

Hortonworks Data Platform

HDP-2.2.0 Release Notes

(Dec 2, 2014)

Hortonworks Data Platform: HDP-2.2.0 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 2.2.0 Release Notes	1
1.1. Apache Component Support	1
1.1.1. Supported Apache Component Version Information	1
1.1.2. Third-Party Tools	2
1.1.3. Unsupported components	2
1.2. Unsupported Features	2
1.3. Behavioral Changes	2
1.3.1. HDFS Behavioral Changes	2
1.3.2. Hive Behavioral Changes	3
1.3.3. Oozie Behavioral Changes	5
1.3.4. Pig Behavioral Changes	5
1.4. Tech Previews in This Release	5
1.5. What's New in this Release	6
1.6. What's Fixed in HDP 2.2	8
1.7. Apache Patch Information	32
1.7.1. DataFu 1.2.0	32
1.7.2. Falcon 0.6.0	32
1.7.3. Flume 1.5.2	32
1.7.4. Hadoop Common/HDFS 2.6.0	34
1.7.5. HBase 0.98.4	34
1.7.6. Hive 0.14.0	41
1.7.7. Knox 0.6.0	42
1.7.8. Mahout 0.9.0	42
1.7.9. Oozie 4.1.0	45
1.7.10. Phoenix 4.2.0	46
1.7.11. Pig 0.14.0	46
1.7.12. Ranger 0.4.0 (formerly XASecure)	46
1.7.13. Sqoop 1.4.5	47
1.7.14. Storm 0.9.3	47
1.7.15. Tez 0.5.2	47
1.7.16. Zookeeper 3.4.6	48
1.8. Known Issues for HDP 2.2	48
1.8.1. Falcon Known Issues	48
1.8.2. Hadoop Common/HDFS Known Issues	49
1.8.3. HBase Known Issues	50
1.8.4. HDFS Known Issues	50
1.8.5. Hive Known Issues	50
1.8.6. Knox Known Issues	52
1.8.7. Oozie Known Issues	52
1.8.8. Phoenix Known Issues	53
1.8.9. Pig Known Issues	53
1.8.10. Ranger (formerly XA Secure) Known Issues	53
1.8.11. Sqoop Known Issues	54
1.8.12. YARN Known Issues	54
1.8.13. ZooKeeper Known Issues	55

1. HDP 2.2.0 Release Notes

This document provides you with the latest information about the HDP 2.2.0 release and its product documentation.

1.1. Apache Component Support

This section provides supported Apache component version information, a list of Apache components not supported on HDP for Windows, and third-party tool information.

1.1.1. Supported Apache Component Version Information

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.

Official Apache versions for HDP 2.2.0.

- Apache Hadoop 2.6.0
- Apache Ranger 0.4.0
- 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
- Apache Knox 0.5.0
- Apache Mahout 0.9.0
- Apache Oozie 4.1.0
- Apache Pig 0.14.0
- Apache Phoenix 4.2.0
- Apache Sqoop 1.4.5
- Apache Slider 0.60.0
- Apache Tez 0.5.2
- Apache Storm 0.9.3
- Apache Zookeeper 3.4.6

1.1.2. Third-Party Tools

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

1.1.3. Unsupported components

The following components are supported on HDP for Linux, but are not currently supported on HDP for Windows.

- Apache Accumulo 1.6.1
- Hue 2.6.1
- Apache Kafka 0.8.1.1

1.2. Unsupported 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

1.3. Behavioral Changes

Behavioral changes denote a marked change in the code from the previously released version to this version of HDP 2.2. Review all behavioral changes for a component to ensure that you are aware of how components have changed in this release.

1.3.1. HDFS Behavioral Changes

[HDFS-6376](#) distcp can copy data between HA clusters. You can use the new configuration property `dfs.internal.nameservices` to explicitly specify the name services belonging to the local cluster, while continue using the configuration property `dfs.nameservices` to specify all the name services in the local and remote clusters.

For example, if you perform distcp between two clusters, for example c1 and c2 and both of which run in the HA mode, make the following configuration changes on the cluster that runs the distcp job:

1. Set `dfs.nameservices` to both c1 and c2.
2. Modify the configuration to include the settings for both clusters. See the comments in [HDFS-6376](#) for more details.
3. Set `dfs.internal.nameservice` to the nameservice id that corresponds to the cluster that runs the distcp job.

1.3.2. Hive Behavioral Changes

[HIVE-7397](#) Hive has changed the following default settings:

```
hive.fetch.task.conversion - from 'minimal' to 'more'
hive.fetch.task.conversion.threshold - from '-1' to '1073741824' (1 GiB)
```

For example, queries such as:

```
select * from foo_table limit 1;
```

which would previously have executed without a MapReduce stage to now require a MapReduce stage. The change to these settings allows more types of queries to be evaluated locally, including queries with "where" clauses. The threshold is applied to the input size of the query (raw table size). Because 'more' allows filtered queries, it's prudent to set a limit on the input size.

[HIVE-6432](#): Hive 0.14 removes all `org.apache.hcatalog.*` classes in favor of `org.apache.hive.hcatalog.*` classes.

The following previously deprecated features were also removed in this release as part of this change:

- `HdfsAuthorizationProvider` - the `HiveStorageBasedAuthorizationProvider` replaces it.
- `HBaseHCatStorageHandler` - Hive's native `HBaseStorageHandler` is now directly usable from `HCatalog`.

[HIVE-6455](#) If you store Hive data in a columnar format, such as ORC, verify that the `hive.optimize.sort.dynamic.partition` property is set to true. It was the default value for Hive 0.13.x, but is now disabled by default in Hive 0.14.x. Enabling dynamic partitions improves Hive performance when using columnar formats.

```
<property>
  <name>hive.optimize.sort.dynamic.partition</name>
  <value>true</value>
</property>
```

([HIVE-5940](#), [HIVE-6252](#)) Added `SHOW GRANT` and `SHOW ROLE GRANT`. Users will be able to list only their own privileges on a table/role unless they are in an admin role. For `grant/revoke` statements, if a user is granted a privilege `WITH GRANT OPTION` on a table or view, then the user can also grant that privilege other users.

[HIVE-6972](#) The JDBC URI for `HiveServer2` has the form:

```
jdbc:hive2://<host1>:<port1>,<host2>:<port2>/dbName;sess_var_list?
hive_conf_list#hive_var_list
```

Ideally, `sess_var_list` is supposed to be used to pass parameters that will be used within the JDBC Driver. However, some of the http mode parameters were being passed in the `hive_conf_list` which is fixed now (in a backward-compatible way).

