

Hortonworks Data Platform

HDP-2.2.6 Release Notes for Windows

(Jun 9, 2015)

Hortonworks Data Platform: HDP-2.2.6 Release Notes for Windows

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 for Windows	1
1.1. Unsupported Apache Features	1
1.2. Behavioral Changes	2
1.3. Tech Previews in This Release	2
1.4. Apache Patch Information	3
1.4.1. Falcon 0.6.0	3
1.4.2. Hadoop Common/HDFS 2.6.0	3
1.4.3. HBase 0.98.4	5
1.4.4. Hive 0.14.0	6
1.4.5. Knox 0.5.0	8
1.4.6. Oozie 4.1.0	8
1.4.7. Phoenix 4.2.0	8
1.4.8. Pig 0.14.0	10
1.4.9. Ranger 0.4.0	11
1.4.10. Slider 0.61.0	11
1.4.11. Storm 0.9.3	11
1.4.12. Tez 0.5.2	11
1.4.13. ZooKeeper 3.4.6	13
1.5. Fixed in HDP 2.2.6	13
1.6. Known Issues for HDP 2.2.6	17

1. HDP 2.2.6 Release Notes for Windows

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 DataFu 1.2.0
- Apache Falcon 0.6.0
- Apache Flume 1.52
- 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 1.2.1
- Apache Sqoop 1.4.5
- Apache Storm 0.9.3
- Apache Tez 0.5.2
- Apache ZooKeeper 3.4.6

1.1. Unsupported Apache Features

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

- NameNode Federation ([HDFS-1052](#))
- viewFS ([HADOOP-7257](#))

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
- YARN support for Docker on Windows

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

- Spark Standalone
- BlinkDB
- GraphX

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. Tech Previews in This Release

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.

- HDFS Transparent Data at Rest Encryption
- Falcon Graph View of Dependencies
- Phoenix support for local secondary indexes

1.4. Apache Patch Information

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

1.4.1. 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

1.4.2. 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.

1.4.3. 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

1.4.4. 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 templeton 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.

1.4.5. 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.

1.4.6. 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.

1.4.7. 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 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

1.4.8. 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-4538](#): Pig script fail with CNF in follow up MR job
- [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

1.4.9. 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

1.4.10. 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)

1.4.11. 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.

1.4.12. 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 HPD 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.

1.4.13. 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

1.5. Fixed in HDP 2.2.6

The following features and fixes were contributed back to Apache with the release of HDP 2.2.6:

Potential Data Loss

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

Security

Component	Apache JIRA	Key	Summary
HBase	HBASE-13239	BUG-33070	HBASE grants at specific column level does not work for Groups

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

Incorrect Results

Component	Apache JIRA	Key	Summary
Ranger		BUG-36124	First name and last name fields should support underscores, dashes and space characters
Ranger	HDFS-8219	BUG-36123	setStoragePolicy with folder behavior is different after cluster restart
Tez	TEZ-2397	BUG-35923	Tez jobs can fail on a cluster with HDFS HA

Stability

Component	Apache JIRA	Key	Summary
Hadoop Common, HDFS	HADOOP-11333	BUG-34682	Fix deadlock in DomainSocketWatcher when the notification pipe is full
Hadoop Common, HDFS	HDFS-4882	BUG-34388	Prevent the Namenode's LeaseManager from looping forever in checkLeases
HBase	HBASE-13515 , HBASE-13169 , HBASE-13469 , HBASE-13518	BUG-36599	Fixes to region replicas
Hive	HIVE-10208	BUG-34033	Hive job via WebHCat now requires /tez-client/conf/tez-site.xml

