

Hortonworks Data Platform

HDP-2.2.8 Release Notes

(Sep 4, 2015)

Hortonworks Data Platform: HDP-2.2.8 Release Notes

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

The Hortonworks Data Platform, powered by Apache Hadoop, is a massively scalable and 100% open source platform for storing, processing and analyzing large volumes of data. It is designed to deal with data from many sources and formats in a very quick, easy and cost-effective manner. The Hortonworks Data Platform consists of the essential set of Apache Hadoop projects including MapReduce, Hadoop Distributed File System (HDFS), HCatalog, Pig, Hive, HBase, ZooKeeper and Ambari. Hortonworks is the major contributor of code and patches to many of these projects. These projects have been integrated and tested as part of the Hortonworks Data Platform release process and installation and configuration tools have also been included.

Unlike other providers of platforms built using Apache Hadoop, Hortonworks contributes 100% of our code back to the Apache Software Foundation. The Hortonworks Data Platform is Apache-licensed and completely open source. We sell only expert technical support, [training](#) and partner-enablement services. All of our technology is, and will remain, free and open source.

Please visit the [Hortonworks Data Platform](#) page for more information on Hortonworks technology. For more information on Hortonworks services, please visit either the [Support](#) or [Training](#) page. Feel free to [contact us](#) directly to discuss your specific needs.



Except where otherwise noted, this document is licensed under
Creative Commons Attribution ShareAlike 3.0 License.
<http://creativecommons.org/licenses/by-sa/3.0/legalcode>

Table of Contents

1. HDP 2.2.8 Release Notes	1
1.1. New Features	2
1.2. Unsupported Features	2
1.2.1. Technical Preview Features	2
1.2.2. Community Features	3
1.3. Upgrading from HDP 2.2.x to HDP 2.2.8	3
1.3.1. Before you begin	4
1.3.2. Optional: Spark Manual Upgrade Procedure	4
1.3.3. Upgrade Procedure	5
1.3.4. Optional: Spark Manual Downgrade Procedure	8
1.4. Behavioral Changes	8
1.5. Apache Patch Information	9
1.5.1. Accumulo 1.6.1	9
1.5.2. Falcon 0.6.0	9
1.5.3. Hadoop Common/HDFS 2.6.0	9
1.5.4. HBase 0.98.4	12
1.5.5. Hive 0.14.0	12
1.5.6. Knox 0.5.0	19
1.5.7. Oozie 4.1.0	19
1.5.8. Phoenix 4.2.0	19
1.5.9. Pig 0.14.0	21
1.5.10. Ranger 0.4.0	21
1.5.11. Slider 0.61.0	21
1.5.12. Storm 0.9.3	21
1.5.13. Tez 0.5.2	21
1.5.14. ZooKeeper 3.4.6	23
1.6. Common Vulnerabilities and Exposures	24
1.7. Third-party Licenses	25
1.8. Fixed Issues	25
1.9. Known Issues for HDP 2.2.8	34

1. HDP 2.2.8 Release Notes

The official Apache versions of most HDP 2.2 components are unchanged from HDP 2.2.0.0, with the exception of Spark. Spark is upgraded from 1.2.1 to 1.3.1. See more details of Spark 1.3.1 in the [New Features](#) section. All HDP 2.2 components listed here are official Apache releases of the most recent stable versions available.

Hortonworks' philosophy is to provide patches only when absolutely necessary to assure the interoperability of the components. Unless you are explicitly directed by Hortonworks Support to take a patch update, each of the HDP 2.2 components needs to remain at the following package version levels to ensure a certified and supported copy of HDP 2.2.

- Apache Accumulo 1.6.1
- Apache DataFu 1.2.0
- Apache Falcon 0.6.0
- Apache Flume 1.5.2
- Apache Hadoop 2.6.0
- Apache HBase 0.98.4
- Apache Hive 0.14.0
- Apache Kafka 0.8.1.1
- Apache Knox 0.5.0
- Apache Mahout 0.9.0
- Apache Oozie 4.1.0
- Apache Phoenix 4.2.0
- Apache Pig 0.14.0
- Apache Ranger 0.4.0
- Apache Slider 0.61.0
- Apache Spark 1.3.1
- Apache Sqoop 1.4.5
- Apache Storm 0.9.3
- Apache Tez 0.5.2
- Apache ZooKeeper 3.4.6
- Hue 2.6.1

As of HDP 2.2, the following third-party tools are deprecated, and will be removed in a future release:

- Ganglia 3.5.0

- Nagios 3.5.0
- Ganglia Web 3.5.7

1.1. New Features

This section highlights the only new feature in HDP 2.2.8.

Spark 1.3.1 is the default Spark version with HDP 2.2.8. With a new HDP 2.2.8 cluster install, Spark 1.3.1 is installed. With the upgrade of an existing HDP cluster to 2.2.8 using Ambari, the Spark component is automatically upgraded from 1.2.1 to 1.3.1. For clusters not managed with Ambari, Spark can be manually upgraded from 1.2.1 to Spark 1.3.1; see the [Optional Spark Manual Upgrade Procedure](#) for instructions.

If you have upgraded to Spark 1.3.1 and want to downgrade to Spark 1.2.1, follow the [Optional Spark Manual Downgrade Procedure](#). The downgrade to Spark 1.2.1 is only available as a manual step.

- In HDP 2.2.8, Spark 1.3.1 is now available as a certified component.
 - Spark Core
 - Spark on YARN
 - Spark History Server
 - Spark on YARN on Kerberos-enabled clusters
 - Spark MLLib
 - Support for Hive 0.13.1, including the collect_list UDF

1.2. Unsupported Features

Some features exist within HDP 2.2.8, but Hortonworks does not currently support these specific capabilities.

1.2.1. Technical Preview Features

The following features are available within HDP 2.2.8, but are not ready for production deployment. We encourage you to explore these technical preview features in non-production environments and provide feedback on your experiences through the Hortonworks Community Forums.

Component	Feature
HDFS	• HDFS Transparent Data at Rest Encryption
Phoenix	• Phoenix support for local secondary indexes
YARN	• YARN support for Docker
Spark	• Spark SQL • Spark Streaming • Spark DataFrame API

Component	Feature
	<ul style="list-style-type: none"> • ORC file support • SparkSQL thrift JDBC/ODBC Server • Dynamic Executor Allocation

1.2.2. Community Features

The following features are developed and tested by the community, but are not officially supported by Hortonworks. There are variety of reasons that these features are excluded, including: insufficient reliability or incomplete test case coverage, declaration of non-production readiness by the community at large, feature deviates from Hortonworks best practices, and more. Do not use them in your production environments.

Component	Feature
HDFS	<ul style="list-style-type: none"> • NameNode Federation (HDFS-1052) • viewFS (HADOOP-7257) • block-volume device choosing (HDFS-1804)
Falcon	<ul style="list-style-type: none"> • Prism Server
YARN	<ul style="list-style-type: none"> • Fair Scheduler • MapReduce Uber AM • MapReduce Eclipse Plug-in
Spark	<ul style="list-style-type: none"> • Spark Standalone • BlinkDB • GraphX

1.3. Upgrading from HDP 2.2.x to HDP 2.2.8

HDP 2.2.8 is a maintenance release of HDP 2.2.x.

If you already have HDP 2.2.x installed, upgrading your cluster to HDP 2.2.8 means:

- Keeping the same configuration files you used for HDP 2.2.x
- Keeping the same data and metadata in the same location you used for HDP 2.2.x
- Installing any new components (added for the first time in HDP 2.2.8) side-by-side with existing components

The following table summarizes HDP 2.2.x-to-2.2.8 upgrade options:

Cluster Management	Supporting Doc	Notes
Cluster managed manually	Use these Release Notes	If you have an earlier version of HDP (such as HDP 2.0 or HDP 2.1), see the HDP 2.2.8 Manual Upgrade Guide
Cluster managed via Ambari 1.7.0	Use these Release Notes	
Cluster managed via Ambari 2.0	Use the Upgrading Ambari Guide	Ambari 2.0 supports rolling upgrade between HDP 2.2.x and HDP 2.2.8 When upgrading to HDP 2.2.8 using Ambari, Spark 1.2.1 will be automatically upgraded to 1.3.1. If you

Cluster Management	Supporting Doc	Notes
		wish to return to using 1.2.1, use the Spark Manual Downgrade Procedure .
Cluster managed via Ambari 2.1	Use the Upgrading Ambari Guide	Ambari 2.1 supports rolling upgrade between HDP 2.2.x and HDP 2.2.8 When upgrading to HDP 2.2.8 using Ambari, Spark 1.2.1 will be automatically upgraded to 1.3.1. If you wish to return to using 1.2.1, use the Spark Manual Downgrade Procedure .

1.3.1. Before you begin

Before You Begin

- Make sure you know what HDP components need to be upgraded at your installation
- Think about whether you are going to upgrade using a [local repository](#) or a [remote repository](#)

1.3.2. Optional: Spark Manual Upgrade Procedure

Spark 1.3.1 is newly available in HDP 2.2.8. You can choose to upgrade Spark from 1.2.1 to 1.3.1, or continue using Spark 1.2.1.

If you wish to upgrade to Spark 1.3.1, perform the following steps as root:

1. Install Spark:

- Enter `wget -nv http://public-repo-1.hortonworks.com/HDP/centos6/2.x/updates/2.2.8.0/hdp.repo -O /etc/yum.repos.d/hdp228.repo`.
- Enter `yum search spark`.

The prompt returns:

```
spark.noarch : spark Distro virtual package
spark-master.noarch : spark-master Distro virtual package
spark-python.noarch : spark-python Distro virtual package
spark-worker.noarch : spark-worker Distro virtual package
spark_2_2_8_0_3150-master.noarch : Server for Spark master
spark_2_2_8_0_3150-python.noarch : Python client for Spark
spark_2_2_8_0_3150-worker.noarch : Server for Spark worker
spark_2_2_8_0_3150-master.noarch : Server for Spark master
spark_2_2_8_0_3150-python.noarch : Python client for Spark
spark_2_2_8_0_3150-worker.noarch : Server for Spark worker
spark_2_2_8_0_3150.noarch : Lightning-Fast Cluster Computing
spark_2_2_8_0_3150.noarch : Lightning-Fast Cluster Computing
```

- To install Spark 1.3.1, select `spark_2_2_8_0_3150`: `yum install spark_2_2_8_0_3150-master`.
- Enter `hdp-select set spark-client 2.2.8.0-3150`.
- Enter `hdp-select set spark-historyserver 2.2.8.0-3150`.

