



HDP 2.3

Release Notes

July 2015

Md5 **VMware** Virtual Appliance - 3de9d6fc535c786db882c06bba63d2e3

Md5 **Virtualbox** Virtual Appliance- f36568f92127736d85f8d85037e910d1

Md5 **HyperV** Virtual Appliance- 9b5ecdb90a32f62d03c77df29a594fb7

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 2.3 Build 2557 Release Notes

http://docs.hortonworks.com/HDPDocuments/HDP2/HDP-2.3.0/bk_HDP_RelNotes/content/ch_relnotes_v230.html

Ambari 2.1 Build 1470 Release Notes

http://docs.hortonworks.com/HDPDocuments/Ambari-2.1.0.0/bk_releasenotes_ambari_2.1.0.0/content/ch_relnotes-ambari-2.1.0.0.html

New Features - Improvements

RMP-3346 - Sandbox splash screen updated for ease of use

RMP-3996 – Shellinabox added for ease of use accessing console via web vs. ssh client

RMP-4292 - NfsGateway is now controlled by Ambari

RMP-3504 – Solr is packaged with HDP-Search

RMP-3652 - Turn off THP setting in the OS

RMP-4136 – Local files view added in Views

New Features – Serviceability

RMP-3179 – Sandbox version output (execute `sandbox-version` in the console)

Known Issues

BUG-41404 – solr displays with error with Banana “Error Collection not found at /solr/collection1”

BUG-41405 - Vmware startup for DAL has failed for portreserve

Fixed Issues

BUG-41026 – Disable Ipv6

BUG-40220 - Running pig on tez job on Sandbox generates the following error
`java.lang.OutOfMemoryError: Java heap space"`

BUG-37937 - Warnings showing up on the startup console virtualbox

BUG-38264 - Upgrade mysql-connector-java.jar to 5.1.29+

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-3553 - Ranger is not administered by Ambari

Workaround: You can use the following commands to start and stop Ranger and Knox services by executing the commands within the console or ssh.

Ranger

```
start Apache Ranger Admin: service ranger-admin start
stop Apache Ranger Admin: service ranger-admin stop
start Ranger UserSync: service ranger-usersync start
stop Ranger UserSync: service ranger-usersync start
```

Ldap Service (For Knox)

For Knox you need to start the Ldap service.

```
start: su - knox -c '/usr/hdp/current/knox-server/bin/ldap.sh start'
stop: su - knox -c '/usr/hdp/current/knox-server/bin/ldap.sh stop'
```

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 6.6 (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_85"
OpenJDK Runtime Environment (rhel-2.6.1.3.el6_6-x86_64 u85-b01)
OpenJDK 64-Bit Server VM (build 24.85-b03, mixed mode)

**Updated from previous version*

Image File Sizes

VMware – 8.0 GB
Virtualbox – 7.9 GB
HyperV – 7.7 GB

Tech Preview Packages

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

- Ambari Views – storm-view-0.1.0.

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 Atlas

atlas-metadata_2_3_0_0_2557-hive-plugin-0.5.0.2.3.0.0-2557.el6.noarch

atlas-metadata_2_3_0_0_2557-0.5.0.2.3.0.0-2557.el6.noarch

Hadoop Search

lucidworks-hdpsearch-2.3-.noarch (contains Apache Solr and Banana)

Installed Packages

These are a list of all the HDP installed packages in the Sandbox.

Apache Ambari

ambari-metrics-collector-2.1.0-1470.x86_64

ambari-metrics-monitor-2.1.0-1470.x86_64
ambari-metrics-hadoop-sink-2.1.0-1470.x86_64
ambari-server-2.1.0-1470.x86_64
ambari-agent-2.1.0-1470.x86_64

Note: *ambari-log4j-2.0.0.1238-1.noarch* is no longer provided

Apache Ambari Views

- capacity-scheduler-2.1.0.1470
- files-2.1.0.1470
- pig-2.1.0.1470
- tez-view-2.1.0.1470

Apache Hadoop (HDFS, YARN, Mapreduce)

*hadoop_2_3_0_0_2557-2.7.1.2.3.0.0-2557.el6.x86_64
hadoop_2_3_0_0_2557-hdfs-2.7.1.2.3.0.0-2557.el6.x86_64
hadoop_2_3_0_0_2557-libhdfs-2.7.1.2.3.0.0-2557.el6.x86_64
hadoop_2_3_0_0_2557-yarn-2.7.1.2.3.0.0-2557.el6.x86_64
hadoop_2_3_0_0_2557-client-2.7.1.2.3.0.0-2557.el6.x86_64
hadoop_2_3_0_0_2557-mapreduce-2.7.1.2.3.0.0-2557.el6.x86_64

*Previous version where there were separate packages are now bundled into one package.

