

Hortonworks Data Platform

HDP-2.2.6 Release Notes

(Jun 9, 2015)

Hortonworks Data Platform: HDP-2.2.6 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 2.2.6 Release Notes | 1 |
| 1.1. Unsupported Features | 2 |
| 1.1.1. Technical Preview Features | 2 |
| 1.1.2. Community Features | 2 |
| 1.2. Behavioral Changes | 3 |
| 1.3. Apache Patch Information | 3 |
| 1.3.1. Accumulo 1.6.1 | 3 |
| 1.3.2. Falcon 0.6.0 | 3 |
| 1.3.3. Flume 1.5.2 | 4 |
| 1.3.4. Hadoop Common/HDFS 2.6.0 | 5 |
| 1.3.5. HBase 0.98.4 | 8 |
| 1.3.6. Hive 0.14.0 | 8 |
| 1.3.7. Knox 0.5.0 | 10 |
| 1.3.8. Oozie 4.1.0 | 11 |
| 1.3.9. Phoenix 4.2.0 | 11 |
| 1.3.10. Pig 0.14.0 | 13 |
| 1.3.11. Ranger 0.4.0 | 13 |
| 1.3.12. Slider 0.61.0 | 13 |
| 1.3.13. Storm 0.9.3 | 13 |
| 1.3.14. Tez 0.5.2 | 13 |
| 1.3.15. ZooKeeper 3.4.6 | 15 |
| 1.4. Upgrading from HDP 2.2.x to HDP 2.2.6 | 15 |
| 1.5. Fixed Issues | 18 |
| 1.6. Common Vulnerabilities and Exposures | 22 |
| 1.7. Known Issues for HDP 2.2.6 | 23 |

1. HDP 2.2.6 Release Notes

The official Apache versions of all HDP 2.2 components are unchanged from HDP 2.2.0.0. 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 Hadoop 2.6.0
- Apache Accumulo 1.6.1
- Apache DataFu 1.2.0
- Apache Falcon 0.6.0
- Apache Flume 1.5.2
- Apache HBase 0.98.4
- Apache Hive 0.14.0
- Hue 2.6.1
- 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 1.2.1
- Apache Sqoop 1.4.5
- Apache Storm 0.9.3
- Apache Tez 0.5.2
- Apache ZooKeeper 3.4.6

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. Unsupported Features

Some features exist within HDP 2.2.6, but Hortonworks does not currently support these specific capabilities.

1.1.1. Technical Preview Features

The following features are available within HDP 2.2.6, 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.

- HDFS Transparent Data at Rest Encryption
- Falcon Graph View of Dependencies
- Phoenix support for local secondary indexes
- YARN support for Docker
- Create a Hive Sink based on the new Hive Streaming support ([FLUME-1734](#))

1.1.2. 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.

The following features are shipped as part of HDP 2.2 HDFS, but are not supported:

- NameNode Federation ([HDFS-1052](#))
- viewFS ([HADOOP-7257](#))
- block-volume device choosing ([HDFS-1804](#))
- Disable File channel dual checkpointing on Windows ([FLUME-2224](#))
- Disable Morphline Sink and Kite Dataset sinks on Windows Builds ([FLUME-2467](#))

The following Apache components are shipped as part of HDP 2.2 YARN, but are not supported:

- Fair Scheduler

- MapReduce Uber AM
- MapReduce Eclipse Plug-in

The following components are shipped as part of HDP 2.2 Spark, but not supported:

- Spark Standalone
- BlinkDB
- GraphX

The following components are shipped as part of HDP 2.2 Kafka, but not supported:

- Delete Topic

1.2. Behavioral Changes

In HDP 2.2.6, behavioral changes affect the following Hadoop components:

- Oozie

Previous releases of the HDP Oozie client loaded the server-side configuration file, where JAVA_HOME is specific. The current version of the client does not load that server-side file. The workaround requires the following changes in its environmental setup:

- The Oozie system environment requires JDK 1.6 or greater
- The machine where the users will run the Oozie command line requires \$JAVA_HOME to be set to version 1.6 or greater.

1.3. Apache Patch Information

This section describes Apache JIRAs that were addressed in HDP 2.2.6.

1.3.1. Accumulo 1.6.1

HDP 2.2.6 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.

1.3.2. Falcon 0.6.0

HDP 2.2.6 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

1.3.3. Flume 1.5.2

HDP 2.2.6 provides Apache Falcon 0.6.0 and the following additional Apache patches:

- [FLUME-1734](#): Create a Hive Sink based on the new Hive Streaming support.
- [FLUME-1334](#): Flume startup script for Windows.
- [FLUME-2442](#): Need an alternative to providing clear text passwords in flume config.
- [FLUME-2224](#): Disable File channel dual checkpointing on Windows.
- [FLUME-2467](#): Disable Morphline Sink and Kite Dataset sinks on Windows Builds.
- [FLUME-2226](#): Refactor BlobHandler out of morphline sink and into HTTP source.
- [FLUME-2227](#): Move BlobDeserializer from Morphline Sink to FLUME-ng-core.
- [FLUME-2337](#): export JAVA_HOME in FLUME-env.sh.template and increase heap size.
- [FLUME-2511](#): Allow configuration of enabled protocols in Avro source and Rpc client.
- [FLUME-2586](#): HDFS Sink should have an option to try rename even if close fails.
- [FLUME-2655](#): Update documentation for hdfs.closeTries based on FLUME-2586.
- [FLUME-2095](#): JMS source with TIBCO (patch-1).
- [FLUME-2451](#): HDFS Sink Cannot Reconnect After NameNode Restart.
- [FLUME-2068](#): File Channel issue: recovering from BadCheckpoint exception on Windows.
- [FLUME-2224](#): Disable File channel dual checkpointing on Windows.
- [FLUME-2407](#): Spillable Channel sometimes fails on reconfigure.
- [FLUME-2145](#): TestCheckpointRebuilder.testFastReplay fails on Windows.
- [FLUME-2218](#): TestFileChannelIntegrityTool tests failing on Windows.
- [FLUME-2219](#): Windows: Flume tests need to know location of Hadoop native libraries (hadoop.dll).
- [FLUME-2137](#): Fix StagedInstall.java to invoke the correct startup script on Windows.
- [FLUME-2058](#): TestFlumeEventQueue in FileChannel fails on Windows.
- [FLUME-2136](#): Fix intermittent test failure in TestMonitoredCounterGroup on Windows.
- [FLUME-2150](#): Fix TestFileChannelEncryption failure on Windows.
- [FLUME-2358](#): File Channel needs to close BackingStore & EventQueue before deleting files in checkpoint directory.
- [FLUME-2359](#): TestFileChannelIntegrityTool throws exception on class teardown on Windows.

- [FLUME-2402](#): Warning seen when overflow is disabled for Spillable Channel.
- [FLUME-2412](#): Improve Logging in Spillable Channel.
- [FLUME-2122](#): Minor cleanups of User guide.
- [FLUME-2175](#): Update Developer Guide with notes on how to upgrade Protocol Buffer version.
- [FLUME-2123](#): Morphline Solr sink missing short type name.
- [FLUME-2162](#): TestHDFSEventSinkOnMiniCluster.maxUnderReplicationTest fails on hadoop2.
- [FLUME-2146](#): Windows: Tmp file creation in TestBodyTextEventSerializer.java needs fixing.
- [FLUME-2151](#): Windows: Update TestExecSource to use native commands on Windows.
- [FLUME-2501](#): Updating HttpClient lib version to ensure compat with Solr.
- [FLUME-2508](#): LineDeserializer causes incorrect behavior in SpoolDir Source on Windows.
- [FLUME-2530](#): Resource leaks found by Coverity tool.
- [FLUME-2532](#): Windows : TestReliableSpoolingFileEventReader : need to close filereader.
- [FLUME-2534](#): Windows : Windows : TestFlumeEventQueue has many failures.
- [FLUME-2541](#): Bug in TestBucketWriter.testSequenceFileCloseRetries.
- [FLUME-2441](#): Unit test TestHTTPSource.java failed with IBM JDK 1.7.
- [FLUME-2520](#): HTTP Source should be able to block a prefixed set of protocols.
- [FLUME-2533](#): HTTPS tests fail on Java 6.
- [FLUME-2544](#): Windows: Incorrect Path Separator used in HDFS path (HDFS Sink).

1.3.4. Hadoop Common/HDFS 2.6.0

HDP 2.2.6 provides Apache Hadoop Core 2.6.0 and the following additional Apache patches:

- [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.
- [HADOOP-926](#): Do not fail job history iteration when encountering missing directories.
- [HADOOP-941](#): Addendum patch.
- [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 KMSSClientProvider.

- [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.
- [MAPREDUCE-6230](#): Fixed RMContainerAllocator to update the new AMRMToken service name properly.
- [YARN-570](#): Time strings are formatted in different timezone.
- [YARN-2246](#): Made the proxy tracking URL always be http(s)://proxy addr:port/proxy/<appld> 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.

1.3.5. HBase 0.98.4

HDP 2.2.6 provides HBase 0.98.4 and the following additional Apache patches:

- [HBASE-212](#): Addendum. Fixes a unit test
- [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-12575](#): Sanity check table coprocessor classes are loadable
- [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-12791](#): HBase does not attempt to clean up an aborted split when the regionserver shutting down
- [HBASE-12714](#): RegionReplicaReplicationEndpoint should not set the RPC Codec
- [HBASE-12958](#): SSH doing hbase:meta get but hbase:meta not assigned
- [HBASE-13120](#): Allow disabling hadoop classpath and native library lookup

1.3.6. Hive 0.14.0

HDP 2.2.6 provides 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-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
- [HIVE-9832](#): Merge join followed by union and a map join in hive on tez fails.

1.3.7. Knox 0.5.0

HDP 2.2.6 provides Knox 0.5.0 and the following additional Apache patches:

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

1.3.8. Oozie 4.1.0

HDP 2.2.6 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.

1.3.9. Phoenix 4.2.0

HDP 2.2.6 includes Apache Phoenix 4.2.0 and the following additional Apache patches:

- [PHOENIX-18](#): Several tests are failing with `SequenceNotFoundException` error
- [PHOENIX-32](#): Some phoenix tests are failing with different assertion errors
- [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](#): Correct arg count in FunctionParseNode
- [PHOENIX-1485](#): Add timezone awareness
- [PHOENIX-1486](#): Fix SpooledTmpFileDeleteIT by having it use a different directory for pooling 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

1.3.10. Pig 0.14.0

HDP 2.2.6 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-4377-2.patch)
- [PIG-4377](#): Skewed outer join produce wrong result in some cases

1.3.11. Ranger 0.4.0

HDP 2.2.6 provides Ranger 0.4.0 and the following additional Apache patch:

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

1.3.12. Slider 0.61.0

HDP 2.2.6 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)

1.3.13. Storm 0.9.3

HDP 2.2.6 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.