2. Validate the Spark installation:

- a. Enter `su spark`
- b. Enter `cd /usr/hdp/current/spark-client.`
- c. Enter `./bin/spark-submit --class org.apache.spark.examples.SparkPi --master yarn-client --num-executors 3 --driver-memory 512m --executor-memory 512m --executor-cores 1 lib/spark-examples*.jar 10.`

3. *Optional:* Run Spark History Server:

- a. Enter `su spark.`
- b. Enter `cd /usr/hdp/current/spark-historyserver/.`
- c. Enter `./sbin/start-history-server.sh.`

1.3.3. Upgrade Procedure

To upgrade your cluster from HDP 2.2.x to HDP 2.2.8:

1. Download the appropriate HDP 2.2.8 hdp.repo file for your OS:

Operating System	Repository Location
RHEL/CentOS/Oracle LINUX 5	http://public-repo-1.hortonworks.com/HDP/centos5/2.x/updates/2.2.8.0/hdp.repo
RHEL/CentOS/Oracle LINUX 6	http://public-repo-1.hortonworks.com/HDP/centos6/2.x/updates/2.2.8.0/hdp.repo
SLES 11 SP1	http://public-repo-1.hortonworks.com/HDP/sles11sp1/2.x/updates/2.2.8.0/hdp.repo
SLES 11 SP3/SP4	http://public-repo-1.hortonworks.com/HDP/suse11sp3/2.x/updates/2.2.8.0/hdp.repo
Ubuntu 12	http://public-repo-1.hortonworks.com/HDP/ubuntu12/2.x/updates/2.2.8.0/hdp.list
Debian 6	http://public-repo-1.hortonworks.com/HDP/debian6/2.x/updates/2.2.8.0/hdp.list

or

Download the HDP RPM single repository tarball. (For information on how to install the repositories, see the [local repository](#) instructions.)

Operating System	Tarball Location
RHEL/CentOS/Oracle LINUX 5	http://public-repo-1.hortonworks.com/HDP/centos5/2.x/updates/2.2.8.0/HDP-2.2.8.0-centos5-rpm.tar.gz
RHEL/CentOS/Oracle LINUX 6	http://public-repo-1.hortonworks.com/HDP/centos6/2.x/updates/2.2.8.0/HDP-2.2.8.0-centos6-rpm.tar.gz
SLES 11 SP1	http://public-repo-1.hortonworks.com/HDP/sles11sp1/2.x/updates/2.2.8.0/HDP-2.2.8.0-sles11sp1-rpm.tar.gz

Operating System	Tarball Location
SLES 11 SP3/SP4	http://public-repo-1.hortonworks.com/HDP/suse11sp3/2.x/updates/2.2.8.0/HDP-2.2.8.0-suse11sp3-rpm.tar.gz
Ubuntu 12	http://public-repo-1.hortonworks.com/HDP/ubuntu12/2.x/updates/2.2.8.0/HDP-2.2.8.0-ubuntu12-rpm.tar.gz
Debian 6	http://public-repo-1.hortonworks.com/HDP/debian6/2.x/updates/2.2.8.0/HDP-2.2.8.0-debian6-rpm.tar.gz

2. Run an update:

```
apt-get update
```

3. Install the HDP 2.2.8 bits:

Operating System	Commands
RHEL/CentOS/Oracle LINUX	Install HDP 2.2.8 components on relevant nodes, according to the services that run on those hosts: <pre>yum install "hadoop_2_2_8_0_3150*" "oozie_2_2_8_0_3150*" "pig_2_2_8_0_3150*" "sqoop_2_2_8_0_3150*" "zookeeper_2_2_8_0_3150*" "hbase_2_2_8_0_3150*" "hive_2_2_8_0_3150*" "tez_2_2_8_0_3150*" "storm_2_2_8_0_3150*" "falcon_2_2_8_0_3150*" "flume_2_2_8_0_3150*" "phoenix_2_2_8_0_3150*" "accumulo_2_2_8_0_3150*" "mahout_2_2_8_0_3150*"</pre>
SLES	Install HDP 2.2.8 components on relevant nodes, according to the services that run on those hosts: <pre>zypper install "hadoop_2_2_8_0_3150*" "oozie_2_2_8_0_3150*" "pig_2_2_8_0_3150*" "sqoop_2_2_8_0_3150*" "zookeeper_2_2_8_0_3150*" "hbase_2_2_8_0_3150*" "hive_2_2_8_0_3150*" "tez_2_2_8_0_3150*" "storm_2_2_8_0_3150*" "falcon_2_2_8_0_3150*" "flume_2_2_8_0_3150*" "phoenix_2_2_8_0_3150*" "accumulo_2_2_8_0_3150*" "mahout_2_2_8_0_3150*"</pre>
Ubuntu/Debian	Install HDP 2.2.8 components on relevant nodes, according to the services that run on those hosts: <pre>apt-get install "hadoop_2_2_8_0_3150*" "oozie_2_2_8_0_3150*" "pig_2_2_8_0_3150*" "sqoop_2_2_8_0_3150*" "zookeeper_2_2_8_0_3150*" "hbase_2_2_8_0_3150*" "hive_2_2_8_0_3150*" "tez_2_2_8_0_3150*" "storm_2_2_8_0_3150*" "falcon_2_2_8_0_3150*" "flume_2_2_8_0_3150*" "phoenix_2_2_8_0_3150*" "accumulo_2_2_8_0_3150*" "mahout_2_2_8_0_3150*"</pre>

4. Stop all HDP 2.2.x Services.

If you are managing your cluster manually, stop all HDP 2.2.x Services. See "Controlling HDP Services Manually" in the *HDP Reference Guide*.

If you are managing your cluster with Ambari 1.7.0, do the following:

- a. Open Ambari Web

- b. Browse to **Services**
- c. Use **Service Actions** to stop each service

For all services, switch the active version to HDP 2.2.8.

On each host in the cluster, use `hdps-select` to switch all services to the HDP 2.2.8 version:

```
hdps-select set all <hdp2.2.8 version>
```

5. (Ambari 1.7.0-managed clusters only) Complete the Stack Upgrade.

If you are managing your cluster with Ambari 1.7.0, update the repository Base URLs to use the HDP 2.2.8 repositories for HDP and HDP-UTILS:

- a. Open Ambari Web
- b. Browse to **Admin > Repositories**
- c. Edit the Base URLs

6. Start all HDP 2.2.8 services, in the following order:

a. ZooKeeper

```
su - zookeeper -c "export ZOOCFGDIR=/usr/hdp/current/
zookeeper-server/conf ; export ZOOCFG=zoo.cfg; source /usr/
hdps/current/zookeeper-server/conf/zookeeper-env.sh ; /usr/hdp/
current/zookeeper-server/bin/zkServer.sh start"
```

b. (HA NameNode upgrade only) ZooKeeper Failover Controller Daemons

```
/usr/hdp/current/hadoop-hdfs-namenode/./hadoop/sbin/hadoop-
daemon.sh start zkfc
```

c. (HA NameNode upgrade only) JournalNodes

```
su - hdfs /usr/hdp/current/hadoop-hdfs-journalnode/./hadoop/
sbin/hadoop-daemon.sh start journalnode
```

d. HDFS NameNode(s)

Start the HDFS NameNode(s). Because there is no metadata schema update for this upgrade, start the NameNode(s) in normal mode:

```
su - hdfs /usr/hdp/current/hadoop-hdfs-namenode/./hadoop/
sbin/hadoop-daemon.sh start namenode
```

e. Remaining Services

Start the rest of the HDP services. On each host in the cluster, start the services that are relevant to that cluster. To identify the start commands for all services, see "Controlling HDP Services Manually" in the *HDP Reference Guide*.

You now have an upgraded cluster. Ensure that your workloads run correctly on this upgraded cluster.

1.3.4. Optional: Spark Manual Downgrade Procedure

When upgrading to HDP 2.2.8 using Ambari, Spark 1.2.1 is automatically upgraded to 1.3.1. However, if you wish to return to using 1.2.1:

1. Remove Spark 1.3.1 from your HDP cluster using Ambari:

```
curl -u admin:admin -H "X-Requested-By: ambari" -X DELETE
```

```
http://<AMBARI_HOST>:8080/api/v1/clusters/<CLUSTER_NAME>/services/SPARK
```

2. Manually install Spark 1.2.1 with [HDP 2.2.6 Installing HDP Manually: Installing and Configuring Apache Spark](#).

1.4. Behavioral Changes

In HDP 2.2.8, behavioral changes affect the following Hadoop components:

Component	Summary	Details
HDFS	HDFS-8549 : Abort the balancer if an upgrade is in progress	As of HDP 2.2.8, the balancer cannot run during upgrade. Running the balancer during an ongoing upgrade has a negative affect, since DNs do not actually delete blocks. This means the balancer is making lots of extra replicas and not actually reducing the disk utilization of over-utilized nodes.
Slider	SLIDER-931 : Security permissions on set up ZK path are too lax	Zookeeper paths created for applications such as HBase are created with stricter permissions.
MapReduce	MAPREDUCE-6454 : A hive query run from a hive client running from a mapreduce job doesn't pick up the classpath of the hive client	Added jars in distributed cache to \$HADOOP_CLASSPATH (previously only added to CLASSPATH). Now, when Hive spawns a local child task, it can add them on classpath automatically (as Hadoop shell only takes HADOOP_CLASSPATH env).
Hue	Hue now shows log messages for create and insert statements.	N/A
Hue	Support dfs.umaskmode and fs.permissions.umask-mode when creating files or directories	When creating files or directories in File Browser, Hue does not respect fs.permissions.umask-mode and defaults always to 022. With this change now it is possible to set umask property in Hue configuration file with an octal value that will be used as default one for files and directories creation.
YARN	YARN-3626 : If we set mapreduce.job.user.classpath.first as true, the localized resources should appear before system classpath. But	In response to the mapreduce.job.user.classpath.first setting the classpath is ordered differently so that localized resources will appear before system classpath

Component	Summary	Details
	in windows, it does not work. We have fixed this in YARN-3626 .	resources when tasks execute. On Windows this does not work because the localized resources are not linked into their final location when the classpath jar is created. To compensate for that localized jar resources are added directly to the classpath generated for the jar rather than being discovered from the localized directories. Unfortunately, they are always appended to the classpath, and so are never preferred over system resources. Workaround: To resolve this, the situation should be detected and, when applicable, localized resources should be put at the beginning of the classpath rather than the end.
YARN	YARN-3978 : Users of the Application Timeline Service can now turn off collection of detailed container-level metrics by the use of a new configuration parameter <code>yarn.generic-application-history.save-non-am-container-meta-info</code> and setting it to false.	The default value of <code>yarn.generic-application-history.save-non-am-container-meta-info</code> is true, so there is no change in behavior. For clusters with more than 100 nodes, we recommend this configuration value be set to false to reduce the load on the Application Timeline Service.

1.5. Apache Patch Information

This section describes Apache JIRAs that were addressed in HDP 2.2.8.

1.5.1. Accumulo 1.6.1

HDP 2.2.8 provides Accumulo 1.6.1 and the following Apache patches:

- [ACCUMULO-3312](#) Fix incorrect path generation during clone table.
- [ACCUMULO-3600](#) Fix deleterows to properly compute file counts.