Additionally, JDBC Driver has a naming convention for the parameters in the JDBC URI which were not followed to the `sas.qop` parameter. It is renamed it in a backward compatible way.

- HTTP Mode parameters have changed. As you can see in the following examples, `hive.server2.transport.mode` has been renamed to `transportMode`, `hive.server2.thrift.http.path` has been renamed to `httpPath` and both have been moved from `hive_conf_list` to `sess_var_list`.

- Old example URI: `jdbc:hive2://<host>:<port>/<db>;user=<username>;password=<password>?hive.server2.transport.mode=http;hive.server2.thrift.http.path=<http_endpoint>`

- New example URI: `jdbc:hive2://<host>:<port>/<db>;user=<username>;password=<password>;transportMode=http;httpPath=<http_endpoint>hive.server2.logging.operation.enabled=false`

- Sasl qop renamed to `saslQop`:

- Old example URI: `jdbc:hive2://<host>:<port>/<db>;principal=<hiveserver2_kerberos_principal>;sas.qop=<qop_value>`

- New example URI: `jdbc:hive2://<host>:<port>/<db>;principal=<hiveserver2_kerberos_principal>;saslQop=<qop_value>`

All changes are backward-compatible and print a deprecation warning message like the following:

```
0: jdbc:hive2://localhost:10007> !connect
jdbc:hive2://localhost:10007//;sas.qop=auth?hive.server2.transport.mode=h
ttp;hive.server2.thrift.http.path=cliservice user user
org.apache.hive.jdbc.HiveDriver
Connecting to
jdbc:hive2://localhost:10007//;sas.qop=auth?hive.server2.transport.mode=h
ttp;hive.server2.thrift.http.path=cliservice
14/10/07 16:22:24 INFO jdbc.Utils: Supplied authorities: localhost:10007
14/10/07 16:22:24 WARN jdbc.Utils: ***** JDBC param deprecation *****
14/10/07 16:22:24 WARN jdbc.Utils: The use of sas.qop is deprecated.
14/10/07 16:22:24 WARN jdbc.Utils: Please use saslQop like so:
jdbc:hive2://<host>:<port>/dbName;saslQop=<qop_value>
```

```
14/10/07 16:22:24 WARN jdbc.Utills: ***** JDBC param deprecation *****
14/10/07 16:22:24 WARN jdbc.Utills: The use of hive.server2.transport.mode
is deprecated.
14/10/07 16:22:24 WARN jdbc.Utills: Please use transportMode like so:
jdbc:hive2://<host>:<port>/dbName;transportMode=<transport_mode_value>
14/10/07 16:22:24 WARN jdbc.Utills: ***** JDBC param deprecation *****
14/10/07 16:22:24 WARN jdbc.Utills: The use of
hive.server2.thrift.http.path is deprecated.
14/10/07 16:22:24 WARN jdbc.Utills: Please use httpPath like so:
jdbc:hive2://<host>:<port>/dbName;httpPath=<http_path_value>
Connected to: Apache Hive (version 0.14.0-SNAPSHOT)
Driver: Hive JDBC (version 0.14.0-SNAPSHOT)
Transaction isolation: TRANSACTION_REPEATABLE_READ
```

Be aware that `org.apache.hcatalog.pig.HCatLoader` was deprecated in Hive 0.13.x and removed in Hive 0.14.x. You must now use `org.apache.hive.hcatalog.pig.HCatLoader`.

1.3.3. Oozie Behavioral Changes

The Oozie system environment must have JDK 1.6 or higher. And on the machine where you the run oozie command line, set `$JAVA_HOME` to a java version of 1.6 or higher.

Calling log4j has changed in Oozie for HDP 2.2. If you are using custom configurations or Ambari to install, add the following line to `oozie-log4j.properties`:

```
log4j.appender.oozie.layout.ConversionPattern=%d{ISO8601} %5p %c{1}:%L -
SERVER[${oozie.instance.id}] %m%n
```

where `${oozie.instance.id}` is determined by Oozie, automatically.

1.3.4. Pig Behavioral Changes

Be aware that `org.apache.hcatalog.pig.HCatLoader` was deprecated in Hive 0.13.x and removed in Hive 0.14.x. You must now use `org.apache.hive.hcatalog.pig.HCatLoader`.

1.4. 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 Encryption
- HDFS read/write storage tier
- [HDFS-6932](#) HDFS Blocking Mover/Balancer from moving files/blocks to/from transient storage.
- Falcon Data Lineage
- Falcon Graph View of Dependencies
- Phoenix support for local secondary indexes
- YARN support for Docker

1.5. What's New in this Release

HDP 2.2 includes the following new features:

- Falcon
 - Authorization
 - Lineage Enhancement
 - HCat Replication / Retention
 - Archive to Cloud
- Flume
 - Flume streaming to Hive for secure and unsecure clusters
- HBase
 - HBase HA: Timeline-consistent replicas with realtime replication
 - HBase block cache compression
 - HBase on YARN via Slider
- HDFS
 - Heterogenous Storage: Support for SSD Tier
 - Heterogenous Storage: Support for Archival Tier
 - Operating secure DataNode without requiring root access
 - AES support for faster wire encryption
- Hive
 - Support for SQL Transactions with ACID semantics
 - Support for SQL Temporary Tables

- Better optimize queries using Cost Based Optimizer
- Add performance improvements for various queries
- Security additions such as Grant, Revoke with a choice of using Native Security or Ranger security with both integrated completely in Hive
- Knox
 - Support for HDFS HA
 - Installation and configuration with Apache Ambari
 - Service-level authorization with Apache Ranger
 - YARN REST API access
- Oozie
 - Oozie HA on secure clusters
- Phoenix
 - Subquery support
 - Robust secondary indexes
 - Build secondary index during bulk import
- Pig
 - Pig on Tez
 - Including DataFu for use with Pig
- Ranger
 - Storm Authorization and Auditing
 - Knox Authorization and Auditing
 - Deeper integration with HDP stack: Hive Auth API support for grant/revoke commands, grant/revoke commands in HBase, and Windows support
 - REST API's for policy manager, local audit log storage in HDFS, and support for Oracle DB for policy store and audits
- Sqoop
 - Support for all Hive types via HCatalog import/export
 - Support for multiple partition keys
 - Integration with Hadoop Credential Management Framework
- Storm