Apache Falcon

falcon_2_3_0_0_2557-doc-0.6.1.2.3.0.0-2557.el6.noarch
falcon_2_3_0_0_2557-0.6.1.2.3.0.0-2557.el6.noarch

Apache Hive

hive_2_3_0_0_2557-hcatalog-1.2.1.2.3.0.0-2557.el6.noarch
hive_2_3_0_0_2557-server-1.2.1.2.3.0.0-2557.el6.noarch
hive_2_3_0_0_2557-1.2.1.2.3.0.0-2557.el6.noarch
hive_2_3_0_0_2557-webhcat-1.2.1.2.3.0.0-2557.el6.noarch
hive_2_3_0_0_2557-hcatalog-server-1.2.1.2.3.0.0-2557.el6.noarch
hive_2_3_0_0_2557-server2-1.2.1.2.3.0.0-2557.el6.noarch
hive_2_3_0_0_2557-jdbc-1.2.1.2.3.0.0-2557.el6.noarch
hive_2_3_0_0_2557-webhcat-server-1.2.1.2.3.0.0-2557.el6.noarch
hive_2_3_0_0_2557-metastore-1.2.1.2.3.0.0-2557.el6.noarch

Apache Hbase

hbase_2_3_0_0_2557-thrift2-1.1.1.2.3.0.0-2557.el6.noarch
hbase_2_3_0_0_2557-master-1.1.1.2.3.0.0-2557.el6.noarch
hbase_2_3_0_0_2557-1.1.1.2.3.0.0-2557.el6.noarch

hbase_2_3_0_0_2557-rest-1.1.1.2.3.0.0-2557.el6.noarch
hbase_2_3_0_0_2557-thrift-1.1.1.2.3.0.0-2557.el6.noarch
hbase_2_3_0_0_2557-doc-1.1.1.2.3.0.0-2557.el6.noarch
hbase_2_3_0_0_2557-regionserver-1.1.1.2.3.0.0-2557.el6.noarch

Apache Flume

flume_2_3_0_0_2557-1.5.2.2.3.0.0-2557.el6.noarch
flume_2_3_0_0_2557-agent-1.5.2.2.3.0.0-2557.el6.noarch

Apache Hue

hue-common-2.6.1.2.3.0.0-2557.el6.x86_64
hue-oozie-2.6.1.2.3.0.0-2557.el6.x86_64
hue-sandbox-1.2.1-88.noarch
hue-beeswax-2.6.1.2.3.0.0-2557.el6.x86_64
hue-pig-2.6.1.2.3.0.0-2557.el6.x86_64
hue-server-2.6.1.2.3.0.0-2557.el6.x86_64
hue-tutorials-1.2.1-88.noarch
hue-hcatalog-2.6.1.2.3.0.0-2557.el6.x86_64
hue-2.6.1.2.3.0.0-2557.el6.x86_64

Apache Kafka

kafka_2_3_0_0_2557-0.8.2.2.3.0.0-2557.el6.noarch

Apache Knox

knox_2_3_0_0_2557-0.6.0.2.3.0.0-2557.el6.noarch

Apache Oozie

oozie_2_3_0_0_2557-4.2.0.2.3.0.0-2557.el6.noarch
oozie_2_3_0_0_2557-client-4.2.0.2.3.0.0-2557.el6.noarch