1.5.2. Falcon 0.6.0

HDP 2.2.8 provides Apache Falcon 0.6.0 and the following additional Apache patches:

- [FALCON-1005](#): In `ProcessInstanceSuspendTest` clean test dir after each method execution
- [FALCON-1007](#): Improve test output and variable names.
- [FALCON-1072](#): Dumping of Oozie info should use os specific newline

1.5.3. Hadoop Common/HDFS 2.6.0

HDP 2.2.8 provides Apache Hadoop Core 2.6.0 and the following additional Apache patches:

- [HADOOP-11321](#): `copyToLocal` cannot save a file to an SMB share unless the user has Full Control permissions.

- [HADOOP-11368](#): Fix SSLFactory truststore reloader thread leak in KMSClientProvider.
- [HADOOP-11381](#): Fix findbugs warnings in hadoop-distcp, hadoop-aws, hadoop-azure, and hadoop-openstack
- [HADOOP-11381](#): Revert. (Fix findbugs warnings in hadoop-distcp, hadoop-aws, hadoop-azure, and hadoop-openstack.)
- [HADOOP-11412](#): POMs mention "The Apache Software License" rather than "Apache License".
- [HADOOP-11412](#): Revert. (POMs mention "The Apache Software License" rather than "Apache License")
- [HADOOP-11490](#): Expose truncate API via FileSystem and shell command.
- [HADOOP-11509](#): change parsing sequence in GenericOptionsParser to parse -D parameters before -files.
- [HADOOP-11510](#): Expose truncate API via FileContext.
- [HADOOP-11523](#): StorageException complaining " no lease ID" when updating FolderLastModifiedTime in WASB.
- [HADOOP-11579](#): Documentation for truncate.
- [HADOOP-11595](#): Add default implementation for AbstractFileSystem#truncate.
- [HADOOP-926](#): Do not fail job history iteration when encountering missing directories.
- [HADOOP-941](#): Addendum patch.
- [HDFS-3107](#): Introduce truncate.
- [HDFS-7009](#): Active NN and standby NN have different live nodes.
- [HDFS-7056](#): Snapshot support for truncate.
- [HDFS-7058](#): Tests for truncate CLI
- [HDFS-7263](#): Snapshot read can reveal future bytes for appended files.
- [HDFS-7425](#): NameNode block deletion logging uses incorrect appender.
- [HDFS-7443](#): Datanode upgrade to BLOCKID_BASED_LAYOUT fails if duplicate block files are present in the same volume
- [HDFS-7470](#): SecondaryNameNode need twice memory when calling reloadFromImageFile.
- [HDFS-7489](#): Incorrect locking in FsVolumeList#checkDirs can hang datanodes.
- [HDFS-7503](#): Namenode restart after large deletions can cause slow processReport
- [HDFS-7606](#): Fix potential NPE in INodeFile.getBlocks().

- [HDFS-7634](#): Disallow truncation of Lazy persist files.
- [HDFS-7638](#): Small fix and few refinements for FSN#truncate.
- [HDFS-7643](#): Test case to ensure lazy persist files cannot be truncated.
- [HDFS-7655](#): Expose truncate API for Web HDFS.
- [HDFS-7659](#): Truncate should check negative value of the new length
- [HDFS-7676](#): Fix TestFileTruncate to avoid bug of HDFS-7611.
- [HDFS-7677](#): DistributedFileSystem#truncate should resolve symlinks.
- [HDFS-7707](#): Edit log corruption due to delayed block removal again.
- [HDFS-7714](#): Simultaneous restart of HA NameNodes and DataNode can cause DataNode to register successfully with only one NameNode.
- [HDFS-7733](#): NFS: readdir/readdirplus return null directory attribute on failure.
- [HDFS-7738](#): Revise the exception message for recover lease; add more truncate tests such as truncate with HA setup, negative tests, truncate with other operations and multiple truncates.
- [HDFS-7760](#): Document truncate for WebHDFS.
- [HDFS-7831](#): Fix the starting index and end condition of the loop in FileDiffList.findEarlierSnapshotBlocks().
- [HDFS-7843](#): A truncated file is corrupted after rollback from a rolling upgrade.
- [HDFS-7885](#): Datanode should not trust the generation stamp provided by client.
- [MAPREDUCE-6230](#): Fixed RMContainerAllocator to update the new AMRMToken service name properly.
- [YARN-2246](#): Made the proxy tracking URL always be http(s)://proxy addr:port/proxy/<appId> to avoid duplicate sections.
- [YARN-2571](#): RM to support YARN registry
- [YARN-2683](#): registry config options: document and move to core-default
- [YARN-2837](#): Support Timeline server to recover delegation token when restarting.
- [YARN-2917](#): Fixed potential deadlock when system.exit is called in AsyncDispatcher
- [YARN-2964](#): RM prematurely cancels tokens for jobs that submit jobs (oozie).
- [YARN-3103](#): AMRMClientImpl does not update AMRM token properly.
- [YARN-3207](#): Secondary filter matches entites which do not have the key being filtered for.

- [YARN-3227](#): Timeline renew delegation token fails when RM user's TGT is expired.
- [YARN-3239](#): WebAppProxy does not support a final tracking url which has query fragments and params.
- [YARN-3251](#): Fixed a deadlock in CapacityScheduler when computing absoluteMaxAvailableCapacity in LeafQueue.
- [YARN-3269](#): Yarn.nodemanager.remote-app-log-dir could not be configured to fully qualified path.
- [YARN-570](#): Time strings are formatted in different timezone.

1.5.4. HBase 0.98.4

HDP 2.2.8 provides HBase 0.98.4 and the following additional Apache patches:

- [HBASE-10499](#): In write-heavy scenario one of the regions does not get flushed causing RegionTooBusyException
- [HBASE-11569](#): Addendum for not skipping replayed edits for primary region replica
- [HBASE-12238](#): A few exceptions on startup - PARTIAL BACKPORT
- [HBASE-12533](#): staging directories are not deleted after secure bulk load
- [HBASE-12536](#): Reduce the effective scope of GLOBAL CREATE and ADMIN permission
- [HBASE-12562](#): Handling memory pressure for secondary region replicas - ADDENDUM for fixing findbug reported issues
- [HBASE-12575](#): Sanity check table coprocessor classes are loadable
- [HBASE-12714](#): RegionReplicaReplicationEndpoint should not set the RPC Codec
- [HBASE-12791](#): HBase does not attempt to clean up an aborted split when the regionserver shutting down
- [HBASE-12958](#): SSH doing hbase:meta get but hbase:meta not assigned
- [HBASE-13120](#): Allow disabling hadoop classpath and native library lookup
- [HBASE-212](#): Addendum. Fixes a unit test

1.5.5. Hive 0.14.0

HDP 2.2.8 provides Hive 0.14.0 and the following additional Apache patches:

- [HIVE-9984](#): JoinReorder's getOutputSize is exponential
- [HIVE-9975](#): Renaming a nonexisting partition should not throw out NullPointerException
- [HIVE-9953](#): fix NPE in WindowingTableFunction

- [HIVE-9952](#): fix NPE in CorrelationUtilities
- [HIVE-9950](#): fix rehash in CuckooSetBytes and CuckooSetLong
- [HIVE-9936](#): fix potential NPE in DefaultUDAFEvaluatorResolver
- [HIVE-9755](#): Hive built-in "ngram" UDAF fails when a mapper has no matches.
- [HIVE-9720](#): Metastore does not properly migrate column stats when renaming a table across databases
- [HIVE-9716](#): Map job fails when table's LOCATION does not have scheme
- [HIVE-9655](#): Dynamic partition table insertion error
- [HIVE-9648](#): Null check key provider before doing set
- [HIVE-9622](#): Getting NPE when trying to restart HS2 when metastore is configured to use org.apache.hadoop.hive.thrift.DBTokenStore
- [HIVE-9619](#): Uninitialized read of numBitVectors in NumDistinctValueEstimator
- [HIVE-9592](#): fix ArrayIndexOutOfBoundsException in date_add and date_sub initialize
- [HIVE-9567](#): JSON SerDe not escaping special chars when writing char/varchar data
- [HIVE-9560](#): When hive.stats.collect.rawdatasize=true, 'rawDataSize' for an ORC table will result in value '0' after running 'analyze table TABLE_NAME compute statistics;'
- [HIVE-9526](#): ClassCastException thrown by HiveStatement
- [HIVE-9513](#): NULL POINTER EXCEPTION
- [HIVE-9509](#): Restore partition spec validation removed by HIVE-9445
- [HIVE-9507](#): Make "LATERAL VIEW inline(expression) mytable" tolerant to nulls
- [HIVE-9471](#): Bad seek in uncompressed ORC, at row-group boundary.
- [HIVE-9445](#): Revert HIVE-5700 - enforce single date format for partition column storage
- [HIVE-9437](#): Beeline does not add any existing HADOOP_CLASSPATH
- [HIVE-9397](#): SELECT max(bar) FROM foo is broken after ANALYZE ... FOR COLUMNS
- [HIVE-9377](#): UDF in_file() in WHERE predicate causes NPE.
- [HIVE-9347](#): Bug with max() together with rank() and grouping sets
- [HIVE-9322](#): Make null-checks consistent for MapObjectInspector subclasses.
- [HIVE-9309](#): schematool fails on Postgres 8.1
- [HIVE-9301](#): Potential null dereference in MoveTask#createTargetPath
- [HIVE-9228](#): Problem with subquery using windowing functions

- [HIVE-9217](#): UnionProcessor misses results for multi-insert when hive.optimize.union.remove=true
- [HIVE-9177](#): Fix child operator references after NonBlockingOpDeDupProc
- [HIVE-9155](#): HIVE_LOCKS uses int instead of bigint hive-txn-schema-0.14.0.mssql.sql
- [HIVE-9113](#): Explain on query failed with NPE
- [HIVE-9060](#): Fix child operator references after NonBlockingOpDeDupProc
- [HIVE-9024](#): NullPointerException when starting webhcat server if templeton.hive.properties is not set
- [HIVE-8926](#): Projections that only swap input columns are identified incorrectly as identity projections
- [HIVE-8889](#): JDBC Driver ResultSet.getXXXXXX(String columnLabel) methods Broken
- [HIVE-8874](#): Error Accessing HBase from Hive via Oozie on Kerberos 5.0.1 cluster
- [HIVE-8869](#): RowSchema not updated for some ops when columns are pruned
- [HIVE-8863](#): Cannot drop table with uppercase name after compute statistics for columns
- [HIVE-8784](#): Querying partition does not work with JDO enabled against PostgreSQL
- [HIVE-8706](#): Table statistic collection on counter failed due to table name character case.
- [HIVE-8680](#): Set Max Message for Binary Thrift endpoints
- [HIVE-8627](#): Compute stats on a table from impala caused the table to be corrupted
- [HIVE-8594](#): Wrong condition in SettableConfigUpdater#setHiveConfWhiteList()
- [HIVE-8518](#): Compile time skew join optimization returns duplicated results
- [HIVE-8518](#): Compile time skew join optimization returns duplicated results
- [HIVE-8450](#): Create table like does not copy over table properties
- [HIVE-8448](#): Union All might not work due to the type conversion issue
- [HIVE-8374](#): schematool fails on Postgres versions less than 9.2
- [HIVE-8330](#): HiveResultSet.findColumn() parameters are case sensitive
- [HIVE-8326](#): Using DbTxnManager with concurrency off results in run time error
- [HIVE-8295](#): Add batch retrieve partition objects for metastore direct sql
- [HIVE-8099](#): IN operator for partition column fails when the partition column type is DATE
- [HIVE-7997](#): Potential null pointer reference in ObjectInspectorUtils#compareTypes()
- [HIVE-7351](#): ANALYZE TABLE statement fails on postgres metastore