1.3.14. Tez 0.5.2

HDP 2.2.6 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 2.2.6 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.

1.3.15. ZooKeeper 3.4.6

HDP 2.2.6 provides ZooKeeper 3.4.6 and the following additional Apache patches:

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

1.4. Upgrading from HDP 2.2.x to HDP 2.2.6

HDP 2.2.6 is a maintenance release of HDP 2.2.x. If you already have HDP 2.2.x installed, upgrading your cluster to HDP 2.2.6 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 2.2.6) side-by-side with existing components

The following table summarizes HDP 2.2.x-to-2.2.6 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 v. 2.2.6 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 2.2.6 |

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](#)

Upgrade Procedure

To upgrade your cluster from HDP 2.2.x to HDP 2.2.6:

1. Download the appropriate HDP 2.2.6 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/2.2.6.0/hdp.repo |
| RHEL/CentOS/Oracle LINUX 6 | http://public-repo-1.hortonworks.com/HDP/centos6/2.x/updates/2.2.6.0/hdp.repo |
| SLES 11 SP1 | http://public-repo-1.hortonworks.com/HDP/sles11sp1/2.x/updates/2.2.6.0/hdp.repo |
| SLES 11 SP3/SP4 | http://public-repo-1.hortonworks.com/HDP/suse11sp3/2.x/updates/2.2.6.0/hdp.repo |
| Ubuntu 12 | http://public-repo-1.hortonworks.com/HDP/ubuntu12/2.x/updates/2.2.6.3/hdp.list |
| Debian 6 | http://public-repo-1.hortonworks.com/HDP/debian6/2.x/updates/2.2.6.3/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/HDP-2.2.6.0-centos5-rpm.tar.gz |
| RHEL/CentOS/Oracle LINUX 6 | http://public-repo-1.hortonworks.com/HDP/centos6/HDP-2.2.6.0-centos6-rpm.tar.gz |
| SLES 11 SP1 | http://public-repo-1.hortonworks.com/HDP/sles11sp1/HDP-2.2.6.0-sles11sp1-rpm.tar.gz |
| SLES 11 SP3/SP4 | http://public-repo-1.hortonworks.com/HDP/suse11sp3/HDP-2.2.6.0-suse11sp3-rpm.tar.gz |
| Ubuntu 12 | http://public-repo-1.hortonworks.com/HDP/ubuntu12/HDP-2.2.6.3-ubuntu12-rpm.tar.gz |

| Operating System | Tarball Location |
|------------------|---|
| Debian 6 | http://public-repo-1.hortonworks.com/HDP/debian6/HDP-2.2.6.3-debian6-rpm.tar.gz |

2. Run an update:

```
apt-get update
```

3. Install the HDP 2.2.6 bits:

| Operating System | Commands |
|--------------------------|---|
| RHEL/CentOS/Oracle LINUX | Install HDP 2.2.6 components on relevant nodes, according to the services that run on those hosts: <pre>yum install "hadoop_2_2_6_2800*" "oozie_2_2_6_2800*" "pig_2_2_6_2800*" "sqoop_2_2_6_2800*" "zookeeper_2_2_6_2800*" "hbase_2_2_6_2800*" "hive_2_2_6_2800*" "tez_2_2_6_2800*" "storm_2_2_6_2800*" "falcon_2_2_6_2800*" "flume_2_2_6_2800*" "phoenix_2_2_6_2800*" "accumulo_2_2_6_2800*" "mahout_2_2_6_2800*"</pre> |
| SLES | Install HDP 2.2.6 components on relevant nodes, according to the services that run on those hosts: <pre>zypper install "hadoop_2_2_6_2800*" "oozie_2_2_6_2800*" "pig_2_2_6_2800*" "sqoop_2_2_6_2800*" "zookeeper_2_2_6_2800*" "hbase_2_2_6_2800*" "hive_2_2_6_2800*" "tez_2_2_6_2800*" "storm_2_2_6_2800*" "falcon_2_2_6_2800*" "flume_2_2_6_2800*" "phoenix_2_2_6_2800*" "accumulo_2_2_6_2800*" "mahout_2_2_6_2800*"</pre> |
| Ubuntu/Debian | Install HDP 2.2.6 components on relevant nodes, according to the services that run on those hosts: <pre>apt-get install "hadoop_2_2_6_2800*" "oozie_2_2_6_2800*" "pig_2_2_6_2800*" "sqoop_2_2_6_2800*" "zookeeper_2_2_6_2800*" "hbase_2_2_6_2800*" "hive_2_2_6_2800*" "tez_2_2_6_2800*" "storm_2_2_6_2800*" "falcon_2_2_6_2800*" "flume_2_2_6_2800*" "phoenix_2_2_6_2800*" "accumulo_2_2_6_2800*" "mahout_2_2_6_2800*"</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 2.2.6.

