Hortonworks Sandbox
Release Notes

May 2015

Md5 **VMware** Virtual Appliance - 470bef8bbec54cc4e2a2e310e174d910
Md5 **Virtualbox** Virtual Appliance- 46f438580fec16deca1c4c2b5574f914
Md5 **HyperVisor** - d9499e7b023a1fc2eccdcb62cf3f8bdc

HDP Stack and Ambari

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

Ambari 2.0 Release Notes

http://docs.hortonworks.com/HDPDocuments/Ambari-2.0.0.0/Ambari_RelNotes_v20/index.html#Item1.1

HDP 2.2.4.2 Release Notes

http://docs.hortonworks.com/HDPDocuments/HDP2/HDP-2.2.4/HDP_RelNotes_v224/index.html

Fixed Issues

List of Issues that have been fixed from previous release.

BUG-34310 - Receiving Error dialog display when trying to modify HBase config within Ambari in Sandbox "500 status code received GET method for API /api/v1/stacks/HDP/versions/2.2/validations"
Component: Ambari

BUG-34455 - No audits for the Knox repository type in Ranger
Component: Sandbox

BUG-34692 - Splash screen of new sandbox hdp 2.2.4.2 is pointing to HDP 2.2
Component: Sandbox

BUG-35290 - Hive policy has the config property for JDBC url "jdbc:hive2://sandbox.hortonworks.com:10000/default;auth=noSasl" which causes multiple threads to the hiveserver2 that cause slowness or use up all the connections.
Component: Ranger

BUG-34485 - Hive service showing it has to be restarted due to configuration changes
Component: Sandbox
Workaround: After starting Ambari Server and Ambari Agent, log into Ambari dashboard and restart Hive service.

BUG-34609 - Virtualbox image still references the name as 2.2.4 which should be 2.2.4.2
Component: Sandbox

Copyright © 2012 - 2015 Hortonworks, Inc. All rights reserved.

Known Issues

List of issues you may encounter when using this version of the Sandbox.

BUG-36371- Hyper V Sandbox is not able to reach the internet
Component: Sandbox (Specific to HyperV Sandbox can only reach the internet if your
Workaround: At the moment, the Hyper Sandbox can only reach the internet if your host machine adapter that the HyperV is using is in the subnet 192.168.56.x

BUG-34671- Ambari shows alerts for some of the services that are started, when ambari is first launched.
Component: Sandbox
Workaround: If you simply wait a few minutes, the alerts will sync up and no longer show as an alert.

BUG-34592 - Broken links for Yarn webui's such as: nodemanager http://127.0.0.1:8042/ which prevents allowing access to the nodemanager info and logs
Component: Sandbox
Workaround: For VMware the workaround is to replace "sandbox.hortonworks.com" with the IP address used for the virtual machine, typically 192.xxx.xxx.xxx

BUG-34485 - The about page for Hue http://127.0.0.1:8000/about/ does not have complete listing of components
Component: Sandbox

BUG-33940 - Jps output is showing as "process information unavailable" in HDP 2.2.4 Ambari 2.0 sandbox (JDK bug https://bugs.launchpad.net/ubuntu/+source/openjdk-6/+bug/1417962)
Component: Sandbox

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

Copyright © 2012 - 2015 Hortonworks, Inc. All rights reserved.

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
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"

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.

RMP-3652 – THP is enabled by default which may degrade performance of Hadoop
Workaround:
How to disable THP?
Red Hat/CentOS: /sys/kernel/mm/redhat_transparent_hugepage/defrag
To see whether transparent hugepage compaction is enabled, run the following command and check the output:
$ cat defrag_file_pathname
always never means that transparent hugepage compaction is enabled.
always never means that transparent hugepage compaction is disabled.