- [HIVE-6727](#): Table level stats for external tables are set incorrectly
- [HIVE-6308](#): COLUMNS_V2 Metastore table not populated for tables created without an explicit column list.
- [HIVE-5664](#): Drop cascade database fails when the db has any tables with indexes
- [HIVE-11613](#): schematool should return non zero exit status for info command, if state is inconsistent
- [HIVE-11493](#): Predicate with integer column equals double evaluates to false
- [HIVE-11442](#): Remove commons-configuration.jar from Hive distribution
- [HIVE-11438](#): Join a ACID table with non-ACID table fail with MR on 1.0.0
- [HIVE-11429](#): Increase default JDBC result set fetch size (# rows it fetches in one RPC call) to 1000 from 50
- [HIVE-11356](#): SMB join on tez fails when one of the tables is empty
- [HIVE-11303](#): Getting Tez LimitExceededException after dag execution on large query
- [HIVE-11285](#): ObjectInspector for partition columns in FetchOperator in SMBJoin causes exception
- [HIVE-11271](#): java.lang.IndexOutOfBoundsException when union all with if function
- [HIVE-11172](#): Vectorization wrong results for aggregate query with where clause without group by
- [HIVE-11147](#): MetaTool doesn't update FS root location for partitions with space in name
- [HIVE-11112](#): ISO-8859-1 text output has fragments of previous longer rows appended
- [HIVE-11104](#): Select operator doesn't propagate constants appearing in expressions
- [HIVE-11095](#): SerDeUtils another bug ,when Text is reused
- [HIVE-11035](#): PPD Orc Split elimination fails because filterColumns=[-1]
- [HIVE-11031](#): ORC concatenation of old files can fail while merging column statistics
- [HIVE-11028](#): Tez: table self join and join with another table fails with IndexOutOfBoundsException
- [HIVE-11027](#): Hive on tez Bucket map joins fail when hashcode goes negative
- [HIVE-11023](#): Disable directSQL if datanucleus.identifierFactory = datanucleus2
- [HIVE-11008](#): webhcat GET /jobs retries on getting job details from history server is too aggressive
- [HIVE-10992](#): WebHCat should not create delegation tokens when Kerberos is not enabled

- [HIVE-10963](#): Hive throws NPE rather than meaningful error message when window is missing
 - [HIVE-10929](#): In Tez mode,dynamic partitioning query with union all fails at moveTask,Invalid partition key and values
 - [HIVE-10841](#): [WHERE col is not null] does not work sometimes for queries with many JOIN statements
 - [HIVE-10816](#): NPE in ExecDriver handleSampling when submitted via child JVM
 - [HIVE-10808](#): Inner join on Null throwing Cast Exception
 - [HIVE-10722](#): external table creation with msck in Hive can create unusable partition
 - [HIVE-10719](#): Hive metastore failure when alter table rename is attempted.
 - [HIVE-10698](#): query on view results fails with table not found error if view is created with subquery alias (CTE).
 - [HIVE-10685](#): Alter table concatenate oparetor will cause duplicate data
 - [HIVE-10607](#): Combination of ReducesinkDedup + TopN optimization yields incorrect result if there are multiple GBY in reducer
 - [HIVE-10595](#): Dropping a table can cause NPEs in the compactor
 - [HIVE-10578](#): update sql standard authorization configuration whitelist
 - [HIVE-10538](#): Fix NPE in FileSinkOperator from hashcode mismatch
 - [HIVE-10450](#): More than one TableScan in MapWork not supported in Vectorization – causes query to fail during vectorization
 - [HIVE-10384](#): RetryingMetaStoreClient does not retry wrapped TTransportExceptions
 - [HIVE-10303](#): HIVE-9471 broke forward compatibility of ORC files
 - [HIVE-10273](#): Union with partition tables which have no data fails with NPE
 - [HIVE-10183](#): [CBO] self-join failing in a test case
 - [HIVE-10140](#): Window boundary is not compared correctly
 - [HIVE-10122](#): Hive metastore filter-by-expression is broken for non-partition expressions
 - [HIVE-10062](#): : HiveOnTez Union followed by Multi-GB followed by Multi-insert loses data
- HDP 2.2.6 provided Hive 0.14.0 and the following additional Apache patches:
- [HIVE-480](#): HDFSCleanup thread holds reference to FileSystem
 - [HIVE-6468](#): HiveServer2 (tcp mode, with SASL layer) OOMs when getting a non-thrift message

- [HIVE-6679](#): HiveServer2 should support TCP Keepalive & Server Socket Timeout on blocking calls
- [HIVE-7175](#): Provide password file option to beeline
- [HIVE-7270](#): Serde info is not shown in show create table statement, but shows in the desc table
- [HIVE-8295](#): Add batch retrieve partition objects for metastore direct sql
- [HIVE-8485](#): SUMMARY-Hive metastore NPE with Oracle DB when there is empty value for string for tblproperties/serdeproperties/etc, table not usable after creation
- [HIVE-8762](#): HiveMetaStore.BooleanPointer should be replaced with an AtomicBoolean
- [HIVE-8791](#) : Hive permission inheritance throws exception S3
- [HIVE-8850](#): SUMMARY-[ObjectStore:: rollbackTransaction()] needs to be looked into further
- [HIVE-8881](#): Receiving json "could not find job" error when web client tries to fetch all jobs from WebHCat but HDFS does not have the data.
- [HIVE-8888](#): Hive on Tez query output duplicate rows when there is explode in subqueries for joins
- [HIVE-8891](#): SUMMARY-[Another possible cause to NucleusObjectNotFoundException from drops/rollback
- [HIVE-8893](#): Implement whitelist for builtin UDFs to avoid untrusted code execution in multiuser mode
- [HIVE-8966](#): Delta files created by hive hcatalog streaming cannot be compacted
- [HIVE-9025](#): join38.q (without map join) produces incorrect result when testing with multiple reducers
- [HIVE-9038](#): Join tests fail on Tez
- [HIVE-9055](#): Tez: union all followed by group by followed by another union all gives error
- [HIVE-9106](#): improve the performance of null scan optimizer when several table scans share a physical path
- [HIVE-9112](#): Query may generate different results depending on the number of reducers
- [HIVE-9141](#): HiveOnTez: mix of union all, distinct, group by generates error
- [HIVE-9155](#): HIVE_LOCKS uses int instead of bigint hive-txn-schema-0.14.0.mssql.sql
- [HIVE-9205](#): Change default Tez install directory to use /tmp instead of /user and create the directory if it does not exist
- [HIVE-9234](#): HiveServer2 leaks FileSystem objects in FileSystem.CACHE

- [HIVE-9235](#): Turn off Parquet Vectorization until all data types work: DECIMAL, DATE, TIMESTAMP, CHAR, and VARCHAR
- [HIVE-9249](#): java.lang.ClassCastException: org.apache.hadoop.hive.serde2.io.HiveVarcharWritable cannot be cast to org.apache.hadoop.hive.common.type.HiveVarchar when joining tables
- [HIVE-9278](#): Cached expression feature broken in one case
- [HIVE-9316](#): TestSqoop tests in WebHCat testsuite hardcode libdir path to hdfs
- [HIVE-9351](#): Running Hive Jobs with Tez cause templetion to never report percent complete
- [HIVE-9359](#): SUMMARY-[Hive export OOM error when table is huge. (32TB data, 4800+ partitions)]
- [HIVE-9382](#) : Query got rerun with Global Limit optimization on and Fetch optimization off
- [HIVE-9390](#): Enhance retry logic wrt DB access in TxnHandler
- [HIVE-9401](#): SimpleFetchOptimizer for limited fetches without filters
- [HIVE-9404](#) NPE in org.apache.hadoop.hive.metastore.txn.TxnHandler.determineDatabaseProduct()
- [HIVE-9436](#): SUMMARY-[RetryingMetaStoreClient does not retry JDOExceptions]
- [HIVE-9446](#): JDBC DatabaseMetadata.getColumns() does not work for temporary tables
- [HIVE-9473](#) : sql std auth should disallow built-in udfs that allow any java methods to be called
- [HIVE-9593](#): ORC reader should ignore new/unknown metadata streams.
- [HIVE-9652](#): STDERR redirection should not use in place updates TEZ UI
- [HIVE-9665](#): Parallel move task optimization causes race condition
- [HIVE-9673](#) : Set operationhandle in ATS entities for lookups
- [HIVE-9683](#): Client TCP keep-alive for Thrift as a workaround for THRIFT-2788
- [HIVE-9684](#): Incorrect disk range computation in ORC because of optional stream kind
- [HIVE-9743](#) Incorrect result set for vectorized left outer join
- [HIVE-9779](#) : ATSHook does not log the end user if doAs=false (it logs the hs2 server user)
- [HIVE-9832](#): Merge join followed by union and a map join in hive on tez fails.
- [HIVE-9836](#): Hive on Tez: fails when virtual columns are present in the join conditions (for e.g. partition columns)
- [HIVE-9841](#): IOException thrown by ORC should include the path of processing file

- [HIVE-9886](#): Tez NPE error when initialize reducer from 2 row_number function on each join side
- [HIVE-9892](#): HIVE schema failed to upgrade from HDP 2.0/2.1 to HDP 2.2 with schematool

1.5.6. Knox 0.5.0

HDP 2.2.8 provides Knox 0.5.0 and the following additional Apache patches:

- [KNOX-492](#): Support service level replayBufferLimit for Oozie, Hive and HBase.

1.5.7. Oozie 4.1.0

HDP 2.2.8 provides Apache Oozie 4.1.0 and the following additional Apache patch:

- [OOZIE-208](#): Adding missing oozie property oozie.service.HadoopAccessorService.hadoop.configurations to oozie install script.