Component	Apache JIRA	Key	Summary
Hive	HIVE-20085	BUG-33904	Laterview on top of a view throws RuntimeException error
Hive		BUG-32795	HIVE schema failed to upgrade from HDP 2.0/2.1 to HDP 2.2 with schematool
Hive		BUG-33402	HiveServer2 fails to start without kerberos ticket in place before starting.
Hive		BUG-33403	Tez job submission fails via WebHCat
Tez	TEZ-2334	BUG-34965	ContainerManagementProtocolProxy modifies IPC timeout configuration without making a backup copy
YARN	YARN-2340	BUG-35602	Fixed NPE when queue is stopped during RM restart
YARN	YARN-2414	BUG-35604	RM web UI: app page will crash if app is failed before any attempt has been created
YARN	YARN-2816	BUG-35600	NM fail to start with NPE during container recovery
YARN	YARN-2874	BUG-35601	Dead lock in "DelegationTokenRenewer" which blocks RM to execute any further apps
YARN	YARN-2905	BUG-35605	AggregatedLogsBlock page can infinitely loop if the aggregated log file is corrupted
YARN	YARN-2992	BUG-34388	ZKRMStateStore crashes due to session expiry
YARN	YARN-3393	BUG-33546	Getting application(s) goes wrong when app finishes before starting the attempt

Upgrade

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

Usability

Component	Apache JIRA	Key	Summary
Ranger		BUG-36124	First name and last name fields should support underscores, dashes and space characters
Ranger	HDFS-8219	BUG-36123	setStoragePolicy with folder behavior is different after cluster restart

Component	Apache JIRA	Key	Summary
TEZ	TEZ-2397	BUG-35923	Tez jobs can fail on a cluster with HDFS HA

Performance

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

Other

Component	Apache JIRA	Key	Summary
Flume	FLUME-2095	BUG-35597	JMS source setup fails with "org.apache.flume.FlumeException: Could not lookup ConnectionFactory"
Flume	FLUME-2586	BUG-34523	HDFS Sink should have an option to try to rename even if close fails
Hadoop Common	HADOOP-11710	BUG-35343	Make CryptoOutputStream behave like DFSOutputStream with respect to synchronization
Hadoop Common	HADOOP-11730	BUG-35340	Running Jobs using S3n results in 'java.net.SocketTimeoutException' when using large >2G data sets
HDFS	HADOOP-11482	BUG-35643	Use correct UGI when KMSSClientProvider is called by a proxy user
HDFS	HDFS-8219	BUG-36123	. setStoragePolicy with folder behavior is different after cluster restart
Hive	HIVE-10421	BUG-35670	StorageBasedAuthorizationProvider can't drop partitioned table with database prefix
Hive		BUG-35266	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
Phoenix		BUG-33674	Support cross joins and comma-separated table list
Phoenix		BUG-34505	Unable to load tab delimited data via Phoenix
Ranger		BUG-34504	Remove condition to copy DB driver jar file
Ranger		BUG-33342	Setup Nexus proxy within openstack
Ranger, Windows		BUG-33815	Validation of Ranger settings is missing in MSI

1.6. Known Issues for HDP 2.2.6

Key	Apache JIRA	Summary	Component/s
BUG-26944		E0803: IO error, null during scheduling feed on a cluster with wire enc on	Falcon, Oozie
BUG-27479		falcon idempotency is being violated	Falcon
BUG-28664		DEV - Fix layout for user generated inputs	Falcon
BUG-27639		DEV - As a user, i need a way to logout when logged in the standalone version of Falcon	Falcon
BUG-27638		DEV - As a user i need a login page when i run Falcon as a standalone app	Falcon
BUG-26319		As a user, I need to be able to replicate the action of inserting rows into a partition of a table at the target Hive DB	Falcon
BUG-26318		As a user, I need to be able to replicate the action of altering a partition in a table at the target Hive DB	Falcon
BUG-26329		As a user, I need to be able to replicate dynamic partitions at the target Hive DB	Falcon
BUG-31676		Falcon jar uploading occasionally fails	Falcon
BUG-27737		DEV - When loading the dashboard, the system should show a spinning loader for tables that are updating, and just a plain message for columns without entities	Falcon
BUG-27734		DEV - Column titles should have 14px size	Falcon
BUG-28348		DEV - Increase vertical spacing in the space for the action buttons	Falcon
BUG-28394		DEV - Correct font sizes in forms	Falcon
BUG-27929		DEV - Put a label before the input and add archive option in step 4 of feed setup	Falcon
BUG-27928		DEV - The form needs a "Target Clusters" title and "add cluster" button in step 4 of feed setup	Falcon
BUG-27927		DEV - Correct the steps information	Falcon