To disable transparent hugepage compaction, add the following command to /etc/rc.local :
echo never > defrag_file_pathname
You can also disable transparent hugepage compaction interactively (The setting do not persists over a reboot).
To disable transparent hugepage compaction temporarily as root:
echo 'never' > defrag_file_pathname
To disable transparent hugepage compaction temporarily using sudo:
$ sudo sh -c "echo 'never' > defrag_file_pathname"

System Information

Operating System and Java versions that the Sandbox uses in this version.

Copyright © 2012 - 2015 Hortonworks, Inc. All rights reserved.

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_75"
OpenJDK Runtime Environment (rhel-2.5.4.0.el6_6-x86_64 u75-b13)
OpenJDK 64-Bit Server VM (build 24.75-b04, mixed mode)

**Updated from previous version*

Image File Sizes
VMware – 5.4 GB
Virtualbox – 5.4 GB
HyperV – 5.2 GB

Tech Preview Packages

These are list of packages installed in the Sandbox but are still in tech preview and currently not supported. These packages are included to allow an early look as to what we are working on.

- Ambari Views - Capacity Scheduler 0.3
- Ambari Views - Files 0.1
- Ambari Views - Hive 0.2
- Ambari Views - Pig 0.1
- Ambari Views - Tez 0.5.2.2.2.2.0 -1238

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

Copyright © 2012 - 2015 Hortonworks, Inc. All rights reserved.

- 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

Newly Added HDP Supported Packages

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

Apache Spark
spark_2_2_4_2_2-python-1.2.1.2.2.4.2-2.el6.noarch
spark_2_2_4_2_2-1.2.1.2.2.4.2-2.el6.noarch
spark_2_2_4_2_2-worker-1.2.1.2.2.4.2-2.el6.noarch
spark_2_2_4_2_2-master-1.2.1.2.2.4.2-2.el6.noarch

Installed Packages

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

***NOTE**: Ganglia and Nagios are no longer installed*

Apache Ambari
*ambari-metrics-collector-2.0.0-151.x86_64
*ambari-metrics-monitor-2.0.0-151.x86_64
*ambari-metrics-hadoop-sink-2.0.0-151.x86_64

ambari-log4j-2.0.0.151-1.noarch
ambari-server-2.0.0-151.noarch
ambari-agent-2.0.0-151.x86_64
***New packages that replace functionality of Ganglia and Nagios in past builds*

Apache Ambari Views

- Slider Apps View 1.0

Copyright © 2012 - 2015 Hortonworks, Inc. All rights reserved.

Apache Hadoop (HDFS, YARN, Mapreduce)

hadoop_2_2_4_2_2-mapreduce-2.6.0.2.2.4.2-2.el6.x86_64
hadoop_2_2_4_2_2-hdfs-2.6.0.2.2.4.2-2.el6.x86_64
hadoop_2_2_4_2_2-libhdfs-2.6.0.2.2.4.2-2.el6.x86_64
hadoop_2_2_4_2_2-2.6.0.2.2.4.2-2.el6.x86_64
hadoop_2_2_4_2_2-yarn-2.6.0.2.2.4.2-2.el6.x86_64
hadoop_2_2_4_2_2-client-2.6.0.2.2.4.2-2.el6.x86_64
*hadoop_2_2_4_2_2-yarn-nodemanager-2.6.0.2.2.4.2-2.el6.x86_64
*hadoop_2_2_4_2_2-httpfs-2.6.0.2.2.4.2-2.el6.x86_64
*hadoop_2_2_4_2_2-yarn-proxyserver-2.6.0.2.2.4.2-2.el6.x86_64
*hadoop_2_2_4_2_2-hdfs-datanode-2.6.0.2.2.4.2-2.el6.x86_64
*hadoop_2_2_4_2_2-hdfs-secondarynamenode-2.6.0.2.2.4.2-2.el6.x86_64
*hadoop_2_2_4_2_2-yarn-resourcemanager-2.6.0.2.2.4.2-2.el6.x86_64
*hadoop_2_2_4_2_2-mapreduce-historyserver-2.6.0.2.2.4.2-2.el6.x86_64
*hadoop_2_2_4_2_2-hdfs-namenode-2.6.0.2.2.4.2-2.el6.x86_64
*hadoop_2_2_4_2_2-hdfs-fuse-2.6.0.2.2.4.2-2.el6.x86_64
*hadoop_2_2_4_2_2-doc-2.6.0.2.2.4.2-2.el6.x86_64
*hadoop_2_2_4_2_2-source-2.6.0.2.2.4.2-2.el6.x86_64
*hadoop_2_2_4_2_2-hdfs-zkfc-2.6.0.2.2.4.2-2.el6.x86_64
*hadoop_2_2_4_2_2-conf-pseudo-2.6.0.2.2.4.2-2.el6.x86_64
*hadoop_2_2_4_2_2-hdfs-journalnode-2.6.0.2.2.4.2-2.el6.x86_64