- Storm clusters can now be provisioned via Ambari
- Storm can now run on YARN based clusters using Apache Slider
- Storm now supports kerberos based authentication and pluggable authorization
- Pre-built Spouts for JMS, HBase lookups and Bolts for Kafka, Hive.
- Real-time visualization of running topologies and its associated metrics
- REST APIs for topology stats
- Tez
 - Tez Debug Tooling & UI
- YARN
 - Support long running services: handling of logs, containers not killed when AM dies, secure token renewal, YARN Labels for tagging nodes for specific workloads
 - Support for CPU Scheduling and CPU Resource Isolation through CGroups
 - Work-preserving restarts of ResourceManager and NodeManager
 - Support node labels during scheduling
 - Global, shared cache for application artifacts
 - REST API for YARN application submission and termination
 - Application Timeline Server is supported in a Secure (Kerberized) cluster

1.6. What's Fixed in HDP 2.2

Apache JIRA	Key	Summary
ARGUS-182	BUG-27686	Date format in Audit log is incorrect
CALCITE-448 HIVE-8656	BUG-26264	[HIVE-8656] CBO: auto_join_filters fails
DATAFU-58 DATAFU-59	RMP-1885	RMP-1508 Package and certify Apache DataFu Pig
FALCON-166 FALCON-474 FALCON-591 FALCON-640	BUG-20294	Add Bulk APIs to drive the dashboard needs
FALCON-229	BUG-17238	Enable SSL for Falcon REST API
FALCON-303	BUG-12650	Falcon API changes required for Mirroring UI
FALCON-434	BUG-17533	champlain CI: falcon doesn't compile with hive trunk

Apache JIRA	Key	Summary
FALCON-466	BUG-20283	REST APIs must add the entity owner as an implicit filter
FALCON-469	BUG-20295	Add ability to filter by tags in REST API
FALCON-470	RMP-1699	Enhancements to tagging capabilities
FALCON-613		
FALCON-482	BUG-19015	Concurrently update/delete a process fails in falcon
FALCON-598	BUG-21520	ProcessHelper throws NPE if the process has no inputs OR no outputs defined
FALCON-669	BUG-22213	Falcon should have a default value for workflow.execution.listeners configuration
FALCON-678	BUG-25181	Falcon's default port is 15443 if TLS is enabled, 15000 otherwise.
FALCON-678	BUG-25180	Falcon's default port is 15443 if TLS is enabled, 15000 otherwise.
FALCON-728	BUG-22317	Linux Falcon Unit Test Failures
FALCON-740	BUG-22668	Scheduled feed or process cannot be deleted
FALCON-777	BUG-24148	champlain falcon unit test failures
FALCON-789	BUG-24629	Retention failed during context serialization
FALCON-800	BUG-24317	Falcon lineage is not working
FALCON-819	BUG-24389	Problem with invalid Catalog server while submitting a cluster
FALCON-831	BUG-25639	Operation on non existing entity throws internal server error
FALCON-839	BUG-26064	Falcon ACL validation is not working
FALCON-840	BUG-26106	SOURCECLUSTER field in -filterBy parameter doesn't work for /api/instance/list
FALCON-851	BUG-26595	In falcon - process with bad acl can be submitted
FALCON-857	BUG-26594	In falcon fetching status of deleted feed throws internal server error
FALCON-858	BUG-26578	In falcon listing processes with custom filter fails
FALCON-859	BUG-26590	falcon entities/summary/feed/ endpoint gives internal server error
FALCON-862	BUG-26367	/list/entities with tag also return entities without tags
FALCON-874	BUG-27261	In falcon instance/list endpoint orderBy=STATUS is case-sensitive
FALCON-875	BUG-27254	Unexpected behaviour with entity summary endpoint
FALCON-880	BUG-26394	hive job via oozie failed read timeouts to the metastore
FALCON-889	BUG-27495	Windows azure replication fails with we can't create it using anonymous credentials
FALCON-902	BUG-28027	Documentation of falcon parameters

Apache JIRA	Key	Summary
FLUME-1734	BUG-14846	Hive Sink Unit test needs a way to interact with Hive and check streamed data
FLUME-2442	RMP-1570	RMP-1511 encrypt ssl truststore / keystore passwords within flume config file
FLUME-2451	BUG-23556	Flume HDFS Sink needs to be able to handle namenode restart
FLUME-2491	BUG-20730	test_Basic failed with org.apache.flume.FlumeException: Unable to load sink type: org.apache.flume.sink.solr.morphline.MorphlineSolrSink, class: org.apache.flume.sink.solr.morphline.MorphlineSolrSink
FLUME-2508	BUG-25196	Windows: Hive Streaming table has extra carriage feed (CR) in the last column when using Spooling directory source
FLUME-2530	BUG-24367	Review Coverity scan for Flume
HADOOP-10345	RMP-1498	Sanitize the the inputs (groups and hosts) for the proxyuser configuration
HADOOP-10416	BUG-24211	Hadoop authentication should not switch to anonymous user when a clearly defined user has an expired token
HADOOP-10472	BUG-16293	KerberosAuthenticator writes security logs on the client
HADOOP-10503	BUG-16795	update junit dependency for all components to 4.11
HADOOP-10508	BUG-16985	EBay Code Review - Please review and commit HADOOP-10508
HADOOP-10522 HADOOP-10527	BUG-15796	sigsegv in mapred history server due to segfault in JniBasedUnixGroupsMapping
HADOOP-10966	BUG-21241	Hadoop build Failure on Windows platform only.
HADOOP-10995	BUG-21709	[HADOOP-10995] HBase cannot run correctly with Hadoop trunk
HADOOP-11063	BUG-22197	KMS breaks all champlain windows installs
HADOOP-11068	BUG-26121	oozie webconsole broken in secure clusters
HADOOP-11109	BUG-23294	[HADOOP-11109] champlain linux bustage on hadoop-common-project/hadoop-kms/src/site/apt/index.apt.vm
HADOOP-11177	BUG-23750	[HADOOP-11177]MR tar ball for distributed cache should not include version number after unzip.
HADOOP-11181 YARN-2656	BUG-23891	[HADOOP-11181][YARN-2656] YARN REST call fails via proxy user in both secure windows and secure linux env's
HADOOP-11207 HIVE-8384 YARN-2676	BUG-24017	[HADOOP-11207][YARN-2676] Webhcat tests failing with org.apache.hadoop.security.authentication.client.AuthenticationException