Apache Phoenix

phoenix_2_3_0_0_2557-4.4.0.2.3.0.0-2557.el6.noarch

Apache Pig

pig_2_3_0_0_2557-0.15.0.2.3.0.0-2557.el6.noarch

Apache Ranger

ranger_2_3_0_0_2557-hive-plugin-0.5.0.2.3.0.0-2557.el6.x86_64
ranger_2_3_0_0_2557-debuginfo-0.5.0.2.3.0.0-2557.el6.x86_64
ranger_2_3_0_0_2557-yarn-plugin-0.5.0.2.3.0.0-2557.el6.x86_64
ranger_2_3_0_0_2557-kafka-plugin-0.5.0.2.3.0.0-2557.el6.x86_64
ranger_2_3_0_0_2557-knox-plugin-0.5.0.2.3.0.0-2557.el6.x86_64
ranger_2_3_0_0_2557-usersync-0.5.0.2.3.0.0-2557.el6.x86_64
ranger_2_3_0_0_2557-admin-0.5.0.2.3.0.0-2557.el6.x86_64
ranger_2_3_0_0_2557-kms-0.5.0.2.3.0.0-2557.el6.x86_64

ranger_2_3_0_0_2557-solr-plugin-0.5.0.2.3.0.0-2557.el6.x86_64
ranger_2_3_0_0_2557-hdfs-plugin-0.5.0.2.3.0.0-2557.el6.x86_64
ranger_2_3_0_0_2557-storm-plugin-0.5.0.2.3.0.0-2557.el6.x86_64

Apache Solr (Included in the Hadoop Search package)

Solr 5.21 lucidworks-hdsearch

Apache Slider

slider_2_3_0_0_2557-0.80.0.2.3.0.0-2557.el6.noarch

Apache Spark

spark_2_3_0_0_2557-python-1.3.1.2.3.0.0-2557.el6.noarch
spark_2_3_0_0_2557-1.3.1.2.3.0.0-2557.el6.noarch
spark_2_3_0_0_2557-worker-1.3.1.2.3.0.0-2557.el6.noarch
spark_2_3_0_0_2557-master-1.3.1.2.3.0.0-2557.el6.noarch

Apache Sqoop

sqoop_2_3_0_0_2557-metastore-1.4.6.2.3.0.0-2557.el6.noarch
sqoop_2_3_0_0_2557-1.4.6.2.3.0.0-2557.el6.noarch

Apache Storm

storm_2_3_0_0_2557-slider-client-0.10.0.2.3.0.0-2557.el6.x86_64
storm_2_3_0_0_2557-0.10.0.2.3.0.0-2557.el6.x86_64

Apache Tez

tez_2_3_0_0_2557-0.7.0.2.3.0.0-2557.el6.noarch

Apache Zookeeper

zookeeper_2_3_0_0_2557-3.4.6.2.3.0.0-2557.el6.noarch

Other Packages

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

Python

python-devel-2.6.6-52.el6.x86_64
python-pycurl-7.19.0-8.el6.x86_64
python-iniparse-0.3.1-2.1.el6.noarch
python-libs-2.6.6-52.el6.x86_64
python-urlgrabber-3.9.1-9.el6.noarch
newt-python-0.52.11-3.el6.x86_64

```
rpm-python-4.8.0-38.el6_6.x86_64  
python-lxml-2.2.3-1.1.el6.x86_64  
python-2.6.6-52.el6.x86_64
```

*difference from previous version, not shipping python-pip and python-setuptools

mysql

```
mysql-server-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-5.1.73-5.el6_6.x86_64
```

Postgres

```
postgresql-server-8.4.20-3.el6_6.x86_64  
postgresql-8.4.20-3.el6_6.x86_64  
postgresql-libs-8.4.20-3.el6_6.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*

HDFS

SecondaryNameNode -
org.apache.hadoop.hdfs.server.namenode.SecondaryNameNode
*Portmap - org.apache.hadoop.portmap.Portmap
NameNode - org.apache.hadoop.hdfs.server.namenode.NameNode
DataNode - org.apache.hadoop.hdfs.server.datanode.DataNode
*Nfs3 - org.apache.hadoop.hdfs.nfs.Nfs3

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

HIVE

*RunJar - webhcatt - 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

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

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.

Atlas

Main - org.apache.atlas.Main

Falcon

Main - org.apache.falcon.Main

Flume

Application - org.apache.flume.node.Application

HBase

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

HMMaster - 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

Ranger (XaSecure)

*UnixAuthenticationService - com.xasecure.authentication.UnixAuthenticationService
EmbededServer - com.xasecure.server.tomcat.EmbededServer

**Run as root user*

Spark

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

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
 - HBase
 - Falcon
 - Storm
 - Flume
 - Ambari Metrics
 - Atlas
 - Kafka
 - Knox
 - Spark