
falcon_2_4_0_0_169-doc-0.6.1.2.4.0.0-169.el6.noarch

Apache Hive

hive_2_4_0_0_169-webhcat-server-1.2.1000.2.4.0.0-169.el6.noarch
hive_2_4_0_0_169-server-1.2.1000.2.4.0.0-169.el6.noarch
hive_2_4_0_0_169-1.2.1000.2.4.0.0-169.el6.noarch
hive_2_4_0_0_169-webhcat-1.2.1000.2.4.0.0-169.el6.noarch
hive_2_4_0_0_169-hcatalog-server-1.2.1000.2.4.0.0-169.el6.noarch
hive_2_4_0_0_169-metastore-1.2.1000.2.4.0.0-169.el6.noarch
hive_2_4_0_0_169-hcatalog-1.2.1000.2.4.0.0-169.el6.noarch
hive_2_4_0_0_169-server2-1.2.1000.2.4.0.0-169.el6.noarch
hive_2_4_0_0_169-jdbc-1.2.1000.2.4.0.0-169.el6.noarch

Apache Hbase

hbase_2_4_0_0_169-rest-1.1.2.2.4.0.0-169.el6.noarch
hbase_2_4_0_0_169-doc-1.1.2.2.4.0.0-169.el6.noarch
hbase_2_4_0_0_169-regionserver-1.1.2.2.4.0.0-169.el6.noarch
hbase_2_4_0_0_169-thrift-1.1.2.2.4.0.0-169.el6.noarch
hbase_2_4_0_0_169-thrift2-1.1.2.2.4.0.0-169.el6.noarch
hbase_2_4_0_0_169-1.1.2.2.4.0.0-169.el6.noarch
hbase_2_4_0_0_169-master-1.1.2.2.4.0.0-169.el6.noarch

Apache Flume

flume_2_4_0_0_169-agent-1.5.2.2.4.0.0-169.el6.noarch
flume_2_4_0_0_169-1.5.2.2.4.0.0-169.el6.noarch

Apache Kafka

kafka_2_4_0_0_169-0.9.0.2.4.0.0-169.el6.noarch

Apache Knox

knox_2_4_0_0_169-0.6.0.2.4.0.0-169.el6.noarch

Apache Oozie

oozie_2_4_0_0_169-client-4.2.0.2.4.0.0-169.el6.noarch

oozie_2_4_0_0_169-4.2.0.2.4.0.0-169.el6.noarch

Apache Phoenix

phoenix_2_4_0_0_169-4.4.0.2.4.0.0-169.el6.noarch

Apache Pig

pig_2_4_0_0_169-0.15.0.2.4.0.0-169.el6.noarch

Apache Ranger

ranger_2_4_0_0_169-hive-plugin-0.5.0.2.4.0.0-169.el6.x86_64

ranger_2_4_0_0_169-yarn-plugin-0.5.0.2.4.0.0-169.el6.x86_64

ranger_2_4_0_0_169-kafka-plugin-0.5.0.2.4.0.0-169.el6.x86_64

ranger_2_4_0_0_169-knox-plugin-0.5.0.2.4.0.0-169.el6.x86_64

ranger_2_4_0_0_169-usersync-0.5.0.2.4.0.0-169.el6.x86_64

ranger_2_4_0_0_169-hdfs-plugin-0.5.0.2.4.0.0-169.el6.x86_64

ranger_2_4_0_0_169-admin-0.5.0.2.4.0.0-169.el6.x86_64

ranger_2_4_0_0_169-storm-plugin-0.5.0.2.4.0.0-169.el6.x86_64

ranger_2_4_0_0_169-hbase-plugin-0.5.0.2.4.0.0-169.el6.x86_64

ranger_2_4_0_0_169-debuginfo-0.5.0.2.4.0.0-169.el6.x86_64

ranger_2_4_0_0_169-kms-0.5.0.2.4.0.0-169.el6.x86_64

ranger_2_4_0_0_169-solr-plugin-0.5.0.2.4.0.0-169.el6.x86_64

Apache Solr (Included in the Hadoop Search package)

Solr 5.21 lucidworks-hdpsearch-2.3-4.noarch

Apache Slider

slider_2_4_0_0_169-0.80.0.2.4.0.0-169.el6.noarch

storm_2_4_0_0_169-slider-client-0.10.0.2.4.0.0-169.el6.x86_64

Apache Spark

spark_2_4_0_0_169-1.6.0.2.4.0.0-169.el6.noarch

spark_2_4_0_0_169-yarn-shuffle-1.6.0.2.4.0.0-169.el6.noarch

spark_2_4_0_0_169-python-1.6.0.2.4.0.0-169.el6.noarch

Apache Sqoop

sqoop_2_4_0_0_169-metastore-1.4.6.2.4.0.0-169.el6.noarch

sqoop_2_4_0_0_169-1.4.6.2.4.0.0-169.el6.noarch

Apache Storm

storm_2_4_0_0_169-slider-client-0.10.0.2.4.0.0-169.el6.x86_64
storm_2_4_0_0_169-0.10.0.2.4.0.0-169.el6.x86_64

Apache Tez

tez_2_4_0_0_169-0.7.0.2.4.0.0-169.el6.noarch

Apache Zookeeper

zookeeper_2_4_0_0_169-server-3.4.6.2.4.0.0-169.el6.noarch
zookeeper_2_4_0_0_169-3.4.6.2.4.0.0-169.el6.noarch

Other Packages

These are some of the installed packages in the Sandbox that the HDP components may depend on.