Apache JIRA	Key	Summary
HADOOP-11217 HDFS-7391	BUG-27680	[HADOOP-11217] [HDFS-7391] Pull in SSLv2Hello fixes into Champlain
HADOOP-11286	BUG-27363	Pull HADOOP-11286 into Champlain
HADOOP-7657	BUG-20082	Included support for LZ4 compression in Hadoop
HADOOP-8069	BUG-20184	Have the tcpnodelay flag set to true in the hadoop configurations?
HADOOP-8943	RMP-1796	Hadoop should support groups in a composite from both Linux and LDAP
HADOOP-9640	RMP-1480	RPC Congestion Control with FairCallQueue
HADOOP-9888	BUG-18256	Apache HADOOP-9888 KerberosName static initialization gets default realm, which causes big slow-down in non-secure deployment.
HBASE-10123	BUG-16257	Hbase master fails to start due to BindException
HBASE-10641	RMP-1632	RMP-1507 Bucket cache improvements
HBASE-10785	BUG-15466	HBase Bloomberg HA Replica Integration Test get/multi-get fails
HBASE-10817 HBASE-10818	BUG-12398	Create bulkload integration Test with replica option
HBASE-11137 HIVE-6584	RMP-1895	RMP-1507 Allow Hive queries over snapshotted HBase tables
HBASE-11149	BUG-17710	Wire encryption is broken
HBASE-11331	RMP-1631	RMP-1507 Block cache compression
HBASE-11344	RMP-1577	RMP-1507 HBase to adhere to security checklist
HBASE-11367 HBASE-11511 HBASE-11512 HBASE-11567 HBASE-11568 HBASE-11569 HBASE-11571 HBASE-11842	RMP-1628	RMP-1507 HBase HA: Support < 5s staleness reading from secondaries
HBASE-11569 HBASE-12274	BUG-24676	ChaosMonkey IntegrationTestIngest fails with balancer and multithread reader errors
HBASE-11574 HBASE-12317 HBASE-12335	BUG-20978	Integration test class IntegrationTestRegionReplicaPerf fails complaining about region replication enabled performance not being as fast enough as without replication
HBASE-11580	RMP-1627	RMP-1507 HBase HA: Proper recovery from failed secondary
HBASE-11716	BUG-25867	IntegrationTestIngestWithVisibilityLabels failed with verification error

Apache JIRA	Key	Summary
HBASE-11886	BUG-21798	Regression: In case of secured cluster, user with only 'C' create permission was able to insert new data prior to Champlain, but not anymore in Champlain
HBASE-11903	BUG-22290	test_ReplicaSplitTable fails as the expected replica number after compaction is not the same as actual replica number
HBASE-11908 HBASE-12561	BUG-24651	IntegrationTestRegionReplicaReplication fails with "Load failed with error code 1"
HBASE-11940	BUG-20046	Provide snapshot-all and restore-all scripts to support rolling upgrade
HBASE-11972	BUG-22601	Hbase secure model is not working
HBASE-11973	BUG-22148	Mapreduce job cannot commit due to FileNotFoundException of container_token
HBASE-11974	BUG-22195	In Champlain when a disabled table is being copied, it throws NotServingRegionException instead of TableNotEnabledException
HBASE-11976	BUG-25958	ChaosMonkey tests failed to write
HBASE-11989	BUG-23028	IntegrationTestLoadAndVerify cannot be configured anymore on distributed mode
HBASE-12040	BUG-23574	Scan with filters has regression in performance
HBASE-12052	BUG-23095	hbase champlain linux unittest failure
HBASE-12052	BUG-21357	CsvBulkLoadTool gets blocked if executed with hrt_qa user
HBASE-12052	BUG-25168	BUG-22192 TestLoadIncrementalHFilesUseSecurityEndPoint fails
HBASE-12055 HBASE-12074	BUG-22303	BUG-22192 Windows Hbase Unit Test Failures
HBASE-12056	BUG-22281	IntegrationTestRegionReplicaReplication fails
HBASE-12072	BUG-23370	Integration tests stuck on windows and linux
HBASE-12072	BUG-23531	Hbase tests get stuck when master is not running
HBASE-12099	BUG-22418	copyTable tests fail with error java.lang.NoSuchMethodError
HBASE-12099	BUG-22603	BUG-22192 TestScannerModel fails consistently
HBASE-12123	BUG-23753	BucketCache assertion error
HBASE-12162	BUG-23956	Hbase HA: LoadTestTool fails immediately after hbase master is killed
HBASE-12189	BUG-24369	Review Coverity scan result for HBase
HBASE-12277	BUG-24508	Integration Tests on secure cluster has Error in LoadIncrementalHFiles and Error in running command-line tool

Apache JIRA	Key	Summary
HBASE-12278	BUG-25170	BUG-22192 TestSecureLoadIncrementalHFilesSplitRecovery fails due to TableNotFoundException
HBASE-12319	BUG-25580	Inconsistencies during region recovery due to close/open of a region during recovery
HBASE-12356	BUG-25848	phoenix champlain Linux unit-test failures
HBASE-12367	BUG-23957	IntegrationTestTimeBoundedMultiGetRequestsWithRegionRecovery calls Linux specific commands
HBASE-12377	BUG-26051	hbase:meta not online for IntegrationTestManyRegions
HBASE-12398	BUG-26242	Race condition in region opening
HBASE-12402	BUG-26432	test_hbaseLoadTestTool gets stuck complaining about AccessDeniedException
HBASE-12449	BUG-26592	Multiple tests are failing because of System time goes back frequently on SLES
HBASE-12478	BUG-27844	HBASE-12478 HBASE-10141 and MIN_VERSIONS are not compatible
HBASE-12485	BUG-25162	"Online" Sequence number value of a region at the time of opening is unexpected
HBASE-12485	BUG-26043	test_replicacopytable_verifycopytable test fails with IOException error: 'Failed to create SeqId file' in RS log
HBASE-12485	BUG-25870	test_ReplicaCopyTable_verifyCopyTable[tableone] fails with org.apache.hadoop.fs.FileAlreadyExistsException in RS log
HBASE-9445	BUG-13215	HBase Snapshot Input Format IT Cannot Run with Default Number of Regions
HDFS-172	BUG-7101	HDFS-172: uploading file larger than the spaceQuota limit results in a partial file.
HDFS-4793		
HDFS-2856	BUG-18711	HDFS DataNode must not require to be started as ROOT (HDFS-2856)
HDFS-2882	BUG-15109	HDFS-2882: Receiving NPE Errors when starting datanode due to a mount not being accessible
HDFS-4211		
HDFS-5147	BUG-8486	dfsadmin commands should interact with the active namenode in ha setup
HDFS-6507		
HDFS-573	BUG-12699	libHDFS native library for Windows
HDFS-6376	BUG-21172	Distcp data between two HA clusters fail with error
HDFS-6455	BUG-18257	NFS: Exception should be added in NFS log for invalid separator in allowed.hosts (HDFS-6455)
HDFS-6570	BUG-19999	add api that enables checking if a user has certain permissions on a file
HDFS-6581	BUG-18704	Support memory as storage tier