1.5.8. Phoenix 4.2.0

HDP 2.2.8 includes Apache Phoenix 4.2.0 and the following additional Apache patches:

- [PHOENIX-1248](#): CsvBulkLoadTool is failing with IAE when local index specified for –index-table parameter
- [PHOENIX-1336](#): Exception when select from local index:Cache of region boundaries are out of date
- [PHOENIX-1346](#): the choice of index by phoenix is not correct where do querying
- [PHOENIX-1362](#): Min/max aggregate query on CHAR and BINARY types always return null
- [PHOENIX-1427](#): Reduce work in StatsCollector
- [PHOENIX-1428](#): Queries with limit against tenant views is broken
- [PHOENIX-1429](#): Cancel queued threads when limit reached
- [PHOENIX-1431](#): DELETE using Subqueries
- [PHOENIX-1436](#) : The choice of index by phoenix is not correct where do querying
- [PHOENIX-1446](#): Add/fix logging for LIMIT optimization
- [PHOENIX-1447](#): Increase guidepost width default to take into account FAST_DIFF compression
- [PHOENIX-1448](#): Fix resource leak when work rejected by thread executor
- [PHOENIX-1449](#): Fix PropertiesUtil.deepCopy()
- [PHOENIX-1455](#): Replace org.xerial.snappy with org.iq80.snappy pure Java snappy implementation

- [PHOENIX-1456](#): Incorrect query results caused by reusing buffers in SpoolingResultIterator
- [PHOENIX-1462](#): Create unit test for COUNT DISTINCT using compression (Ram)
- [PHOENIX-1463](#): phoenix.query.timeoutMs doesn't work as expected
- [PHOENIX-1466](#): Prevent multiple scans when query run serially
- [PHOENIX-1467](#): Upgrade to 4.12 Junit and update tests by removing @Category annotation
- [PHOENIX-1468](#): Add org.iq80.snappy to client and server assemblies
- [PHOENIX-1469](#): Binary columns do not work correctly for indexing
- [PHOENIX-1472](#): Fix salt bucket for Integer.MIN_VALUE
- [PHOENIX-1473](#): Connecting with Phoenix client when Phoenix is not deployed on region server(s) takes down region server(s)
- [PHOENIX-1474](#): NPE when RVC between combined with key part comparison
- [PHOENIX-1476](#): Immediately open scanner for parallel queries
- [PHOENIX-1477](#): Phoenix 4.0 -> 4.2 upgrade doesn't work
- [PHOENIX-1480](#): Incorrect query results may occur when VIEW uses indexes from physical table
- [PHOENIX-1484](#): Index creation failed due to specifying DEFAULT_COLUMN_FAMILY option
- [PHOENIX-1485](#): Add timezone awareness
- [PHOENIX-1485](#): Correct arg count in FunctionParseNode
- [PHOENIX-1486](#): Fix SpooledTmpFileDeleteIT by having it use a different directory for spooling files
- [PHOENIX-1498](#): Turn KEEP_DELETED_CELLS off by default
- [PHOENIX-1500](#): Disallow mutations and queries using using PhoenixPreparedStatement.executeQuery() and executeUpdate() respectively
- [PHOENIX-1503](#): - Compilation on Mac OS fails on maven-assembly-plugin
- [PHOENIX-1532](#): Phoenix tarball assembly does not include server jar
- [PHOENIX-1533](#): Last key part not taken into child/parent optimization
- [PHOENIX-1537](#): Set reuseForks to false for integration tests
- [PHOENIX-1551](#): Upgrading from Phoenix 4.0.x to 4.2.2 and throw can't find SYSTEM.STATS

- [PHOENIX-1571](#): Replace hard-coded date
- [PHOENIX-1575](#): Identical sql type values for UNSIGNED_TIMESTAMP and UNSIGNED_DATE
- [PHOENIX-1592](#): ORDER BY not closing ResultIterator properly
- [PHOENIX-1596](#): Turning tracing on causes region servers to crash
- [PHOENIX-1601](#): Performance.py script not working in binary distribution
- [PHOENIX-18](#): Several tests are failing with SequenceNotFoundException error
- [PHOENIX-32](#): Some phoenix tests are failing with different assertion errors

1.5.9. Pig 0.14.0

HDP 2.2.8 provides Apache Pig 0.14.0 and the following additional Apache patches:

- [PIG-156](#): Pig command fails because the input line is too long
- [PIG-4334](#): PigProcessor does not set pig.datetime.default.tz
- [PIG-4342](#): Pig 0.14 cannot identify the uppercase of DECLARE and DEFAULT
- [PIG-4377](#): Skewed outer join produce wrong result in some cases
- [PIG-4377](#): Skewed outer join produce wrong result in some cases (PIG-4377-2.patch)

1.5.10. Ranger 0.4.0

HDP 2.2.8 provides Ranger 0.4.0 and the following additional Apache patch:

- [RANGER-188](#): Added LSB headers to Ranger Admin/Usersync init.d scripts

1.5.11. Slider 0.61.0

HDP 2.2.8 provides Slider 0.61.0 and the following additional Apache patches:

- [SLIDER-769](#): Modify Slider app packages for seamless integration with AMS provider (Ambari Metrics Service)

1.5.12. Storm 0.9.3

HDP 2.2.8 provides Apache Storm 0.9.3 and the following additional patches:

- [STORM-586](#): TridentKafkaEmitter should catch updateOffsetException.
- [STORM-682](#): Supervisor should handle worker state corruption gracefully.

1.5.13. Tez 0.5.2

HDP 2.2.8 provides Apache Tez 0.5.2 and the following Apache patches:

- [TEZ-1642](#): TestAMRecovery sometimes fails.
- [TEZ-1775](#): Allow setting log level per logger.
- [TEZ-1800](#): Integer overflow in ExternalSorter.getInitialMemoryRequirement()
- [TEZ-1836](#): Provide better error messages when tez.runtime.io.sort.mb, spill percentage is incorrectly configured.
- [TEZ-1851](#): FileSystem counters do not differentiate between different FileSystems
- [TEZ-1852](#): Get examples to work in Local Mode.
- [TEZ-1861](#): Fix failing test: TestOnFileSortedOutput.
- [TEZ-1878](#): Task-specific log level override not working in certain conditions
- [TEZ-1924](#): Tez AM does not register with AM with full FQDN causing jobs to fail in some environments.
- [TEZ-1931](#): Publish tez version info to Timeline.
- [TEZ-1934](#): TestAMRecovery may fail due to the execution order is not determined.
- [TEZ-1942](#): Number of tasks show in Tez UI with auto-reduce parallelism is misleading.
- [TEZ-1949](#): Whitelist TEZ_RUNTIME_OPTIMIZE_SHARED_FETCH for broadcast edges
- [TEZ-1962](#): Fix a thread leak in LocalMode.
- [TEZ-2024](#): Compilation error due to conflict.
- [TEZ-2024](#): TaskFinishedEvent may not be logged in recovery.
- [TEZ-2037](#): Should log TaskAttemptFinishedEvent if taskattempt is recovered to KILLED
- [TEZ-2135](#): ACL checks handled incorrectly in AMWebController.

As part of HDP 2.2.8 Hortonworks is providing a Tez Debugging User Interface. This interface does not impact the behavior or function of jobs that leverage Tez, and its use is optional. Patches added to facilitate the Tez Debugging User Interface include:

- [TEZ-1990](#): Tez UI: DAG details page shows Nan for end time when a DAG is running.
- [TEZ-2031](#): Tez UI: horizontal scrollbars do not appear in tables, causing them to look truncated.
- [TEZ-2038](#): TEZ-UI DAG is always running in tez-ui when the app is failed but no DAGFinishedEvent is logged.
- [TEZ-2043](#): Tez UI: add progress info from am webservice to dag and vertex views.
- [TEZ-2052](#): Tez UI: log view fixes, show version from build, better handling of ats url config.

- [TEZ-2056](#): Tez UI: fix VertexID filter, show only tez configs by default, fix appattemptid.
- [TEZ-2063](#): Tez UI: Flaky log url in tasks table.
- [TEZ-2065](#): Setting up tez.tez-ui.history-url.base with a trailing slash can result in failures to redirect correctly.
- [TEZ-2068](#): Tez UI: Dag view should use full window height, disable webuiservice in localmode.
- [TEZ-2069](#): Tez UI: appld should link to application in dag details view.
- [TEZ-2077](#): Tez UI: No diagnostics on Task Attempt Details page if task attempt failed
- [TEZ-2078](#): Tez UI: Task logs url use in-progress url causing various errors.
- [TEZ-2079](#): Tez UI: trailing slash in timelineBaseUrl in ui should be handled.
- [TEZ-2092](#): Tez UI history url handler injects spurious trailing slash.
- [TEZ-2098](#): Tez UI: Dag details should be the default page for dag, fix invalid time entries for failed Vertices.
- [TEZ-2101](#): Tez UI: Issues on displaying a table.
- [TEZ-2102](#): Tez UI: DAG view has hidden edges, dragging DAG by holding vertex causes unintended click.
- [TEZ-2106](#): TEZ UI: Display data load time, and add a refresh button for items that can be refreshed.
- [TEZ-2112](#): Tez UI: fix offset calculation, add home button to breadcrumbs.
- [TEZ-2114](#): Tez UI: task/task attempt status is not available when its running.
- [TEZ-2116](#): Tez UI: dags page filter does not work if more than one filter is specified.
- [TEZ-2134](#): TEZ UI: On request failure, display request URL and server name in error bar.
- [TEZ-2136](#): Some enhancements to the new Tez UI.
- [TEZ-2142](#): TEZ UI: Breadcrumb border color looks out of place in wrapped mode.
- [TEZ-2158](#): TEZ UI: Display dag/vertex names, and task/attempt index in breadcrumb
- [TEZ-2160](#): Tez UI: App tracking URL should support navigation back.
- [TEZ-2165](#): Tez UI: DAG shows running status if killed by RM in some cases.

1.5.14. ZooKeeper 3.4.6

HDP 2.2.8 provides ZooKeeper 3.4.6 and the following additional Apache patches:

- [ZOOKEEPER-1506](#): Re-try DNS hostname -> IP resolution if node connection fails

1.6. Common Vulnerabilities and Exposures

HDP 2.2.8 provides fixes for the following information-security vulnerabilities and exposures (CVEs):

- **CVE-2014-0227:** Apache - Tomcat - HTTP Request Smuggling Issue

Severity: Important

Vendor: Hortonworks

Versions Affected: All HDP 2.2.x releases prior to 2.2.8

Users Affected: Users who use HttpFS, Hadoop KMS, Hadoop Auth, and Oozie.

Impact: See [BUG-41101](#) and [HADOOP-12232](#). Allows unauthorized modification; Allows disruption of service. `java/org/apache/coyote/http11/filters/ChunkedInputFilter.java` in Apache Tomcat 6.x before 6.0.42, 7.x before 7.0.55, and 8.x before 8.0.9 does not properly handle attempts to continue reading data after an error has occurred, which allows remote attackers to conduct HTTP request smuggling attacks or cause a denial of service (resource consumption) by streaming data with malformed chunked transfer coding.

Recommended Action: Upgrade to HDP 2.2.8.

- **CVE-2015-5167:** Restrict REST API data access for non-admin users

Severity: Important

Vendor: Hortonworks

Versions Affected: All HDP 2.2.x releases prior to 2.2.8

Users Affected: All users of Ranger Policy Admin tool.

Impact: See [BUG-41603](#) and [RANGER-630](#). Data access restrictions via REST API are not consistent with restrictions in policy admin UI. Non-admin users can access some ranger data restricted for admin users by calling REST API.

Recommended Action: Upgrade to HDP 2.2.8+.