Key	Apache JIRA	Summary	Component/s
BUG-27936		DEV - Correct vertical spacings between labels and inputs in all setups	Falcon
BUG-27933		DEV - Each cluster or input container must have it's title	Falcon
BUG-24591		BUG-24589 UI Mock for setting up HDFS Replication	Falcon
BUG-24590		DR - set up replication for Table	Falcon
BUG-24598		DR - delete replication for Table	Falcon
BUG-24594		BUG-24590 UI Mock for setup table replication	Falcon
BUG-26324		As a user, I need to be able to replicate Regex partitions at the target Hive DB	Falcon
BUG-24589		DR - set up replication for HDFS	Falcon
BUG-24597		DR - delete replication for HDFS	Falcon
BUG-26320		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	Falcon
BUG-27708		DEV - When uploading an XML of a cluster that already exists (same name) the system allows the operation	Falcon
BUG-28344		DEV - XML editor should only be able to change it's vertical size, not horizontal	Falcon
BUG-26307		As a user, I need to be able to replicate the action of creating new database at the target Hive DB	Falcon
BUG-26309		As a user, I need to be able to replicate the action of dropping a database at the target Hive DB	Falcon
BUG-26328		As a user, I need to be able to replicate Avro partitions at the target Hive DB	Falcon
BUG-26323		As a user, I need to be able to replicate Site partitions at the target Hive DB	Falcon
BUG-24595		BUG-24590 Rest API design for set up table replication	Falcon
BUG-24592		BUG-24589 Rest API design for HDFS replication	Falcon
BUG-26326		As a user, I need to be able to replicate multiple tables at once in a database at the target Hive DB	Falcon

Key	Apache JIRA	Summary	Component/s
BUG-26311		As a user, I need to be able to replicate the action of creating new table at the target Hive DB	Falcon
BUG-26325		As a user, I need to be able to replicate custom non standard partitions at the target Hive DB	Falcon
BUG-26313		As a user, I need to be able to replicate the action of altering a table at the target Hive DB	Falcon
BUG-26312		As a user, I need to be able to replicate the action of dropping a table at the target Hive DB	Falcon
BUG-26327		As a user, I need to be able to replicate multiple table names (defined by regular expressions) at the target Hive DB	Falcon
BUG-26322		As a user, I need to be able to replicate Date partitions at the target Hive DB	Falcon
BUG-26330		As a user, I need to be able to replicate hive roles at the target Hive DB	Falcon
BUG-32902		Falcon pom.xml should not use exact HDP versions for Hive,Oozie and Hadoop	Falcon
BUG-26315		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	Falcon
BUG-26310		As a user, I need to be able to replicate the action of altering database at the target Hive DB	Falcon
BUG-26314		As a user, I need to be able to replicate the action of inserting rows into unpartitioned table at the target Hive DB	Falcon
BUG-32768	FALCON-954	CLONE - Secure Kerberos setup : Falcon should periodically revalidate security credentials.	Falcon
BUG-19930		Document entity management dashboard	Falcon
BUG-19912		Document how to configure Falcon on Kerberized HDP for Windows	Falcon
BUG-19911		Ensure falcon windows installer takes secure cluster installation into account	Falcon