Apache JIRA	Key	Summary
HDFS-6584	BUG-23196	Umbrella jira for tracking in-progress Archival Storage work
HDFS-6678	BUG-14645	MiniDFSCluster leaves CacheReplicationMonitor thread running causing heap thrashing and eventual failure of Oozie unittests
HDFS-6689	BUG-19268	NFS: can't access file under directory with 711 access right as other user
HDFS-6715	BUG-16913	[HDFS-6715] webhdfs wont fail over when it gets java.io.IOException: Namenode is in startup mode
HDFS-6830	BUG-21033	NameNode must handle block migration between storages on the same DN
HDFS-6836	BUG-20769	HDFS INFO logging is verbose & uses file appenders
HDFS-6898	BUG-20738	HDFS Fails to write 1 replica data when local disks fill up
HDFS-6932	BUG-23481	Block Mover/Balancer from moving file blocks on transient storage
HDFS-6934	BUG-24729	Disable checksum when writing to Memory
HDFS-6981	BUG-22233	Support DN layout changes with rolling upgrade
HDFS-6988	BUG-23480	Make RAM Disk eviction thresholds configurable
HDFS-6996	BUG-22222	Snapshot diff report has IndexOutOfBoundsException.
HDFS-6998	BUG-22267	Windows: warning message 'ssl.client.truststore.location has not been set' gets printed for hftp command
HDFS-7073	BUG-22562	Inter Cluster Distcp fails
HDFS-7129	BUG-23545	Adding metrics to track usage of memory for writes
HDFS-7167	BUG-23764	NPE while moving file between states with Mover
HDFS-7172	BUG-23526	HDFS test files checked out of git with incorrect line endings.
HDFS-7178	BUG-24071	Additional unit tests for single replica writes in disk full situation
HDFS-7230	BUG-24189	HDFS Rolling downgrade - need to figure out how to use HDFS's -downgrade flag
HDFS-7237	BUG-24416	hadoop namenode -rollingUpgrade should print correct Error message
HDFS-7305	BUG-25722	[HDFS-7305] HDFS HA: Seeing NPE when running slive test
HDFS-7340	BUG-26796	[HDFS-7340] rollingUpgrade prepare command does not have retry cache support
HDFS-7359	BUG-25781	When in HA mode, dfs.namenode.secondary.http-address is still being loaded and used

Apache JIRA	Key	Summary
HDFS-7359	BUG-26809	In HA mode, if dfs.namenode.secondary.http-address has value "null" it prevents namenode to start properly and generates mixmatch txid
HDFS-7364	BUG-26963	HDFS Rebalance Always Shows 0 Bytes Moved
HDFS-7382	BUG-27327	NullPointerException in SaslDataTransferServer.receive
HDFS-7383	BUG-27326	NullPointerException in DataNode.requestShortCircuitFdsForRead
HDFS-7387	BUG-27281	NFS: File gets downloaded partially
HIVE-4627	BUG-10529	Add support for total ordering of output
HIVE-4629	BUG-21620	Beeline should show HS2 logs
HIVE-7615		
HIVE-538	RMP-1370	Produce a single JDBC connector JAR to HiveServer2
HIVE-5775	BUG-21859	Clone hive cbo branch in hwx github repo
HIVE-6507	BUG-14013	hcatalog should automatically work with new tableproperties in ORC
HIVE-6549	BUG-14969	remove hcatalog/bin/hive-config.sh
HIVE-6626	BUG-16596	Hive does not expand the DOWNLOADED_RESOURCES_DIR path
HIVE-6715	BUG-25020	HiveServer2 user name not being passed in NOSASL authentication setting
HIVE-6928	BUG-20173	Beeline truncates output rows longer than 170 characters
HIVE-7052	BUG-17521	Optimizing tez split calculation time
HIVE-7068	RMP-1886	RMP-1512 Hive / Accumulo Connector
HIVE-7071	BUG-15733	Schema evolution is broken on Tez
HIVE-7158	RMP-1965	More intelligent reducer allocation
HIVE-7159	RMP-1966	Filter nulls early in equi-joins
HIVE-7219	BUG-18933	upload/checkin unrolling code
HIVE-7231	BUG-18934	fix stripe padding (64M - 256M blocks)
HIVE-7250	BUG-17231	Problems loading wide tables into ORC
HIVE-7295	BUG-20374	automatically set ownership of new files/dirs to Administrators group instead of individual user
HIVE-7297	BUG-20372	org.apache.hadoop.hive ql.WindowsPathUtil is not easy to be reused across different hive projects
HIVE-7324	BUG-19033	Unit test framework addition to allow import of metadata/explain only queries
HIVE-7342	BUG-16918	SQL std auth - should be possible to configure it using hive configuration files alone
HIVE-7361	RMP-1979	Changes in Hive needed to support XA Secure integration