- **CVE-2015-1836:** Apache HBase remote denial of service, information integrity, and information disclosure vulnerability

Severity: Important

Vendor: Hortonworks

Versions Affected: All HDP 2.2.x releases prior to 2.2.8

Users Affected: Users who use HBase in secure environments are affected by a logic error caused HBase in most secure configuration deployments to handle its coordination state in ZooKeeper via insecure ACLs. Providing inappropriate access by any user with authenticated access to the network.

Impact: See [BUG-38465](#) and [HBASE-13768](#). Inappropriate access by any user with authenticated access to the network.

Recommended Action: HBase users should update to the latest hotfix release of their respective version (e.g. HDP-2.2.8.0, HDP-2.3.0.0) to ensure newly written coordination information has the correct ACLs.

1.7. Third-party Licenses

Component	Subcomponents	License
Accumulo	JCommander	JCommander
Falcon	cern.colt*, cern.jet*, cern.dhep	CERN
Knox	ApacheDS, Groovy	ANTLR
Knox	SL4J	MIT
Knox	Jetty and Jerico	EPL
Knox	ApacheDS	Bouncy Castle
Phoenix		EPL
Storm	Logback	EPL

1.8. Fixed Issues

Fixed issues represents selected issues that were previously logged via Hortonworks Support, but are now addressed in the current release. These issues may have been reported in previous versions within the Known Issues section; meaning they were reported by customers or identified by Hortonworks Quality Engineering team.

Incorrect Results

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-36095	HIVE-10607	Hive	Combination of ReducesinkDedup + TopN optimization yields incorrect result if there are multiple GBY in reducer
BUG-36685		Hive	(2.2) - Need fix for HDP 2.2.4.2-2 - select operator missing column, wrong results in multiple left outer join with sub queries
BUG-38406	PIG-4541	Pig	(2.2) - Skewed full outer join does not return records if any relation is empty. Outer join does not return any record if left relation is empty
BUG-38852	YARN-3764	YARN	CapacityScheduler should forbid moving LeafQueue from one parent to another
BUG-39519		Ranger	UserSync Process didn't sync the group when groups are added to the user at later time (2.2.4.2 backport)

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-39832	HIVE-11027	Hive	[2.2] Hive Incorrect query results with Bucket Map joins vs Shuffle Joins.
BUG-40282	HIVE-11104	Hive	(2.2, patch) - Insert overwrite query does not return expected value
BUG-41821	HIVE-10303 HIVE-9471	Hive	Bad seek in uncompressed ORC, at row-group boundary
BUG-42173		Hive	query wrong results, filter dropped in multi-way joins that contain join on same table twice+
BUG-42695	HBASE-12865	HBase	Back Port HBASE-12865 : WALs may be deleted before they are replicated to peers
BUG-42891	HIVE-11493	Hive	Backport HIVE-11493 (Predicate with integer column equals double evaluates to false)

Other

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-37634		RelEng	(2.2.4.10) - hdp-select script trips up on 'lost+found' directory
BUG-39010		RelEng	2.2-maint apt-get install fails for Zookeeper, Knox, Hbase due to duplicate files in deb's
BUG-39972		RelEng	https is broken in HDP 2.2.6.0
BUG-41480	YARN-3832	YARN	Resource Localization fails on a cluster due to existing cache directories
BUG-42123		Storm	Nimbus Client Misconfiguration
BUG-42616		Kafka	kafka-run-class.sh should have client_jmx_port and broker_jmx_port

Performance

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-37834	HADOOP-10597	HDFS	RPC Server signals backoff to clients when all request queues are full
BUG-38033	HIVE-9984	Hive	joinReorder getOutputSize exponential
BUG-38150	HDFS-7398 HDFS-7609	HDFS	Improve NameNode EditLog load time by fixing retry cache collision
BUG-38893	ACCUMULO-3890	Accumulo, Slider	Use of CredentialProvider results in a lot of NameNode operations

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-39833	HIVE-11035	Hive	[2.2] PPD: Orc Split elimination fails because filterColumns=[-1]
BUG-41358		Ranger	(2.2) - UserSync process performance issue due to user group cache inconsistency.
BUG-42199	HIVE-11429	Hive	(2.2) - Increase default JDBC result set fetch size (# rows it fetches in one RPC call) to 1000 from 50
BUG-42210		Hive	Replace use of distcp with FileUtils.copy(), during Hive.renameFile()
BUG-42284	HDFS-8856	HDFS	Improve HDFS Lease Manager performance
BUG-42364	HIVE-6727	Hive	table fast stats numFiles, totalSize set to 0 when table has non-default location
BUG-42810	YARN-3641	YARN	NodeManager: stopRecoveryStore() shouldn't be skipped when exceptions happen in stopping NM's sub-services.
BUG-43083	TEZ-2628	Tez	[ATS V1.5] Backport initial patches from TEZ-2628 (ATS v1.5 tez plugin) into 2.2-maint

Potential Data Loss

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-39580	HIVE-10685	Hive	[2.2] Backport HIVE-10685 : Alter ORC table concatenate operator can cause duplicate/corrupt data
BUG-41424	FLUME-2754	Flume	Flume Hive Sink skipping first transaction in each TxnBatch
BUG-42323	STORM-960	Storm	While using a storm/hive/bolt to populate an ACID ORC table, it causes data loss and scan jobs fails with 'java.io.EOFException' error
BUG-42680	HBASE-14054	HBase	Back port HBASE-14054 : Acknowledged writes may get lost if regionserver clock is set backwards

Query Failure

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-36758	HIVE-10273	Hive	Union with partition tables which have no data fails with NPE
BUG-38986	HIVE-11028	Hive	table self join and join with another table fails with IndexOutOfBoundsException

Hortonworks Bug ID	Apache JIRA	Component	Summary
			constant propagation wrong plan
BUG-38996		Hive	(2.2) - In Tez mode, dynamic partitioning query with union all fails at moveTask, Invalid partition key & values
BUG-39431	HIVE-9950	Hive	(2.2) - Hive query using Cuckoo hashing causing NullPointerException
BUG-39432		Hive	(2.2) - Hive Alter table, concatenate partition files, throws error
BUG-40150	HIVE-10538	Hive	[2.2] Fix NPE in FileSinkOperator from hashcode mismatch (HIVE-10538 backport)
BUG-40704	HIVE-11172	Hive	Vectorization NPE error with aggregation with sub function on query with where clause
BUG-40858		Hive	(2.2-maint) vectorization with ACID table, count(1) fails with ArrayIndexOutOfBoundsException
BUG-41237	HIVE-11285	Hive	In MR mode, SMB join is failing with ClassCastException when a partition key of type CHAR / INT / BIGINT / DECIMAL is read.
BUG-41439		Hive, Tez	Error thrown when Empty tables and SMB (OUTER) JOINS are used
BUG-41663	HIVE-11438	Hive	hive cli starts in MR mode/ CombineHiveInputFormat, query of ACID table join/ union non-ACID fails with NPE map plan not found error.

Security

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-36065		Ranger	PRIVACY case issue in HBase Ranger plugin
BUG-38608		Hue	(2.2) - Hue does not respect Support dfs.umaskmode, fs.permissions.umask-mode when creating files or folders
BUG-41603		Ranger	(2.2) Data consistency across API and UI
BUG-41975		Ranger	(2.2) Ranger enable plugin script should handle file permission with umask 027
BUG-42587	STORM-857	Storm	Supervisor process fails to write log metadata to YAML file when

Hortonworks Bug ID	Apache JIRA	Component	Summary
			"supervisor.run.worker.as.user" is enabled

Stability

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-27650	HBASE-12465	HBase	HBase master start fails due to incorrect file creations
BUG-30184		Hive	(2.2, Patch) - empty ORC file without schema produced by Hive
BUG-31699	HBASE-12070	HBase	Back Port HBASE-12070 : Add an option to hbck to fix ZK inconsistencies
BUG-36081	HIVE-10595	Hive	(2.2) - Deleted tables stalls all HiveMetastore compactions
BUG-36797	YARN-3626	YARN	On Windows localized resources are not moved to the front of the classpath when they should be merge to 2.2-maint
BUG-37096	HADOOP-11328	Hadoop Common	ZKFailoverController does not log Exception when doRun raises errors
BUG-37157	HBASE-13331	HBase	(2.2) - Exceptions from DFS client can cause CatalogJanitor to delete referenced files
BUG-37751	STORM-329 STORM-404	Storm	Worker on one machine crashes due to a failure of another worker on another machine
BUG-37760	PHOENIX-2031	Phoenix	Unable to process timestamp data loaded via Phoenix org.apache.phoenix.pig.PhoenixHBaseLoader
BUG-37915	TEZ-2488	Tez	(2.2) - [TEZ-2488] Tez AM crashes if a submitted DAG is configured to use invalid resource sizes.
BUG-39400		Hive	hive tez dynamic partition pruning hang on to-be-pruned Map Vertex initialization, when tez reuse containers
BUG-39565	HDFS-7871	HDFS	NameNode edit log roller error message printing can cause NameNode to freeze
BUG-39649	TEZ-2189	Tez	[TEZ-2189] Tez UI live AM tracking url only works for localhost addresses
BUG-39773	HIVE-11023	Hive	(2.2) - Disable direct sql if datanucleus.identifierFactory = datanucleus2
BUG-39803	YARN-2922	YARN	ConcurrentModificationException in CapacityScheduler's LeafQueue

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-40133	YARN-3055	YARN	The token is not renewed properly if it's shared by jobs (oozie) in DelegationTokenRenewer
BUG-40195	HIVE-11442	Hive	commons-configuration-1.6.jar is backwards incompatible and regressive
BUG-41126		Hue	Workaround HS2 memory Leak
BUG-41137	HDFS-8767	HDFS	HDFS listStatus on unix pipe file fails with NPE
BUG-41148	TEZ-2290 TEZ-2732	Tez	DefaultSorter throws ArrayIndex exceptions on 2047 Mb size sort buffers
BUG-41179	HBASE-12270	HBase	A bug in the bucket cache, with cache blocks on write enabled
BUG-41385	TEZ-2629	Tez	[TEZ-2629] LimitExceededException in Tez client when DAG has exceeds the default max
BUG-41388	HIVE-11303	Hive	Getting Tez LimitExceededException after dag execution on large query
BUG-41401	STORM-951	Storm	Storm Hive Bolt leaking a transaction for every txnBatch
BUG-41461	TEZ-2636	Tez	[TEZ-2636] Tez needs to allow inputs that don't have data in MultiMRInput
BUG-41536	SQOOP-2387	Sqoop	(2.2) NPE during sqoop import when column contains a special character
BUG-41538	YARN-3487	YARN	CapacityScheduler scheduler lock obtained unnecessarily when calling getQueue
BUG-41656		RelEng, Storm	Nimbus process dies in a 20 node cluster when there are multiple supervisors attempt to download a large jar.
BUG-41920		Falcon	(2.2) Feed replication test failed because of authentication issues
BUG-42198	HBASE-14196	HBase	(2.2) - Thrift server connection idle timeout issue
BUG-42209	YARN-2019 YARN-2136 YARN-2856 YARN-2946 YARN-2992 YARN-3094 YARN-3242 YARN-3369 YARN-3385 YARN-3469 YARN-3753 YARN-3878	YARN	Stabilizing the RM-ZK integration
BUG-42217	TEZ-2635	Tez	Backport TEZ-2635 to 2.2 and 2.3