On each host in the cluster, use `hdps-select` to switch all services to the HDP 2.2.6 version:

```
hdp-select set all $VERSION-$BUILD
```

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 2.2.6 repositories for HDP and HDP-UTILS:

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

6. Start all HDP 2.2.6 services, in the following order:

a. ZooKeeper

```
su - zookeeper export ZOOCFGDIR=/usr/hdp/current/zookeeper-  
server/conf ; export ZOOCFG=zoo.cfg; source /usr/hdp/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.

1.5. 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.

Potential Data Loss

| Hortonworks Bug ID | Apache JIRA | Component | Summary |
|--------------------|---|-----------|--|
| BUG-34432 | HADOOP-11482 , HDFS-7208 | HDFS | NameNode doesn't schedule replication when a DataNode fails |
| BUG-34958 | HDFS-7960 , HDFS-7575 , HDFS-7596 | HDFS | The full block report should prune zombie storages even if they're not empty |

Security

| Hortonworks Bug ID | Apache JIRA | Component | Summary |
|--------------------|-----------------------------|-----------|--|
| BUG-23732 | | YARN | Fix how ZooKeeperSecurity works with ResourceManager |
| BUG-31949 | | Ranger | In hive plugin TABLE policy restriction for all tables {{*}} fails when UDF policy for all functions {{*}} is maintained for a database. |
| BUG-33070 | HBASE-13239 | HBase | HBASE grants at specific column level does not work for Groups |
| BUG-33167 | | Hive | Hive does not prevent non-whitelisted config params from being set under certain conditions |
| BUG-33792 | | Ranger | Bundle oraclejdbc.jar with Ranger |
| BUG-33822 | | Ranger | Do not display the Tomcat SessionId as part of Audit report |
| BUG-34640 | | Ranger | Admin REST API appears to be open to password-less modifications |
| BUG-35442 | | Ranger | UserSync Process didn't sync the group when groups are added to the user at later time |
| BUG-35674 | | Hue | Support LDAP authentication via Hue to Hiveserver2 |

Incorrect Results

| Hortonworks Bug ID | Apache JIRA | Component | Summary |
|--------------------|---------------------------|-----------|--|
| BUG-35606 | YARN-2906 | YARN | CapacitySchedulerPage shows HTML tags for a queue's Active Users |
| BUG-35923 | TEZ-2397 | Tez | Tez jobs can fail on a cluster with HDFS HA |
| BUG-36123 | HDFS-8219 | Ranger | setStoragePolicy with folder behavior is different after cluster restart |

| Hortonworks Bug ID | Apache JIRA | Component | Summary |
|--------------------|-------------|-----------|---|
| BUG-36124 | | Ranger | First name and last name fields should support underscores, dashes and space characters |

Stability

| Hortonworks Bug ID | Apache JIRA | Component | Summary |
|--------------------|---|---------------------|--|
| BUG-32795 | | Hive | HIVE schema failed to upgrade from HDP 2.0/2.1 to HDP 2.2 with schematool |
| BUG-33402 | | Hive | HiveServer2 fails to start without kerberos ticket in place before starting. |
| BUG-33403 | | Hive | Tez job submission fails via WebHCat |
| BUG-33546 | YARN-3393 | YARN | Getting application(s) goes wrong when app finishes before starting the attempt |
| BUG-33904 | HIVE-20085 | Hive | Laterview on top of a view throws RuntimeException error |
| BUG-34033 | HIVE-10208 | Hive | Hive job via WebHCat now requires /tez-client/conf/tez-site.xml |
| BUG-34388 | HDFS-4882 | Hadoop Common, HDFS | Prevent the Namenode's LeaseManager from looping forever in checkLeases |
| BUG-34388 | YARN-2992 | YARN | ZKRMStateStore crashes due to session expiry |
| BUG-34682 | HADOOP-11333 | Hadoop Common, HDFS | Fix deadlock in DomainSocketWatcher when the notification pipe is full |
| BUG-34965 | TEZ-2334 | Tez | ContainerManagementProtocolProxy modifies IPC timeout configuration without making a backup copy |
| BUG-35600 | YARN-2816 | YARN | NM fail to start with NPE during container recovery |
| BUG-35601 | YARN-2874 | YARN | Dead lock in "DelegationTokenRenewer" which blocks RM to execute any further apps |
| BUG-35602 | YARN-2340 | YARN | Fixed NPE when queue is stopped during RM restart |
| BUG-35604 | YARN-2414 | YARN | RM web UI: app page will crash if app is failed before any attempt has been created |
| BUG-35605 | YARN-2905 | YARN | AggregatedLogsBlock page can infinitely loop if the aggregated log file is corrupted |
| BUG-36599 | HBASE-13515 , HBASE-13169 , HBASE-13469 , HBASE-13518 | HBase | Fixes to region replicas |

Upgrade

| Hortonworks Bug ID | Apache JIRA | Component | Summary |
|--------------------|-------------|-----------|---|
| BUG-33818 | | Knox | Persisted service registry is not updated to support HA after upgrade |
| BUG-34770 | | Knox | Persisted service registry is not updated to support HA after upgrade |

Usability

| Hortonworks Bug ID | Apache JIRA | Component | Summary |
|--------------------|---------------------------|-----------|--|
| BUG-34718 | HIVE-9613 | Hive | Select operator missing column, wrong results in multiple left outer join with sub queries |
| BUG-35606 | YARN-2096 | YARN | CapacitySchedulerPage shows HTML tags for a queue's Active Users |
| BUG-35923 | TEZ-2397 | TEZ | Tez jobs can fail on a cluster with HDFS HA |