Key	Apache JIRA	Summary	Component/s
BUG-29287		Backport https://issues.apache.org/jira/browse/FALCON-829	Falcon
BUG-28080	FALCON-885	Merge falcon-885 to internal repo for Champlain-M10	Falcon
BUG-28977	FALCON-959	Unable to update feed's acl to a different user of different group	Falcon
BUG-24136		Special/utf8 chars should not be allowed in pipeline name	Falcon
BUG-28276		Falcon actions intermittently fail with java.net.SocketTimeoutException on windows	Falcon
BUG-28049		lineage is not enabled for Falcon in MSI	Falcon, RelEng
BUG-29864	FALCON-1101	Cluster submission in falcon does not create an owned-by edge	Falcon
BUG-30055		Falcon recipe should have rest api	Falcon
BUG-33764		HDP-Select needs to handle user-created directories in /usr/hdp	HDP stack
BUG-33113		HBase-HA master fail over immediately	HBase
BUG-32195		Need to add hbase-slider RU downgrade	HBase
BUG-31943		test_IntegrationTestRegionReplication region not serving	HBase
BUG-30059		In case of secure HBase slider, hbase shell's grant command throws 'ERROR: DISABLED: Security features are not available'	HBase
BUG-32509		test_runIntegrationTestMTTRWMB due to failure in testKillRsHoldingMeta	HBase DistributedLogs
BUG-31353		IntegrationTestRegionReplication region not serving	HBase
BUG-29903		IntegrationTestMTTR slowness	HBase
BUG-31936		test_hbaseBigLinkedList[Loop]TimeoutException	HBase
BUG-27858		HBase HA tests fail with Table Namespace Manager not ready yet	HBase
BUG-27509	HBASE-12464	test_ReplicaCopyTable_verifyCopy fails after trying several times for 'Failed to find location, tableName=hbase:meta, row='	CopyTable[tableone]

Key	Apache JIRA	Summary	Component/s
BUG-27256		Create mr based Hbase Long running application for Rollingupgrade	HBase
BUG-25961		HBase tests failing with org.apache.hadoop.hbase.NoAllMetaRegionsOnlineException: Timed out; 10000ms	HBase
BUG-23619		On Windows, tests fail complaining about Cannot run program /usr/bin/env	HBase
BUG-23981		IOException: Invalid HFile block magic in Windows unit tests	HBase
BUG-22841		HBase test failed with Insufficient permissions (user=hrt_qa@HWQE.HORTONWORKS.COM error	HBase
BUG-21254		hbck tool fixMeta flag not able to recreate deleted entries from hbase:meta table when just one column is deleted	HBase
BUG-23897		HBase REST server continuously restarting in secure windows env	HBase
BUG-33278	HBASE-13239	HBASE grants at specific column level does not work for Groups	HBase
BUG-27656	HBASE-12472, HBASE-13192	IntegrationTestBulkLoad fails on mapreduce	HBase
BUG-27756		test_IntegrationTestMultiGet slowness	HBaseReplicas
BUG-31624		Dropping a 2000+ partition table takes over 10 minutes	Hive
BUG-31671		HS2 concurr tests intermittently fail with http mode on secure Openstack throwing org.apache.hive.service.auth. Kerberos authentication failed	HttpAuthenticationException:
BUG-30984		Hive compaction tests failing	Hive
BUG-33168		hive authorization checks on configs should not be done if authorization is disabled	Hive
BUG-28218		Hive CLI hangs	Hive
BUG-27636		Oracle: intemittent acid_concurrency test failures due to NoSuchLockException	Hive
BUG-26243		CBO : Q51 fails in explain with Failed to breakup Windowing invocations into Groups.	Hive
BUG-27507		Hive CLI returns 0 even in case of failures	Hive
BUG-29427		Turn off Parquet Vectorization until all data	Hive

Key	Apache JIRA	Summary	Component/s
		types work: DECIMAL, DATE, TIMESTAMP, CHAR, and VARCHAR	
BUG-27582		Beeline Client Connections Errors - Obscure errors and exception swallow - 1	Hive
BUG-25064		Hive : Q51 fails in CBO with "Failed to breakup Windowing invocations into Groups. At least 1 group must only depend on input columns"	Hive
BUG-23260		BUG-22878 insert after alter table issue	Hive
BUG-21468		With CBO enabled queries fail with Failed to breakup Windowing invocations into Groups	Hive
BUG-20856		Column names are missing from join expression in Map join with CBO enabled	Hive
BUG-18450	HIVE-7160	Vectorization Udf: GenericUDFConcat, is not supported for String+Long	Hive
BUG-33338		Hiveserver2 in HTTP mode is not applying auth_to_local rules	Hive
BUG-27094		Knox HBase tests skipped for ambari test config	Knox
BUG-27091		Knox YarnSysTest skipped in secure mode for ambari test config	Knox
BUG-27089		Knox YarnSysTest skipped in secure mode for gsInstaller test config	Knox
BUG-28039		Knox Failed Authentication Auditing	Knox
BUG-28031		Knox rolling upgrade system test support for wire encryption	Knox
BUG-25049		Two WebHdfsHaFuncTest are unstable and should be fixed and enabled	Knox
BUG-23889		Windows security scripting should setup Knox to use ActiveDirectory for LDAP	Knox
BUG-20866		Mahout trainnb command fails with -c option	Mahout
BUG-32763		yum cannot find Xvfb for centos06	Oozie
BUG-31426		Oozie failed to start due to "SA" schema not created	Oozie, RelEng
BUG-32026		test_mapred_via_oozie_ha[oozie-nn-rm-5-jobs] FAILED	Oozie