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-42256	HADOOP-12186 YARN-2605 YARN-2713 YARN-2865 YARN-3351 YARN-3526 YARN-3804 YARN-3999	YARN	Fixes for RM HA stabilization on HDP-2.2.x
BUG-42662	ACCUMULO-3957	Accumulo	Accumulo monitor's getContentSummary calls may cause NN pauses
BUG-42682	HBASE-12449	HBase	Back port HBASE-12449 : Use the max timestamp of current or old cell's timestamp in HRegion.append()
BUG-42690	HBASE-13877	HBase	Back Port HBASE-13877 : Interrupt to flush from TableFlushProcedure causes dataloss in ITBLL
BUG-42692	HBASE-13473	HBase	Back Port HBASE-13473 : deleted cells come back alive after the stripe compaction
BUG-42696	HBASE-12686	HBase	Back Port HBASE-12686 : Failures in split before PONR not clearing the daughter regions from regions in transition during rollback
BUG-42718	HBASE-12891	HBase	Back Port HBASE-12891 : Parallel execution for Hbck checkRegionConsistency
BUG-42733	PHOENIX-1535 PHOENIX-1536 PHOENIX-1570 PHOENIX-1634	Phoenix	Backport local index related issues to 2.2-maint version
BUG-42734	PHOENIX-1448 PHOENIX-1456 PHOENIX-1600 PHOENIX-1610 PHOENIX-1650	Phoenix	Backport resource leak related issues to Phoenix 2.2-maint
BUG-42765	ACCUMULO-3445	Accumulo	Backport ACCUMULO-3445 "Server failure to login using Kerberos doesn't exit the process"
BUG-42766	ACCUMULO-3385	Accumulo	Backport ACCUMULO-3385 "DateLexicoder fails to correctly order dates prior to 1970"
BUG-42767	ACCUMULO-3745	Accumulo	Backport ACCUMULO-3745 "deadlock in SourceSwitchingIterator"
BUG-42768	ACCUMULO-2641	Accumulo	Backport ACCUMULO-2641 "scan id not set in org.apache.accumulo.core.client.admin.ActiveScan"
BUG-42810	YARN-3011	YARN	NM dies because of the failure of resource localizatio
BUG-42810	YARN-3222	YARN	RMNodeImpl#ReconnectNodeTransition should send scheduler events in sequential order
BUG-42810	YARN-3850	YARN	NM fails to read files from full disks which can lead to

Hortonworks Bug ID	Apache JIRA	Component	Summary
			container logs being lost and other issues
BUG-42810	YARN-3238	YARN	Connection timeouts to nodemanagers are retried at multiple levels
BUG-42810	YARN-3055	YARN	The token is not renewed properly if it's shared by jobs (oozie) in DelegationTokenRenewer
BUG-42876	HBASE-14258	HBase	[2.2-maint] Make region_mover.rb case insensitive in regards to hostname
BUG-42933	HBASE-13995	HBase	Backport HBASE-13995 to HDP-2.2
BUG-42947	TEZ-2687	Tez	(2.2) - Tez should release/kill all held containers before stopping services during the shutdown phase
BUG-42973	HBASE-13576	HBase	Back Port HBASE-13576 : HBase enhancement: Failure in checking one region should not fail the entire HBase operation.
BUG-43084	YARN-3978 YARN-2301	YARN	Backport YARN-3978 + YARN-2301 - Reduce ResourceManager memory footprint of Timeline events - Configurably turn off the saving of container info in Generic AHS
BUG-43107	TEZ-1808 TEZ-2734z	Tez	[TEZ-1808] Job can fail since name of intermediate files can be too long in specific situation
BUG-43153	HBASE-13413	HBase	Cherry-pick IntegrationTestReplication (HBASE-13413)
BUG-43258	TEZ-2552	Tez	CLONE[2.2] - backport TEZ-2552 to 2.2 maint

Upgrade

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-38009	HDFS-7533	HDFS	Datanode sometimes does not shutdown on receiving upgrade shutdown command

Usability

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-32732		Sqoop	Netezza skip control characters issue
BUG-35352	STORM-745	Storm	Second Commandline Parameter passed to the main class is skipped when run in windows

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-35987		Ranger	Disabling Hive Ranger does not restore original functionality
BUG-36103		Hue	(2.2) - Can't see Tez job details from Hue job browser while Tez job is running
BUG-36829	HIVE-10719	Hive	alter table rename can not rename table in non-default database
BUG-37223	TEZ-2546	Tez	Tez UI: Backport tez-ui changes to display hive query text to 2.2 maint branch
BUG-37748		Sqoop	(2.2) - Netezza skip control characters issue
BUG-38045	HIVE-9655	Hive	Dynamic partition insert with distribute by clause
BUG-39419	STORM-938	Storm	Add an option to the storm hive-bolt to flush at a certain interval
BUG-40345		Hue	(Hotfix on top of 2.2.4.2) - Hue - log empty for create and insert statements
BUG-40423		Hue	(2.2) - Hue UTF8 Filter Decoding
BUG-41033	HDFS-6673	HDFS	Provide delimited format for the HDFS Offline Image Viewer
BUG-41349	YARN-2513	YARN	[YARN-2513] Host framework UIs in YARN for use with the ATS
BUG-41579	HIVE-11112	Hive	EY: Hive client is not displaying "non-ascii" characters when queried from CLI
BUG-41696	YARN-2513	YARN	Host framework UIs in YARN for use with the ATS
BUG-41849	STORM-512	Storm	KafkaBolt doesn't handle ticks properly
BUG-42145	SPARK-6918	Spark	Secure HBase with Kerberos does not work over YARN
BUG-42149	MAPREDUCE-6238	MapReduce	[MAPREDUCE-6238]MR2 can't run local jobs with -libjars command options which is a regression from MR1
BUG-42717		Ranger	Exceptions related to XA secure audit record
BUG-42810	YARN-3585	YARN	NodeManager cannot exit on SHUTDOWN event triggered and NM recovery is enabled

1.9. Known Issues for HDP 2.2.8

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-28507		Ranger	Add timeout for connection settings
BUG-33182	HIVE-11123 HIVE-11628	Hive	Oracle 12:direct sql is failing while running test_HIVE_8295
BUG-33843		Tez	Log warnings or fail early if there are too many tasks for a job
BUG-33845		Tez	Backport TEZ-1796 to Baikal, Champlain
BUG-36322		Tez	Sometimes Tez Client keeps on showing old status even if application is finished
BUG-36587	HBASE-12366	HBase	Inclusion of HBASE-12366 in HDP 2.2-Next
BUG-36601		HBase	Backport scanner keyvalue size / partial results related fixes to 2.2
BUG-36823		HBase	[2.2] Backports of fixes to region replicas for bloomberg
BUG-38391		Ranger	Ranger setup in Ambari does not respect custom service users
BUG-39321		Falcon	falcon regression test for tuture tests .LEExp_FutureAndLatestTest, latestestTest keeps on failing on both 2.1 and 2.2-maint lines saying endtime of feed before startime of feed
BUG-40259		Ranger	<p>Issue: Upgrade to HDP 2.2.6 breaks usersync.</p> <p>Details: Ranger in HDP 2.2.6 introduced authentication between usersync and the admin REST API. If you upgrade to HDP 2.2.6, there will be a previous version of <code>/etc/ranger/usersync/conf/unixauthservice.properties</code> and <code>{{ /etc/ranger/admin/conf/security-applicationContext.xml}}</code> in place. Because <code>setup.sh</code> treats these files as templates, it expects certain base content to be there in order to generate correct configurations. It looks like there is new configuration in the default templates (the ones in <code>/usr/hdp/2.2.6.0-2800/etc/ranger/...</code>) that will not get pulled into the</p>

Hortonworks Bug ID	Apache JIRA	Component	Summary
			<p>live configs. This results in usersync pulling 401 errors every time it tries to add/modify a user.</p> <p>Workaround:</p> <ol style="list-style-type: none"> 1. Apply to the copy files: <pre> /usr/hdp/2.2.6.0-2800/etc/ranger/usersync/conf/unixauthservice.properties -> /etc/ranger/usersync/conf/unixauthservice.properties /usr/hdp/2.2.6.0-2800/etc/ranger/admin/conf/security-applicationContext.xml -> /etc/ranger/admin/conf/security-applicationContext.xml </pre> 2. Rerun setup.sh. 3. Restart admin. 4. Restart usersync.
BUG-41828		Hive	Hiveserver2 failing with OOM PermGen when using temporary functions due to Hadoop reflections utils caching classloaders
BUG-42084		Hive	Prevent misconfiguration when StorageBasedAuthorization is set on hive.security.authorization.manager
BUG-42151		Hive	HS2 JVM settings
BUG-42203		Hive	Hiveserver2 failing with OOM PermGen when using temporary functions due to DataNucleus caching classloaders
BUG-42263		HDFS	Allow reconfiguration without restart for parameters where applicable.
BUG-42275		HDFS	A new healthcheck and heartbeat port
BUG-42498		HDFS	Block received messages put heavy load on namenode
BUG-42674		Hive	HS2 concurr tests intermittently fail via http mode throwing org.apache.hive.service.auth.HttpAuthenticationException
BUG-42698	HADOOP-12318	Hadoop Common	BUG-42808 [2.2] Expose underlying LDAP exceptions in SaslPlainServer (HADOOP-12318)

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-42699	HADOOP-12324	Hadoop Common	BUG-42808 [2.2] Better exception reporting in SaslPlainServer (HADOOP-12324)
BUG-42709	HDFS-7278	HDFS	BUG-42808 [2.2] Add a command that allows sysadmins to manually trigger full block reports from a DN (HDFS-7278)
BUG-42710	HDFS-7433	HDFS	BUG-42808 [2.2] Optimize performance of DatanodeManager's node map (HDFS-7433)
BUG-42713		YARN	Timeline server OOM
BUG-42721	HIVE-11407	Hive	JDBC DatabaseMetaData.getTables with large no of tables call leads to HS2 OOM
BUG-42725	HIVE-10925	Hive	HIVE-10925 : Non-static threadlocals in metastore code can potentially cause memory leak
BUG-42790		Hive	BUG-42658 Gen Plan Changes to support multi col in clause
BUG-42802	HDFS-7923 HDFS-8930	HDFS	BUG-42808 [2.2] HDFS-7923 : The DataNodes should rate-limit their full block reports by asking the NN on heartbeat messages
BUG-42808		HDFS	[2.2] Backport fixes to HDP 2.2-october-next
BUG-42860		Ranger	EAR-1858 - Ranger UserSync seems to be limited to pulling a up to 100 users from Active Directory.
BUG-42890		HBase	Configuration for bulkload needs to be updated
BUG-42916	HIVE-11606	Hive	(2.2) - bucket mapjoin error "capacity must be power of two" when #rows < #buckets
BUG-42937	HDFS-8864	HDFS	BUG-42808 [2.2] Padding is needed in remaining space check (HDFS-8864)
BUG-42938	HDFS-8871	HDFS	BUG-42808 [2.2] Decommissioning of a node with a failed volume may not start (HDFS-8871)
BUG-42939	HDFS-8870	HDFS	BUG-42808 [2.2] Lease is leaked on write failure (HDFS-8870)
BUG-43023	HDFS-5215	HDFS	BUG-42808 [2.2] dfs.datanode.du.reserved is not considered while computing available space (HDFS-5215)
BUG-43027		HBase	Have metrics for blockcache take into account replicas