| BUG-36123 | HDFS-8219 | Ranger | setStoragePolicy with folder behavior is different after cluster restart |
| BUG-36124 | | Ranger | First name and last name fields should support underscores, dashes and space characters |

Performance

| Hortonworks Bug ID | Apache JIRA | Component | Summary |
|--------------------|---------------------------|-----------|---|
| BUG-34296 | HDFS-7531 | HDFS | Improve the concurrent access on FsVolumeList. |
| BUG-35569 | | Falcon | Provide the ability to disable oozie dryrun while scheduling the Falcon process |

Other

| Hortonworks Bug ID | Apache JIRA | Component | Summary |
|--------------------|----------------------------|-----------------|--|
| BUG-33342 | | Ranger | Setup Nexus proxy within openstack |
| BUG-33674 | | Phoenix | Support cross joins and comma-separated table list |
| BUG-33815 | | Ranger, Windows | Validation of Ranger settings is missing in MSI |
| BUG-34504 | | Ranger | Remove condition to copy DB driver jar file |
| BUG-34505 | | Phoenix | Unable to load tab delimited data via Phoenix |
| BUG-34523 | FLUME-2586 | Flume | HDFS Sink should have an option to try to rename even if close fails |

| Hortonworks Bug ID | Apache JIRA | Component | Summary |
|--------------------|------------------------------|---------------|--|
| BUG-35266 | | Hive | Tez 2 sub queries with same group by key, when one set group by key empty, outer join on those group-by keys generate different wrong results depending on #reducers of the join |
| BUG-35340 | HADOOP-11730 | Hadoop Common | Running Jobs using S3n results in 'java.net.SocketTimeoutException' when using large >2G data sets |
| BUG-35343 | HADOOP-11710 | Hadoop Common | Make CryptoOutputStream behave like DFSOutputStream with respect to synchronization |
| BUG-35597 | FLUME-2095 | Flume | JMS source setup fails with "org.apache.flume.FlumeException: Could not lookup ConnectionFactory" |
| BUG-35643 | HADOOP-11482 | HDFS | Use correct UGI when KMSClientProvider is called by a proxy user |
| BUG-35670 | HIVE-10421 | Hive | StorageBasedAuthorizationProvider can't drop partitioned table with database prefix |
| BUG-36123 | HDFS-8219 | HDFS | . setStoragePolicy with folder behavior is different after cluster restart |

1.6. Common Vulnerabilities and Exposures

- **CVE-2015-1772:** Apache Hive Authentication vulnerability in HiveServer2

Severity: Important

Vendor: Hortonworks

Versions Affected: HDP 1.2-2.2; fixed in HDP 2.1.15+ and 2.2.6+.

Users Affected: Users who use LDAP authentication mode in HiveServer2 and also have LDAP configured to allow unauthenticated (anonymous) binds.

Impact: See BUG-41739 and [HIVE-9934](#). LDAP services are sometimes configured to allow simple unauthenticated binds. When HiveServer2 is configured to use LDAP authentication mode (hive.server2.authentication configuration parameter is set to LDAP), with such LDAP configurations, it can allow users without proper credentials to get authenticated. This is more easily reproducible when Kerberos authentication is also enabled in the Apache Hadoop cluster.

Recommended Action: This vulnerability can be addressed in any one of the following ways:

- Upgrade to HDP 2.1 version newer than 2.1.15.0, or HDP 2.2 version newer than 2.2.6.0.

- Configure LDAP service to disallow unauthenticated binds. If the service allows anonymous binds, not having Hive authorization checks enabled can also expose this vulnerability.
- Update the Hive installation to use an authenticator with the fix. HiveServer2 LDAP authenticator is packaged in ldap-fix.tar.gz, which is available at <http://public-repo-1.hortonworks.com/security/ldap-fix.tar.gz>.

Please use the instructions in the README.txt for install and configuration instructions. Once there is a full HDP patch release available, this notice will be replaced with instructions for obtaining the new release.

1.7. Known Issues for HDP 2.2.6

| Hortonworks Bug ID | Apache JIRA | Component/s | Summary |
|--------------------|---------------------------|-------------|---|
| BUG-18450 | HIVE-7160 | Hive | Vectorization Udf: GenericUDFConcat, is not supported for String+Long |
| BUG-19255 | | Zookeeper | Uninstalling the ZooKeeper RPM does not remove all ZooKeeper directories |
| BUG-19911 | | Falcon | Ensure falcon windows installer takes secure cluster installation into account |
| BUG-19912 | | Falcon | Document how to configure Falcon on Kerberized HDP for Windows |
| BUG-19930 | | Falcon | Document entity management dashboard |
| BUG-20856 | | Hive | Column names are missing from join expression in Map join with CBO enabled |
| BUG-20866 | | Mahout | Mahout trainnb command fails with -c option |
| BUG-21166 | | Phoenix | testGetCaseInsensitiveDynCol fails intermittently throwing AssertionError |
| BUG-21254 | | HBase | hbck tool fixMeta flag not able to recreate deleted entries from hbase:meta table when just one column is deleted |
| BUG-21468 | | Hive | With CBO enabled queries fail with Failed to breakup Windowing invocations into Groups |
| BUG-22522 | | YARN | MR Job failed when run via oozie and NN and RM were being killed |
| BUG-22807 | | YARN | Non-secure slider app processes run with yarn user. |