Apache JIRA	Key	Summary
HIVE-7416		
HIVE-7451		
HIVE-7488		
HIVE-7472	BUG-20395	Support for additional filetypes in webhcat java client api
HIVE-7533	BUG-17307	support SQL std authz from hive cli
HIVE-7579	BUG-15881	Error messaging for different authorization failures is not informative
HIVE-7616	BUG-20424	Use cardinality estimate for hive.hashtable.initialCapacity instead of using nextHighestPowerOfTwo(100000)
HIVE-7620	BUG-20697	Hive metastore fails to start in secure mode due to "java.lang.NoSuchFieldError: SASL_PROPS" error
HIVE-7648	BUG-21050	HiveAuthorizer.checkPrivileges(hiveOpType=CREATEDATABASE, CREATETABLE): parameters should include database name/table name respectively
HIVE-7648	BUG-21051	HiveAuthorizer.checkPrivileges(hiveOpType=DROPINDEX): parameters should include database name and table name
HIVE-7648	BUG-21049	HiveAuthorizer.checkPrivileges(hiveOpType=SWITCHDATABASE): parameters should include database name
HIVE-7695	BUG-21850	incorrect count(*) results after insert appending data
HIVE-7877	BUG-22886	BUG-22878 Insert/values with negative values causes a parse exception
HIVE-7914	BUG-23833	Simplify join predicates for CBO to avoid cross products
HIVE-7943	BUG-27427	hive.security.authorization.createtable.owner.grants is ineffective with Default Authorization
HIVE-7950	BUG-22175	AccumuloStorageHandler doesn't work with Hive on Tez
HIVE-7984		
HIVE-7971	BUG-22037	Not able to update decimal type on partitioned data in Hive 13
HIVE-7985	BUG-23832	With CBO enabled cross product is generated when a subquery is present
HIVE-8021	BUG-25413	ctas/insert into queries + cbo
HIVE-8042	BUG-23168	Hive queries execute MoveTasks sequentially which slows down performance significantly in large queries.
HIVE-8045	BUG-22551	External table creation fails when metastore.authorization is set to use MetaStoreAuthzAPIAuthorizerEmbedOnly and sql std auth is enabled from hive-site.xml
HIVE-8045	BUG-22664	add jar, dfs fails from hive cli when hive.security.authorization.manager

Apache JIRA	Key	Summary
		is set to org.apache.hadoop.hive.ql.security.authorization.plugin.sql in hive-site.xml
HIVE-8052	BUG-22528	Vectorization: min() on TimeStamp datatype fails with error "Vector aggregate not implemented: min for type: TIMESTAMP"
HIVE-8062	BUG-22810	Stats collection for columns fails on a partitioned table with null values in partitioning column
HIVE-8078	BUG-22902	ORC Delta encoding corrupts data when delta overflows long
HIVE-8104	BUG-22883	BUG-22878 INSERT VALUES statement runs in vectorized mode NPEs
HIVE-8105	BUG-22880	BUG-22878 add support NULL literal in INSERT VALUES statement
HIVE-8107	BUG-22879	BUG-22878 Improve error msg for UPDATE of invalid table
HIVE-8149	BUG-23092	ReduceSinkDeDup optimization needs to be turned off for ACID
HIVE-8152	BUG-23094	BUG-22878 Updates don't allow expressions in set clause
HIVE-8167	BUG-23169	Champlain compile broken on hive
HIVE-8170	BUG-23195	Hive fails to run on Windows
HIVE-8191	BUG-22882	BUG-22878 attempt to insert/update/delete a non acid table should raise an error
HIVE-8203	BUG-16279	Valid transaction list acquired before locks in Driver.java
HIVE-8203	BUG-23522	BUG-22878 ACID operations result in NPE when run through HS2
HIVE-8203	BUG-23215	table_insert_auth_role.q, table_insert_auth_user.q and HBase_5 tests are failing because of NPE
HIVE-8212	BUG-23111	alter view query doesn't work from hcatalog anymore
HIVE-8221	BUG-22665	Webhcat tests for restricted permissions failing
HIVE-8231	BUG-23923	Unable to read data when inserted into empty table via HS2
HIVE-8257	BUG-22761	WebHcat component showing NoSuchMethodError: Text ClientRMProxy.getRMDelegationTokenService(Configuration)
HIVE-8258	BUG-16281	Cleaner needs to track readers reading its files rather than acquire a lock
HIVE-8263	BUG-23829	CBO : TPC-DS Q64 is item is joined last with store_sales while it should be first as it is the most selective
HIVE-8272	BUG-23734	Query with particular decimal expression causes NPE during execution initialization
HIVE-8277	BUG-23668	IP address returned by HiveAuthzContext.getIpAddress() has a leading "/"

Apache JIRA	Key	Summary
HIVE-8279	BUG-20921	drop table proceeds without authorization check if "db.table" syntax is used
HIVE-8280	BUG-23828	CBO : When filter is applied on dimension table PK/FK code path is not in effect.
HIVE-8290	BUG-23922	Need to add table property to indicate that a table is transactional
HIVE-8290	BUG-23805	Add 'transactional' property to ACID tables in IUD system tests
HIVE-8311	BUG-23924	Txn list information being encoded too late
HIVE-8318	BUG-23982	Queries hit RuntimeException: cannot find field with constant folding always false
HIVE-8332	BUG-23855	IUD tests are failing because of NPE while running through HS2
HIVE-8341	BUG-24044	valid transaction list causes failure in transform clause (aka streaming aka python udf)
HIVE-8367	BUG-24319	delete produces records in wrong order in some cases
HIVE-8367	BUG-24719	BUG-22878 add system test for BUG-24319
HIVE-8368	BUG-24321	compactor writing deletes into base files
HIVE-8372	BUG-24210	ORC : Insert into ORC hits Null Pointer exception in MergeFileRecordProcessor.init(MergeFileRecordProcessor.java
HIVE-8387	BUG-24132	[Rolling upgrades] Validate Zookeeper retries from WebHCat
HIVE-8390	BUG-24303	CBO: Failures are wrapped in Runtime exception
HIVE-8392	BUG-24439	windows: HiveServer2 Operation.close fails
HIVE-8402	BUG-24472	ORC pushing SARGs into delta files resulting in ArrayOutOfBoundsException
HIVE-8403	BUG-24470	hive build failure on windows due to missing dependency .
HIVE-8408	BUG-24464	Hcatalog smoke test is failing with NullPointerException when v2 authorization is enabled in cli
HIVE-8427	BUG-23835	test_SinkHive tests failing with java.util.concurrent.TimeoutException
HIVE-8427	BUG-24480	Hive Streaming client throws NPE
HIVE-8461	BUG-25112	[HIVE-8461] Vectorized query has different Decimal result (Vector_Aggregate_9)
HIVE-8474	BUG-25069	Vectorization and ACID selects fail if not all columns are selected
HIVE-8516	BUG-24721	TestIUD_5 fails because of NPE in tez mode when vectorization is on