*New separate packages that were previously bundled into larger packages in previous builds

Apache Falcon

falcon_2_2_4_2_2-doc-0.6.0.2.2.4.2-2.el6.noarch
falcon_2_2_4_2_2-0.6.0.2.2.4.2-2.el6.noarch

Apache Hive

hive_2_2_4_2_2-webhcat-0.14.0.2.2.4.2-2.el6.noarch
hive_2_2_4_2_2-metastore-0.14.0.2.2.4.2-2.el6.noarch
hive_2_2_4_2_2-hcatalog-0.14.0.2.2.4.2-2.el6.noarch
hive_2_2_4_2_2-webhcat-server-0.14.0.2.2.4.2-2.el6.noarch
hive_2_2_4_2_2-server2-0.14.0.2.2.4.2-2.el6.noarch
hive_2_2_4_2_2-server-0.14.0.2.2.4.2-2.el6.noarch
hive_2_2_4_2_2-0.14.0.2.2.4.2-2.el6.noarch
hive_2_2_4_2_2-hcatalog-server-0.14.0.2.2.4.2-2.el6.noarch
hive_2_2_4_2_2-jdbc-0.14.0.2.2.4.2-2.el6.noarch

Copyright © 2012 - 2015 Hortonworks, Inc. All rights reserved.

Apache Hbase

hbase_2_2_4_2_2-rest-0.98.4.2.2.4.2-2.el6.noarch
hbase_2_2_4_2_2-doc-0.98.4.2.2.4.2-2.el6.noarch
hbase_2_2_4_2_2-master-0.98.4.2.2.4.2-2.el6.noarch
hbase_2_2_4_2_2-thrift2-0.98.4.2.2.4.2-2.el6.noarch
hbase_2_2_4_2_2-thrift-0.98.4.2.2.4.2-2.el6.noarch
hbase_2_2_4_2_2-0.98.4.2.2.4.2-2.el6.noarch
hbase_2_2_4_2_2-regionserver-0.98.4.2.2.4.2-2.el6.noarch

Apache Flume

flume_2_2_4_2_2-agent-1.5.2.2.2.4.2-2.el6.noarch
flume_2_2_4_2_2-1.5.2.2.2.4.2-2.el6.noarch

Apache Hue

hue-hcatalog-2.6.1.2.2.4.2-2.el6.x86_64
hue-2.6.1.2.2.4.2-2.el6.x86_64
hue-beeswax-2.6.1.2.2.4.2-2.el6.x86_64
hue-pig-2.6.1.2.2.4.2-2.el6.x86_64
hue-oozie-2.6.1.2.2.4.2-2.el6.x86_64
hue-tutorials-1.2.1-88.noarch
hue-common-2.6.1.2.2.4.2-2.el6.x86_64
hue-server-2.6.1.2.2.4.2-2.el6.x86_64
hue-sandbox-1.2.1-88.noarch

Apache Kafka

kafka_2_2_4_2_2-0.8.1.2.2.4.2-2.el6.noarch

Apache Knox

knox_2_2_4_2_2-0.5.0.2.2.4.2-2.el6.noarch

Apache Oozie

oozie_2_2_4_2_2-4.1.0.2.2.4.2-2.el6.noarch
oozie_2_2_4_2_2-client-4.1.0.2.2.4.2-2.el6.noarch