Python

python-lxml-2.2.3-1.1.el6.x86_64
python-pycurl-7.19.0-8.el6.x86_64
python-iniparse-0.3.1-2.1.el6.noarch
python-pip-7.1.0-1.el6.noarch
python-libs-2.6.6-64.el6.x86_64
python-urlgrabber-3.9.1-9.el6.noarch
rpm-python-4.8.0-47.el6.x86_64
newt-python-0.52.11-3.el6.x86_64
python-devel-2.6.6-64.el6.x86_64
python-argparse-1.2.1-2.1.el6.noarch
python-setuptools-0.6.10-3.el6.noarch
python-2.6.6-64.el6.x86_64

*difference from previous version there is now python-urlgrabber

mysql

mysql-5.1.73-5.el6_6.x86_64
mysql-libs-5.1.73-5.el6_6.x86_64
mysql-connector-java-5.1.17-6.el6.noarch
mysql-server-5.1.73-5.el6_6.x86_64

Postgres

postgresql-server-8.4.20-4.el6_7.x86_64
postgresql-8.4.20-4.el6_7.x86_64
postgresql-libs-8.4.20-4.el6_7.x86_64

HDP Services Started Automatically on Startup

When the virtual machine is booted up, the following services are started. If not specified, assume all are java processes. The users that launch the process are the corresponding names of the component. The processes are listed with their main class.

Ambari

*AmbariServer - org.apache.ambari.server.controller.AmbariServer

Ambari Agent (non java process)

**Run as root user*

Atlas

Main - org.apache.atlas.Main

HDFS

*Portmap - org.apache.hadoop.portmap.Portmap

NameNode - org.apache.hadoop.hdfs.server.namenode.NameNode

DataNode - org.apache.hadoop.hdfs.server.datanode.DataNode

**Nfs

**Unlike the other processes that are launched by hdfs user, these are run as root user.*

***The nfs process doesn't show up as a name for jps output*

HIVE

*RunJar - webhcat - org.apache.hadoop.util.RunJar

RunJar - metastore - org.apache.hadoop.util.RunJar

RunJar - hiveserver2 - org.apache.hadoop.util.RunJar

**Run as hcat user*

Mapreduce

*JobHistoryServer - org.apache.hadoop.mapreduce.v2.hs.JobHistoryServer

**mapred is the user used to launch this process*

Oozie

Bootstrap - org.apache.catalina.startup.Bootstrap

Ranger (XaSecure)

*UnixAuthenticationService- com.xasecure.authentication.UnixAuthenticationService

EmbeddedServer- com.xasecure.server.tomcat.EmbeddedServer

**Run as root user*

Spark

HistoryServer - org.apache.spark.deploy.history.HistoryServer

YARN

ApplicationHistoryServer -

org.apache.hadoop.yarn.server.applicationhistoryservice.ApplicationHistoryServer

ResourceManager -

org.apache.hadoop.yarn.server.resourcemanager.ResourceManager

NodeManager - org.apache.hadoop.yarn.server.nodemanager.NodeManager

Zookeeper

QuorumPeerMain - org.apache.zookeeper.server.quorum.QuorumPeerMain

Zeppelin

ZeppelinServer - org.apache.zeppelin.server.ZeppelinServer

HDP Services not started automatically

In order to utilize the functionality of these services, they need to be turned on and require more memory to be added to the virtual machine.

HDFS

SecondaryNameNode -

org.apache.hadoop.hdfs.server.namenode.SecondaryNameNode

**Since on a single node, secondary namenode is not needed, it is not started.*

Falcon

Main - org.apache.falcon.Main

Flume

Application - org.apache.flume.node.Application

HBase

HRegionServer - org.apache.hadoop.hbase.regionserver.HRegionServer

HMaster - org.apache.hadoop.hbase.master.HMaster

Kafka

Kafka - kafka.Kafka

Knox

gateway.jar - /usr/hdp/current/knox-server/bin/gateway.jar

*ldap.jar - /usr/hdp/current/knox-server/bin/ldap.jar

**This process is a mini ldap server*

Storm

supervisor - backtype.storm.daemon.supervisor

nimbus - backtype.storm.daemon.nimbus

logviewer - backtype.storm.daemon.logviewer

core - backtype.storm.ui.core
drpc - backtype.storm.daemon.drpc

As Designed

The following are Ambari settings that one should be aware of.

- By default Ambari Metrics service is off.
- Maintenance Mode is on for
 - HDFS
 - HBase
 - Falcon
 - Storm
 - Ambari Metrics
 - Kafka
 - Knox