| BUG-22841 | | HBase | HBase test failed with Insufficient permissions |

| Hortonworks Bug ID | Apache JIRA | Component/s | Summary |
|--------------------|------------------------------|-------------|--|
| | | | (user=hrt_qa@HWQE.HORTONWORKS.COM error |
| BUG-23260 | HIVE-11421 | Hive | BUG-22878 insert after alter table issue |
| BUG-23527 | | Oozie | Oozie config referred to versioned hadoop config |
| BUG-23619 | | HBase | On Windows, tests fail complaining about Cannot run program /usr/bin/env |
| BUG-23655 | | Oozie | Oozie stuck on mapreduce job because workflow state isn't changing from RUNNING to KILLED |
| BUG-23889 | | Knox | Windows security scripting should setup Knox to use ActiveDirectory for LDAP |
| BUG-23897 | | HBase | HBase REST server continuously restarting in secure windows env |
| BUG-23981 | HADOOP-7844= | HBase | IOException: Invalid HFile block magic in Windows unit tests |
| BUG-24136 | | Falcon | Special/utf8 chars should not be allowed in pipeline name |
| BUG-24518 | | Tez | Tez needs to support Application Tag |
| BUG-24589 | | Falcon | DR - set up replication for HDFS |
| BUG-24590 | | Falcon | DR - set up replication for Table |
| BUG-24591 | | Falcon | BUG-24589 UI Mock for setting up HDFS Replication |
| BUG-24592 | | Falcon | BUG-24589 Rest API design for HDFS replication |
| BUG-24594 | | Falcon | BUG-24590 UI Mock for setup table replication |
| BUG-24595 | | Falcon | BUG-24590 Rest API design for set up table replication |
| BUG-24597 | | Falcon | DR - delete replication for HDFS |
| BUG-24598 | | Falcon | DR - delete replication for Table |
| BUG-24831 | | Tez | TEZ windows test framework needs to respect \${fs.defaultFS} in tez configs. |
| BUG-25049 | | Knox | Two WebHdfsHaFuncTest are unstable and should be fixed and enabled |
| BUG-25064 | | Hive | Hive : Q51 fails in CBO with "Failed to breakup Windowing invocations into Groups. At least 1 group must only depend on input columns" |

| Hortonworks Bug ID | Apache JIRA | Component/s | Summary |
|--------------------|-------------|-------------|---|
| BUG-25194 | | Tez | AM running out of memory for partitioned bucketed tables |
| BUG-25380 | | Tez | For hive to work with Virtual columns and paritition names, we need information available within a record-reader |
| BUG-25451 | | Accumulo | [Moved to Dal] Integrate Accumulo in RU Automation Frame work |
| BUG-25961 | | HBase | HBase tests failing with org.apache.hadoop.hbase.NotAllMetaRegionsOnTimed out; 10000ms |
| BUG-26243 | | Hive | CBO : Q51 fails in explain with Failed to breakup Windowing invocations into Groups. |
| BUG-26307 | | Falcon | As a user, I need to be able to replicate the action of creating new database at the target Hive DB |
| BUG-26309 | | Falcon | As a user, I need to be able to replicate the action of dropping a database at the target Hive DB |
| BUG-26310 | | Falcon | As a user, I need to be able to replicate the action of altering database at the target Hive DB |
| BUG-26311 | | Falcon | As a user, I need to be able to replicate the action of creating new table at the target Hive DB |
| BUG-26312 | | Falcon | As a user, I need to be able to replicate the action of dropping a table at the target Hive DB |
| BUG-26313 | | Falcon | As a user, I need to be able to replicate the action of altering a table at the target Hive DB |
| BUG-26314 | | Falcon | As a user, I need to be able to replicate the action of inserting rows into unpartitioned table at the target Hive DB |
| BUG-26315 | | Falcon | As a user, I need to be able to replicate the action of inserting and overwriting rows into unpartitioned table at the target Hive DB |
| BUG-26318 | | Falcon | As a user, I need to be able to replicate the action of altering a partition in a table at the target Hive DB |
| BUG-26319 | | Falcon | As a user, I need to be able to replicate the action of inserting rows into a |

| Hortonworks Bug ID | Apache JIRA | Component/s | Summary |
|--------------------|-------------|---------------|--|
| | | | partition of a table at the target Hive DB |
| BUG-26320 | | Falcon | As a user, I need to be able to replicate the action of inserting and overwriting rows into a partition of a table at the target Hive DB |
| BUG-26322 | | Falcon | As a user, I need to be able to replicate Date partitions at the target Hive DB |
| BUG-26323 | | Falcon | As a user, I need to be able to replicate Site partitions at the target Hive DB |
| BUG-26324 | | Falcon | As a user, I need to be able to replicate Regex partitions at the target Hive DB |
| BUG-26325 | | Falcon | As a user, I need to be able to replicate custom non standard partitions at the target Hive DB |
| BUG-26326 | | Falcon | As a user, I need to be able to replicate multiple tables at once in a database at the target Hive DB |
| BUG-26327 | | Falcon | As a user, I need to be able to replicate multiple table names (defined by regular expressions) at the target Hive DB |
| BUG-26328 | | Falcon | As a user, I need to be able to replicate Avro partitions at the target Hive DB |
| BUG-26329 | | Falcon | As a user, I need to be able to replicate dynamic partitions at the target Hive DB |