Key	Apache JIRA	Summary	Component/s
BUG-31349		oozie job in running state is causing intermittent timeout failures	Oozie
BUG-33180		oozie logging tests fail on MR framework with postgres 9	Oozie
BUG-28614		Oozie property oozie.authentication.kerberos.principal is being used for authentication of webhdfs endpoint	Oozie
BUG-23655		Oozie stuck on mapreduce job because workflow state isn't changing from RUNNING to KILLED	Oozie
BUG-23527		Oozie config referred to versioned hadoop config	Oozie
BUG-21166		testGetCaseInsensitiveDynCol fails intermittently throwing AssertionError	Phoenix
BUG-29261		Tests are failing with ERROR 1200 and then raising an error due to jobManager rejection	Phoenix
BUG-31497		[RollingUpgrade] Phoenix intermittently fails to return every column on sql queries	Phoenix
BUG-26661		Windows runs failed with "region isn't co-located with its data" exception	Phoenix
BUG-28011		testInnerJoinWithPreFilters and testJoinOverSaltedTables are failing with SequenceNotFoundException exception	Phoenix
BUG-27893		testCurrentDate and testCurrentTime are failing due to NTP on windows server	Phoenix
BUG-26428		XaAdmin and XaAgent jobs for Debian failing on Unix clusters.	Ranger
BUG-33056		{start,stop}-thriftserver equivalent script on Windows is missing	Spark
BUG-33054		Spark examples are packaged differently on Windows	RelEng, Spark
BUG-33335		spark make-distribution.sh doesn't evaluate SPARK_HIVE on OSX	Spark
BUG-27769		[Rolling upgrade] [Test issue] Sqoop Long running job finishes before upgrade of all component finishes	Sqoop

Key	Apache JIRA	Summary	Component/s
BUG-32688		kafka/api/OffsetRequest is not packaged in storm-kafka-example jar	Storm
BUG-27513		Storm-slider default config has ganglia error in supervisor.log	Storm
BUG-29730		Storm Hive,Hbase, Hadoop connector should support user specified UGI in non-secure env	Storm
BUG-31522		You cannot search by the name of 'Queue'	Tez
BUG-31565		When you kill a job in the middle of running it, it Diagnostics shows 'killed/failed due to:null'.	Tez
BUG-26419		Tez application fails with diff client and AM version	Tez
BUG-27562		WordCount on Tez with non-nightly failed with NPE	Tez
BUG-24831		TEZ windows test framework needs to respect \${fs.defaultFS} in tez configs.	Tez
BUG-25194		AM running out of memory for partitioned bucketed tables	Tez
BUG-25380		For hive to work with Virtual columns and partition names, we need information available within a record-reader	Tez
BUG-24518		Tez needs to support Application Tag	Tez
BUG-27194		[YARN-2821] Failures encountered on Distributed Shell on windows	YARN
BUG-27590		Distributed shell job does not execute python script in Windows	YARN
BUG-22807		Non-secure slider app processes run with yarn user.	YARN
BUG-22522		MR Job failed when run via oozie and NN and RM were being killed	YARN
BUG-27632		Add feature of retry in zookeeper client	Zookeeper
BUG-30784	ZOOKEEPER-1952	zookeeper.log.file property is not respected; log output goes only to the zookeeper.out	Zookeeper
BUG-19255		Uninstalling the ZooKeeper RPM does not remove all ZooKeeper directories	Zookeeper