Apache Phoenix

Copyright © 2012 - 2015 Hortonworks, Inc. All rights reserved.

phoenix_2_2_4_2_2-4.2.0.2.2.4.2-2.el6.noarch

Apache Pig
pig_2_2_4_2_2-0.14.0.2.2.4.2-2.el6.noarch

Apache Ranger
ranger_2_2_4_2_2-debuginfo-0.4.0.2.2.4.2-2.el6.x86_64
ranger_2_2_4_2_2-hdfs-plugin-0.4.0.2.2.4.2-2.el6.x86_64
ranger_2_2_4_2_2-usersync-0.4.0.2.2.4.2-2.el6.x86_64
ranger_2_2_4_2_2-hbase-plugin-0.4.0.2.2.4.2-2.el6.x86_64
ranger_2_2_4_2_2-hive-plugin-0.4.0.2.2.4.2-2.el6.x86_64
ranger_2_2_4_2_2-knox-plugin-0.4.0.2.2.4.2-2.el6.x86_64
ranger_2_2_4_2_2-admin-0.4.0.2.2.4.2-2.el6.x86_64
ranger_2_2_4_2_2-storm-plugin-0.4.0.2.2.4.2-2.el6.x86_64

Apache Solr
solr-4.10.4 (tarball Install) /opt/solr/solr-4.10.4

Apache Slider
slider_2_2_4_2_2-0.61.0.2.2.4.2-2.el6.noarch

Apache Sqoop
sqoop_2_2_4_2_2-1.4.5.2.2.4.2-2.el6.noarch
sqoop_2_2_4_2_2-metastore-1.4.5.2.2.4.2-2.el6.noarch

Apache Storm
storm_2_2_4_2_2-slider-client-0.9.3.2.2.4.2-2.el6.x86_64
storm_2_2_4_2_2-0.9.3.2.2.4.2-2.el6.x86_64

Apache Tez
tez_2_2_4_2_2-0.5.2.2.2.4.2-2.el6.noarch

Apache Zookeeper
*zookeeper_2_2_4_2_2-server-3.4.6.2.2.4.2-2.el6.noarch
zookeeper_2_2_4_2_2-3.4.6.2.2.4.2-2.el6.noarch

Copyright © 2012 - 2015 Hortonworks, Inc. All rights reserved.

*New separate packages that were previously bundled in previous builds

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-pip-1.3.1-4.el6.noarch
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-setuptools-0.6.10-3.el6.noarch
python-lxml-2.2.3-1.1.el6.x86_64
python-2.6.6-52.el6.x86_64

mysql
mysql-server-5.1.73-3.el6_5.x86_64
mysql-libs-5.1.73-3.el6_5.x86_64
mysql-5.1.73-3.el6_5.x86_64
mysql-connector-java-5.1.17-6.el6.noarch

Postgres
postgresql-server-8.4.20-2.el6_6.x86_64
postgresql-8.4.20-2.el6_6.x86_64
postgresql-libs-8.4.20-2.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.

Copyright © 2012 - 2015 Hortonworks, Inc. All rights reserved.

Falcon
Main - org.apache.falcon.Main

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*

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

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
EmbededServer - com.xasecure.server.tomcat.EmbededServer

Copyright © 2012 - 2015 Hortonworks, Inc. All rights reserved.

**Run as root user*

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.

Ambari
*AmbariServer - org.apache.ambari.server.controller.AmbariServer
Ambari Agent (non java process)
**Run as root user*

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

Storm
supervisor - backtype.storm.daemon.supervisor
nimbus - backtype.storm.daemon.nimbus
logviewer - backtype.storm.daemon.logviewer

Copyright © 2012 - 2015 Hortonworks, Inc. All rights reserved.

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

Copyright © 2012 - 2015 Hortonworks, Inc. All rights reserved.