| BUG-26330 | | Falcon | As a user, I need to be able to replicate hive roles at the target Hive DB |
| BUG-26419 | | Tez | Tez application fails with diff client and AM version |
| BUG-26428 | | Ranger | XaAdmin and XaAgent jobs for Debian failing on Unix clusters. |
| BUG-26661 | | Phoenix | Windows runs failed with "region isn't co-located with its data" exception |
| BUG-26944 | | Falcon, Oozie | E0803: IO error, null during scheduling feed on a cluster with wire enc on |
| BUG-27089 | | Knox | Knox YarnSysTest skipped in secure mode for gsInstaller test config |
| BUG-27091 | | Knox | Knox YarnSysTest skipped in secure mode for ambari test config |

| Hortonworks Bug ID | Apache JIRA | Component/s | Summary |
|--------------------|---|-------------|--|
| BUG-27094 | | Knox | Knox HBase tests skipped for ambari test config |
| BUG-27194 | | YARN | [YARN-2821] Failures encountered on Distributed Shell on windows |
| BUG-27256 | | HBase | Create mr based Hbase Long running application for Rollingupgrade |
| BUG-27479 | | Falcon | falcon idempotency is being violated |
| BUG-27507 | | Hive | Hive CLI returns 0 even in case of failures |
| BUG-27509 | HBASE-12464 | HBase | test_ReplicaCopyTable_verifyCopyTable[tableone] fails after trying several times for 'Failed to find location, tableName=hbase:meta, row=' |
| BUG-27513 | | Storm | Storm-slider default config has ganglia error in supervisor.log |
| BUG-27562 | | Tez | WordCount on Tez with non-nightly failed with NPE |
| BUG-27582 | | Hive | Beeline Client Connections Errors - Obscure errors and exception swallow - 1 |
| BUG-27590 | | YARN | Distributed shell job does not execute python script in Windows |
| BUG-27632 | | Zookeeper | Add feature of retry in zookeeper client |
| BUG-27636 | HIVE-10500 | Hive | Oracle: intermittent acid_concurrency test failures due to NoSuchLockException |
| BUG-27638 | | Falcon | DEV - As a user i need a login page when i run Falcon as a standalone app |
| BUG-27639 | | Falcon | DEV - As a user, i need a way to logout when logged in the standalone version of Falcon |
| BUG-27656 | HBASE-12472 , HBASE-13192 | HBase | IntegrationTestBulkLoad fails on mapreduce |
| BUG-27708 | | Falcon | DEV - When uploading an XML of a cluster that already exists (same name) the system allows the operation |
| BUG-27734 | | Falcon | DEV - Column titles should have 14px size |
| BUG-27737 | | Falcon | DEV - When loading the dashboard, the system should show a spinning loader for tables that are updating, and just a |

| Hortonworks Bug ID | Apache JIRA | Component/s | Summary |
|--------------------|----------------------------|----------------|---|
| | | | plain message for columns without entities |
| BUG-27756 | | HBase | test_IntegrationTestMultiGetRegionReplicas slowness |
| BUG-27769 | | Sqoop | [Rolling upgrade] [Test issue] Sqoop Long running job finishes before upgrade of all component finishes |
| BUG-27858 | | HBase | HBase HA tests fail with Table Namespace Manager not ready yet |
| BUG-27893 | | Phoenix | testCurrentDate and testCurrentTime are failing due to NTP on windows server |
| BUG-27927 | | Falcon | DEV - Correct the steps information |
| BUG-27928 | | Falcon | DEV - The form needs a "Target Clusters" title and "add cluster" button in step 4 of feed setup |
| BUG-27929 | | Falcon | DEV - Put a label before the input and add archive option in step 4 of feed setup |
| BUG-27933 | | Falcon | DEV - Each cluster or input container must have it's title |
| BUG-27936 | | Falcon | DEV - Correct vertical spacings between labels and inputs in all setups |
| BUG-28011 | | Phoenix | testInnerJoinWithPreFilters and testJoinOverSaltedTables are failing with SequenceNotFoundException exception |
| BUG-28031 | | Knox | Knox rolling upgrade system test support for wire encryption |
| BUG-28039 | KNOX-534 | Knox | Knox Failed Authentication Auditing |
| BUG-28049 | | Falcon, RelEng | lineage is not enabled for Falcon in MSI |
| BUG-28080 | FALCON-885 | Falcon | Merge falcon-885 to internal repo for Champlain-M10 |
| BUG-28218 | | Hive | Hive CLI hangs |
| BUG-28276 | | Falcon | Falcon actions intermittently fail with java.net.SocketTimeoutException on windows |
| BUG-28344 | | Falcon | DEV - XML editor should only be able to change it's vertical size, not horizontal |
| BUG-28348 | | Falcon | DEV - Increase vertical spacing in the space for the action buttons |

| Hortonworks Bug ID | Apache JIRA | Component/s | Summary |
|--------------------|--------------------------------|---------------|--|
| BUG-28394 | | Falcon | DEV - Correct font sizes in forms |
| BUG-28614 | | Oozie | Oozie property oozie.authentication.kerberos.principal is being used for authentication of webhdfs endpoint |
| BUG-28664 | | Falcon | DEV - Fix layout for user generated inputs |
| BUG-28977 | FALCON-959 | Falcon | Unable to update feed's acl to a different user of different group |
| BUG-29261 | | Phoenix | Tests are failing with ERROR 1200 and then raising an error due to jobManager rejection |
| BUG-29287 | | Falcon | Backport https://issues.apache.org/jira/browse/FALCON-829 |
| BUG-29427 | HIVE-9235 | Hive | Turn off Parquet Vectorization until all data types work: DECIMAL, DATE, TIMESTAMP, CHAR, and VARCHAR |