Apache JIRA	Key	Summary
HIVE-8534	BUG-25164	update hive configuration whitelist for sql std authorization
HIVE-8543	BUG-25102	Compaction tests failing due to the state getting stuck at ready for cleaning/working
HIVE-8547	BUG-25546	partition_varchar2 fails with CBO after some recent fix
HIVE-8557	BUG-24147	zk delegation token store - use kerberos auth, setup acl, use curator
HIVE-8560	BUG-25581	3rd party SerDes do not get table properties during initialize()
HIVE-8562	BUG-26015	MySQL compaction failures
HIVE-8588	BUG-25099	Webhcat sqoop tests fail to find the JDBC classes
HIVE-8597	BUG-25565	Hive CustomVertexManager should send the same set of splits for the small table in case of SMB joins
HIVE-8605	BUG-25420	Hive metastore backward incompatible, should not require property existing changes to run
HIVE-8615	BUG-23869	additional double quotes in select * complex datatypes output
HIVE-8619	BUG-26052	CBO: ".0" saga, part 3
HIVE-8629	BUG-25453	test_SinkHive tests failing with org.apache.flume.sink.hive.HiveWriter \$ConnectFailure error
HIVE-8631	BUG-26056	ACID / Streaming : When the txn list gets long, the queries on ACID tables are failing
HIVE-8635	BUG-26060	CBO: ambiguous_col negative test no longer fails
HIVE-8637	BUG-25419	Insert into nonbucketedtable select * from bucketedtable does not work
HIVE-8643	BUG-25103	Webhcat tests running as other user failing to execute commands
HIVE-8646	BUG-26331	Oozie Hive action fails with NoClassDefFound for FieldSchema class
HIVE-8660	BUG-26280	hive.exec.dynamic.partition.mode property should be included in the default list for hive.security.authorization.sqlstd.confwhitelist
HIVE-8685	BUG-26443	WebHCat DDL commands in non-secure mode failing
HIVE-8687	BUG-21253	Pig throws java.lang.ClassCastException when storing into Avro partitioned table with HCatLoader.
HIVE-8688	BUG-24291	HS2 concur test intermittently fails with org.apache.hive.com.esotericsoftware.kryo.KryoException: Buffer underflow.
HIVE-8704	BUG-24078	Hive HBaseStorageHandler tests failed because of "NPE Expected Accumulo table name to be provided in job configuration"

Apache JIRA	Key	Summary
HIVE-8711	BUG-26747	TxnHandler not catch deadlocks on Postgres, Oracle, and SQLServer
HIVE-8719	BUG-26803	CLONE - load data or insert into partition changes partition location previously defined.
HIVE-8732	BUG-26926	ORC string statistics are not merged correctly
HIVE-8740	BUG-26792	update query failed due to IndexOutOfBoundsException with hive.optimize.sort.dynamic.partition on
HIVE-8745	BUG-26992	Joins on decimal keys return different results whether they are run as reduce join or map join
HIVE-8754	BUG-26910	Test_Sqoop failing in secure cluster
HIVE-8772	BUG-26885	zookeeper info logs are always printed from beeline
HIVE-8785	BUG-27130	HiveServer2 LogDivertAppender should be more selective for beeline getLogs
HIVE-8797	BUG-27476	concurrent_insert_values_dynamic_partitioned and concurrent_insert_values_partitioned failed due to movetask error on windows
HIVE-8797	BUG-27312	MoveTaskError while running concurrent_insert_values_dynamic_partitioned.q test on a cluster with postgres as metastore
HIVE-8798	BUG-27315	ORACLE: acid_concurrency test failure due to "ORA-08177: can't serialize access for this transaction" error
HIVE-8812	BUG-27546	Fix TestMinimrCliDriver.testCliDriver_auto_sortmerge_join_16 on Windows
HIVE-8830	BUG-27433	HCatalog E2E testruns get stuck
HIVE-8875	BUG-27141	concurrent_delete_dynamic acid concurrency test failed on a cluster with postgres metastore
HIVE-8880	BUG-27356	TestIUD_3 failed because it didn't insert one record
HIVE-8901	BUG-27911	Hive queries failing withavax.jdo.JDODataStoreException: Could not retrieve transaction read-only status server
KNOX-443	BUG-24945	Release Notes for Knox->HS2 with new Connect String and Rolling Upgrades
KNOX-468	BUG-26869	Knox config for Group Look up to be optimized
KNOX-88	RMP-1250	Knox to provide full support for HDFS HA
MAHOUT-1589 MAHOUT-1590	BUG-18849	champlain mahout auto merge is broken
MAPREDUCE-5831	BUG-16650	[MAPREDUCE-5831] Old MR client is not compatible with new MR application

Apache JIRA	Key	Summary
MAPREDUCE-5858	BUG-19114	[MAPREDUCE-5858] MR should make use of the timeline server
MAPREDUCE-5891	BUG-21413	[MAPREDUCE-5891] Improved shuffle error handling across NM restarts
MAPREDUCE-5891 TEZ-1637	BUG-22819	Improved shuffle error handling across NM restarts
MAPREDUCE-5933	BUG-23714	BUG-17687 [Review] [MAPREDUCE-5933] Enable MR AM to post history events to the timeline server
MAPREDUCE-5933	BUG-19116	BUG-17687 [Review] [MAPREDUCE-5933] Enable MR AM to post history events to the timeline server
MAPREDUCE-5963	BUG-19885	[MAPREDUCE-5963] Make ShuffleHandler state versioning in rolling upgrade
MAPREDUCE-6018	BUG-25744	BUG-17687 [Review] [MAPREDUCE-6018] Create a framework specific config to enable timeline server
MAPREDUCE-6052	BUG-21756	[MAPREDUCE-6052] Support overriding log4j.properties per job
MAPREDUCE-6126	BUG-24551	[MAPREDUCE-6126] Rumen tool returns error "java.lang.IllegalArgumentException: JobBuilder.process(HistoryEvent): unknown event type"
MAPREDUCE-6142	BUG-26185	BUG-17687 [MAPREDUCE-6142] Test failure in TestJobHistoryEventHandler and TestMRTimelineEventHandling
MAPREDUCE-6149	BUG-26795	[MAPREDUCE-6149] Documentation our new configuration to log4j.property in mapred-site.xml
MAPREDUCE-6156	BUG-27591	[MAPREDUCE-6156] Fetcher - connect() doesn't handle connection refused correctly
OOZIE-1492	RMP-1882	RMP-1510 Oozie HA needs to work with HCatalog and SLA notifications
OOZIE-1532	RMP-1930	RMP-1510 Purging should remove completed children job for long running jobs
OOZIE-1769	RMP-1881	RMP-1510 Allow updating job properties on existing workflow and coordinator jobs.
OOZIE-1865	RMP-2038	RMP-1510 Oozie HA support in Secure Clusters
OOZIE-2051	BUG-26227	Easier use of Oozie + Tez jobs with hdp.version
PHOENIX-1069	RMP-1894	RMP-1507 Build secondary indexes during Phoenix bulk import
PHOENIX-1072	BUG-15077	sqlline.py: On failed connection, continues to run in the background
PIG-3949	BUG-17902	Pig compilation Failure - Champlain