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-43029		Hive	Explain plan is missing PPD
BUG-43032		Hive	(2.2) Backport HIVE-11502 (Map side aggregation is extremely slow)
BUG-43069		Hue	CLONE - Hue does not respect Support dfs.umaskmode, fs.permissions.umask-mode when creating files or folders
BUG-43070		Hue	CLONE - Select has a 1024 fd limit. If the FD number goes above 1024 for some reason (long running process), then all apps relying on thrift will fail.
BUG-43071	HIVE-9223	Hive, Hue	CLONE - HS2/hue/tez: multiple concurrent queries not allowed in single tez session
BUG-43082		YARN	BUG-42343 [ATS V1.5] backport YARN-3942 : Timeline store to read events from HDFS (ATS V1.5 patch)into 2.2-maint
BUG-43110		Hive	NPE explain in vectorization with empty tables
BUG-43163		Hive, Oozie, Tez	In Tez, fix the same issue of MR that spawn task doesn't take files in distributed cache into classpath.
BUG-43176		Hive	Backport HIVE-5277 (HBase handler skips rows with null valued first cells when only row key is selected)
BUG-43215		Tez	Fix Standalone Tez UI to work with RM HA
BUG-43221		Tez	ClassNotFoundException in InputInitializer causes AM to crash
BUG-43224	HDFS-8155	HDFS	BUG-42808 [2.2] Support OAuth2 in WebHDFS (HDFS-8155)
BUG-43226	HDFS-8435	HDFS	BUG-42808 [2.2] Support CreateFlag in WebHdfs (HDFS-8435)
BUG-43228	HDFS-8885	HDFS	BUG-42808 [2.2] ByteRangeInputStream used in webhdfs does not override available() (HDFS-8885)
BUG-43229	HADOOP-12358	HDFS	BUG-42808 [2.2] FSShell should prompt before deleting directories bigger than a configured size. (HADOOP-12358)

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-43238	HIVE-10528	Hive	(2.2) - Hiveserver2 in HTTP mode is not applying auth_to_local rules
BUG-43246		Ranger	(2.2) Ranger admin should redirect back to login page when login cookies are no longer valid
BUG-43266		Tez	CLONE[2.2] - When you kill a job in the middle of running it, it Diagnostics shows 'killed/failed due to:null'.
BUG-43273		HDFS	HDFS limits: Namenode goes down due to one bad app doing lots of calls
BUG-43278		Hive	Password information shows up as clear text when pulling Tez configs
BUG-43281		HDFS	NN web UI should display storage used for rollback during upgrade
BUG-43297		YARN	Jobclient prints End of File Exception if RM failover when application is running
BUG-43305	HDFS-8809	HDFS	BUG-42808 [2.2] HDFS fsck reports under construction blocks as "CORRUPT" (HDFS-8809)
BUG-43307		Hive	select query on json formatted table with map type column fails
BUG-43330		Tez, YARN	Functional standalone UI (hosted in ATS) for Tez
BUG-43331		YARN	BUG-43330 [YARN-4009] CORS support for RM webservice
BUG-43335		Ranger	Processing done by Audit Shutdown hooks can confuse someone looking at logs to think that shutdown of a service is held up due to Ranger plugin
BUG-43341	HDFS-8965	HDFS	BUG-42808 [2.2] Harden edit log reading code against out of memory errors (HDFS-8965)
BUG-43350		Ranger	Ranger install fails as - ojdbc.jar missing in classpath in ambari2.1.0
BUG-43403	HIVE-11581	Hive	(2.2) - HS2 url parameters should be stored in zookeeper
BUG-43411		HDP / Stack	Need a central doc for calling out all the steps for applying a hot-fix for each and every component
BUG-43416	YARN-4107	YARN	[YARN-4107]Both RM becomes Active if all

Hortonworks Bug ID	Apache JIRA	Component	Summary
			zookeepers can not connect to active RM
BUG-43433	RANGER-629	Hive	[Ranger][hive] security issue with hive & ranger for import table command
BUG-43465		Hive	(2.2) - Hive "SLF4J: Class path contains multiple SLF4J bindings." error
BUG-43466		HBase	(2.2) - Cherry-pick HBASE-14302 TableSnapshotInputFormat should not create back references when restoring snapshot
BUG-43468		HBase	(2.2) - Integrate 'HBASE-14313 After a Connection sees ConnectionClosingException it never recovers'
BUG-43470		HDFS	CLONE(2.2) - HDFS: Track DFSClient name in hdfs-audit.log
BUG-43493		YARN	Active RM can not be started if 2/3 ZKs are killed
BUG-43520		Tez, YARN	[ATS v1.5] Events for APP attempts published to ATS 1.5 and ATS 1.0 differs for TEZ
BUG-43543		HBase	Integrate HBASE-14309 'Allow load balancer to operate when there is region in transition by adding force flag'
BUG-43583	SLIDER-931	HBase, Slider	test_runintegrationtestzkandfspermissions failed on HBase-Slider
BUG-43588		Tez	Auto clean tez history log when meeting hdfs limit of 1048576 (can't start hive cli, etc)
BUG-43591	TEZ-2781	Tez	Unhelpful error message on Hive console from a failing Hive on Tez query
BUG-43592		Hive, Tez	Tez queries stopped running - ApplicationNotFoundException
BUG-43596		Tez	[TEZ-2064] ApplicationNotFound exception should be translated to SessionNotRunning
BUG-43598		HDFS	BUG-42808 [2.2] computeReplicationWork does not record all failure paths
BUG-43603		Hive	hive does not translate \${system:java.io.tmpdir}/\${system:user.name}/operation_logs properly

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-43604		Hive	Exception in FileSink when num_bucket > num_reducer
BUG-43609		Hive	Adjust compression buffer size to avoid creation of too many small stripes
BUG-43625	RANGER-631	Ranger	audit log is not present for deleting user
BUG-43658	HDFS-8995	HDFS	BUG-42808 [2.2] Flaw in registration bookkeeping can make DN die on reconnect (HDFS-8995)
BUG-43660		Hive, Hive View, Tez, YARN	[Umbrella] ATS v 1.5
BUG-43681	RANGER-635	Ranger	delete an user and then add a new user with same id but with diffenent role creates user with old role
BUG-43698		Hive	<p>There are some issues that can arise while upgrading older Hive Metastore versions, in the case that the Metastore tables have been created using datanucleus.autoCreateSchema rather than using the Hive Metastore creation scripts. Make a backup of your Metastore database before attempting any changes</p> <p>It can be the case that autoCreateSchema may not have created all of the Metastore tables. If there are some expected Metastore tables missing, the Metastore upgrade scripts will fail. Workaround: One way to fix this is to look at the Metastore creation script to the current version of your Hive Metastore, and determine which tables are missing in your Metastore compared to the tables that are being created in the Metastore creation script. Then create a new script containing the commands in the creation script relevant to the Metastore tables that are missing in your Hive Metastore, and run this script to ensure that your Metastore has all of the necessary tables.</p> <p>Some auto-created columns may have a different name - for example the COMMENT column in COLUMNS_V2 might instead have the name FCOMMENT. The upgrade scripts will also</p>

Hortonworks Bug ID	Apache JIRA	Component	Summary
			fail in this case as it was expecting the column names to match the names from the creation scripts. Workaround: In this case you can try to rename the column name of the Metastore column, to match the column names as they appear in the Metastore creation script.
BUG-43699		HDFS, Tez	CLONE(2.2) - Log a useful error message when the summary stream cannot be closed when shutting down an AM
BUG-43738		Hive	Ensure Hive does not propagate credentials file property to tasks, plugins
BUG-43756		Phoenix	Data not returned correctly for query on local index columns
BUG-43766	HDFS-8180	HDFS	BUG-42808 [2.2] AbstractFileSystem Implementation for WebHdfs (HDFS-8180)
BUG-43773		Spark	Issue: Spark: Fails to find table Details: Spark + Kerberos + yarn-cluster mode + remote hive meta store for SparkSQL does not work. Workaround: Try this with Spark + Kerberos + yarn-client mode for SparkSQL.
BUG-43789	HDFS-6945	HDFS	BUG-42808 [2.2] BlockManager should remove a block from excessReplicateMap and decrement ExcessBlocks metric when the block is removed (HDFS-6945)
BUG-43799	HDFS-8542	HDFS	BUG-42808 [2.2] WebHDFS getHomeDirectory behavior does not match specification (HDFS-8542)
BUG-43801	HDFS-8939	HDFS	BUG-42808 [2.2] Test(S)WebHdfsFileContextMainOperations failing on branch-2 (HDFS-8939)
BUG-43895	HDFS-8797	HDFS	BUG-42808 [2.2] WebHdfsFileSystem creates too many connections for pread (HDFS-8797)
BUG-43966	YARN-4105	YARN	(2.2) - [YARN-4105] Capacity Scheduler headroom for DRF is wrong

Hortonworks Bug ID	Apache JIRA	Component	Summary
BUG-43970		Hive	smb join query generates empty result in mr mode
BUG-43973		Ranger	Error in test_security_agent_admin.py: "E AttributeError: 'NoneType' object has no attribute 'endswith'"
BUG-43979		YARN	RM UI prints ZK RM connection state="CONNECTED" , if >50% ZKs are disconnected from active RM
BUG-43980		YARN	[UI usability] 'ResourceManager HA zookeeper connection state' should be highlighted if connection between ZK-RM is lost
BUG-43986		HBase	Integrate HBASE-12635 'Delete acl notify znode of table after the table is deleted'
BUG-44026		HDFS	[2.2-maint] Backport HDFS-9019 Adding informative message to sticky bit permission denied exception
BUG-44039		Ranger	(2.2) Data consistency across API and UI to include PUT
BUG-44055	ACCUMULO-3967	Accumulo	(2.2) Backport " ACCUMULO-3967 bulk import loses records when loading pre-split table"
BUG-44108		Flume	Flume agent fails to start after downgrade. Workaround: Downgrade, in this order: 1) clients 2) core.
BUG-50531		Kafka	Kafka file system support Issue: Encrypted file systems such as SafenetFS are not supported for Kafka. Index file corruption can occur. For more information, see: Install Kafka .
RMP-3126		HDFS	Accurate metrics and messages from Balancer
RMP-3532	HDFS-7858	HDFS	Improve HA Namenode Failover detection on the client
RMP-3847		HDFS	Provide way to observe HDFS LeaseManager state