



Technical Preview HDP 2.3

Release Notes

June 2015

Md5 **VMware** Virtual Appliance - d51b833200dce60a67b6456daae3cc02

Md5 **Virtualbox** Virtual Appliance- 52ee9dcd3209202c68cce69d908d38c0

HDP Stack and Ambari

The Sandbox uses the following versions of Ambari and HDP stack.

HDP 2.3 Build 2130

Ambari 2.1 Build 781

Known Issues

Since this is an early preview of a release, you may experience some unexpected behaviors. Below are some issues you may encounter.

BUG-38937 – Pig view not able to show parsing errors, job says completed.

Component: Pig View

BUG-38787 – Opening pig view from hive view shows 500 error (very intermittent) just refresh to workaround

Component: Pig View

BUG-38982 - When click on the Ambari views button on the top right , if you are in the welcome ambari page you, the complete list of views are not listed.

Component: Ambari

Limitations

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

RMP-3553 - Ranger and Knox are 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 stop

Knox

start: su -l knox -c "/usr/hdp/current/knox-server/bin/gateway.sh start"

stop: su -l knox -c "/usr/hdp/current/knox-server/bin/gateway.sh stop"

For knox you also 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_79"

OpenJDK Runtime Environment (rhel-2.5.5.3.el6_6-x86_64 u79-b14)

OpenJDK 64-Bit Server VM (build 24.79-b02, mixed mode)

**Updated from previous version*

Image File Sizes

VMware – 6.6 GB

Virtualbox –7 GB

Tech Preview Packages

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

- Ambari Views – Storm

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 Datafu
- Apache Mahout
- Solr

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_2130-0.1.0.2.3.0.0-2130.el6.noarch

atlas-metadata_2_3_0_0_2130-hive-plugin-0.1.0.2.3.0.0-2130.el6.noarch

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.Nfs3

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

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

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.

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*

Ranger (XaSecure)

*UnixAuthenticationService - com.xasecure.authentication.UnixAuthenticationService

EmbeddedServer - com.xasecure.server.tomcat.EmbeddedServer

**Run as root user*

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

Apache Ambari Notes

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
 - Spark
 - Kafka
 - Ambari Metrics
 - Atlas