| BUG-29730 | | Storm | Storm Hive,Hbase, Hadoop connector should support user specified UGI in non-secure env |
| BUG-29864 | FALCON-1101 | Falcon | Cluster submission in falcon does not create an owned-by edge |
| BUG-29903 | | HBase | IntegrationTestMTTR slowness |
| BUG-30055 | | Falcon | Falcon recipe should have rest api |
| BUG-30784 | ZOOKEEPER-1952 | Zookeeper | zookeeper.log.file property is not respected; log output goes only to the zookeeper.out |
| BUG-30984 | HIVE-10521 | Hive | Hive compaction tests failing |
| BUG-31349 | | Oozie | oozie job in running state is causing intermittent timeout failures |
| BUG-31353 | | HBase | IntegrationTestRegionReplicaPerf region not serving |
| BUG-31426 | | Oozie, RelEng | Oozie failed to start due to "SA" schema not created |
| BUG-31497 | | Phoenix | [RollingUpgrade] Phoenix intermittently fails to return every column on sql queries |
| BUG-31522 | | Tez | You cannot search by the name of 'Queue' |
| BUG-31565 | | Tez | When you kill a job in the middle of running it, it Diagnostics shows 'killed/failed due to:null'. |

| Hortonworks Bug ID | Apache JIRA | Component/s | Summary |
|--------------------|-----------------------------|---------------|--|
| BUG-31624 | | Hive | Dropping a 2000+ partition table takes over 10 minutes |
| BUG-31671 | | Hive | HS2 concurr tests intermittently fail with http mode on secure Openstack throwing org.apache.hive.service.auth.HttpAuthenticationException: Kerberos authentication failed |
| BUG-31676 | | Falcon | Falcon jar uploading occasionally fails |
| BUG-31936 | | HBase | test_hbaseBigLinkedList[Loop] TimeoutException |
| BUG-31943 | | HBase | test_IntegrationTestRegionReplicas region not serving |
| BUG-32026 | | Oozie | test_mapred_via_oozie_hadoopoozie-nn-rm-5-jobs] FAILED |
| BUG-32195 | | HBase | Need to add hbase-slider RU downgrade |
| BUG-32509 | | HBase | test_runIntegrationTestMTTRwithDistributedLogs due to failure in testKillRsHoldingMeta |
| BUG-32688 | | Storm | kafka/api/OffsetRequest is not packaged in storm-kafka-example jar |
| BUG-32763 | | Oozie | yum cannot find Xvfb for centos06 |
| BUG-32768 | FALCON-954 | Falcon | CLONE - Secure Kerberos setup : Falcon should periodically revalidate security credentials. |
| BUG-32902 | | Falcon | Falcon pom.xml should not use exact HDP versions for Hive,Oozie and Hadoop |
| BUG-33054 | | RelEng, Spark | Spark examples are packaged differently on Windows |
| BUG-33056 | | Spark | {start,stop}-thriftserver equivalent script on Windows is missing |
| BUG-33113 | | HBase | HBase-HA master fail over immediately |
| BUG-33168 | | Hive | hive authorization checks on configs should not be done if authorization is disabled |
| BUG-33180 | | Oozie | oozie logging tests fail on MR framework with postgres 9 |
| BUG-33278 | HBASE-13239 | HBase | HBASE grants at specific column level does not work for Groups |
| BUG-33335 | | Spark | spark make-distribution.sh doesn't evaluate SPARK_HIVE on OSX |

| Hortonworks Bug ID | Apache JIRA | Component/s | Summary |
|--------------------|----------------------------|-------------|---|
| BUG-33338 | HIVE-10528 | Hive | Hiveserver2 in HTTP mode is not applying auth_to_local rules |
| BUG-33764 | | HDP stack | HDP-Select needs to handle user-created directories in /usr/hdp |
| BUG-35528 | | Ranger | Users must remove empty user entries from their MYSQL database Error: Internal Exception: java.sql.SQLException: Connections could not be acquired from the underlying database! |
| BUG-35942 | | YARN | Users must manually configure ZooKeeper security with ResourceManager High Availability. |
| BUG-38088 | HDFS-7553 | HDFS | Issue: DataNode restart fails during Rolling Upgrade from 2.2.x-2.2.6 to 2.3.x. Details: Not all writers have responder running, but sendOOB() tries anyway. This causes NPE and the shutdown thread dies, halting the shutdown after only shutting down DataXceiverServer. Workaround: Two possible workarounds: <ul style="list-style-type: none">• Skip the failed datanodes, downgrade them and try to upgrade them again.• Manually restart the failed datanodes. Related bugs: BUG-38088, BUG-36247, BUG-38009 |
| BUG-40259 | | Ranger | Issue: Upgrade to HDP 2.2.6 breaks usersync. 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 /etc/ranger/usersync/conf/unixauthservice.properties and {{ /etc/ranger/admin/conf/security-applicationContext.xml }} in place. Because setup.sh treats these files as templates, it expects certain base content to be there in |

| Hortonworks Bug ID | Apache JIRA | Component/s | Summary |
|--------------------|-------------|-------------|--|
| | | | <p>order to generate correct configurations. It looks like there is new configuration in the default templates (the ones in /usr/hdp/2.2.6.0-2800/etc/ranger/...) that will not get pulled into the 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 data-bbox="1182 646 1427 947">/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-50531 | | Kafka | <p>Kafka file system support</p> <p>Issue: Encrypted file systems such as SafenetFS are not supported for Kafka. Index file corruption can occur.</p> <p>For more information, see: Install Kafka.</p> |