Apache JIRA	Key	Summary
PIG-3989	BUG-15988	Need a way to pass PIG_OPTS to Hadoop for all Hadoop versions
PIG-4019	BUG-19109	champlain pig compile is broken -
PIG-4102	RMP-1981	RMP-1508 Pig should support ORC predicate pushdown
PIG-4141	BUG-22532	Orc tests failed with '/hive-common-jar' does not exist.
PIG-4156	BUG-22451	Null pointer exception in PIG when executing a script remotely on Linux
PIG-4157	BUG-22395	Pig compilation failure due to HIVE-7208
PIG-4170	BUG-22933	Multiquery with different type of key gives wrong result
PIG-4176	BUG-22964	Bloom_* fails on Pig on Tez with File Not found exception
PIG-4184	BUG-22997	Datafu UnitTest failures - linux
PIG-4187	BUG-23340	Orc_3 fails with FileNotFoundException
PIG-4218	BUG-19417	Pig OrcStorage fail to load a map with null key
PIG-4227	BUG-24936	streamingpythonudfs_5 fails with "StreamingUDFException: LINE : Error deserializing output."
PIG-4227	BUG-24934	streamingpythonudfs_12 fails with checksum mismatch
PIG-4250	BUG-24375	Review Coverity scan results for Pig
PIG-4257	BUG-24252	Bloom_3, Union_7,8,13, Join_6,7,8 fails with Local Rearrange exception
PIG-4257	BUG-26662	Rank_*, Union_4,5, Order_18, 20 fails with Can't get Master Kerberos principal for use as renewer
PIG-4314	BUG-27287	Ambari Pig test hangs and make no progress
SLIDER-180	RMP-2115	Requirements to run "HBase on YARN" via Slider
SLIDER-181	RMP-2113	Requirements to run "Storm on YARN" via Slider
SLIDER-182	RMP-2114	Requirements to run "Accumulo on YARN" via Slider
SLIDER-183	RMP-2117	Long lived service support in Slider
SLIDER-184	BUG-22646	Slider support for an HA enabled Cluster
SLIDER-214	RMP-2085	Usability improvements to Slider
SLIDER-267	BUG-20699	Slider AM web UI should give a link to show AM logs
SLIDER-306	BUG-25516	Slider cli list sub-command cannot list stopped applications
SLIDER-334	BUG-21419	slider compilation failure
SLIDER-430	BUG-22968	Slider registry should give better error message on missing command
SLIDER-448	BUG-23610	AMFailuresIT test fails in slider

Apache JIRA	Key	Summary
SLIDER-450	BUG-23591	Rename appConfig.json and resources.json files as default inside the app package
SLIDER-481	BUG-23004	Application level exports (exportGroups) should allow multiple exports of the same name on a per component instance basis
SLIDER-486	BUG-24243	slider bin internally uses -classpath and causes issues in some jdk (notably suse11.1 Teradata)
SLIDER-497	BUG-24581	Slider funtest are failing on secure slider cluster due to AM failure
SLIDER-500	BUG-24642	Slider-Accumulo: test_continuous_ingestion fails throwing org.apache.hadoop.registry.client.exceptions.InvalidPathNameException: Invalid Path element "hrt_qa"
SLIDER-513	BUG-24869	AgentFailures2IT test is failing
SLIDER-516	BUG-24958	Make Slider registry call not default to current user.
SLIDER-520	BUG-24981	secure slider funtests are failing with exitcode=36
SLIDER-522	BUG-24991	Storm-slider security throws NPE at create.
SLIDER-534	BUG-26301	Slider-HA: CommandEnvironmentIT gets Error
SLIDER-540	BUG-25283	Slider tests do not respect configs from yarn-site.xml
SLIDER-543	BUG-25365	champlain CI: slider builds are broken
SLIDER-557	BUG-25651	HBase slider deployment fails on secure cluster
SLIDER-560	BUG-25763	Add default secured appConfig.json file to application packages
SLIDER-581	BUG-26188	Windows: resolvecommandit test failed with wrong error code 0 not 44
SLIDER-581	BUG-26190	Windows: AgentClusterLifecycleIT test failed with error code 0, expected 73
SLIDER-581	BUG-26189	Windows: UnknownClusterOperationsIT test exit code - 0 not 69
SLIDER-588	BUG-26365	AppsThroughAgentIT test failed with 'Application registry is not accessible'
SLIDER-593	BUG-26608	[SLIDER-593] Slider registry client returns IOException returned a response status of 302 Found.
SLIDER-596	BUG-26507	Slider should automatically replace \${USER_NAME} token from config
SLIDER-598	BUG-26469	Slider-Hbase creation on Windows fails due to JsonParse Exception
SLIDER-602	BUG-25915	Tracking ticket to review and commit RM integration changes for Service Registry
SLIDER-603	BUG-26646	AppsThroughAgentIT fails with NPE

Apache JIRA	Key	Summary
SLIDER-605	BUG-26688	slider list command does not return application in the order of most recent first
SLIDER-606	BUG-26707	[SLIDER-606] quicklinks dont return any data for slider apps
SLIDER-613	BUG-26093	'AppsThroughAgentIT-testCreateFlex' Failed on cluster with encryption on
SLIDER-621	BUG-27095	slider AgentClusterLifecycleIT test case timing out after 10 minutes
SLIDER-624	BUG-27128	Couple of slider-hbase tests failed due to delayed application launch
SLIDER-630	BUG-27321	Linux UT failures on CentOS5, Debian, and Suse11
SLIDER-646	BUG-27626	AgentLaunchFailureIT test failed with error code 68
SLIDER-647	BUG-27740	allocation requests not being satisfied when a cluster goes to labels
SLIDER-69	RMP-2118	Failure resiliency
SLIDER-86	RMP-2066	Slider on all Windows platforms supported by HDP
SLIDER-89	RMP-2068	Slider on Kerberos-enabled Cluster
SQOOP-1322	RMP-1876	RMP-1509 Support all Hive types when importing via HCatalog
SQOOP-1322	RMP-1828	RMP-1509 Support multiple partition keys in Sqoop HCat integration
SQOOP-1519	BUG-22849	Sqoop: hive-import and Hbase import fails with – direct option - with oraoop
STORM-10 STORM-490 STORM-9	BUG-22564	Storm Misc issues on Windows
STORM-346	BUG-20302	(Security) Oozie style delegation tokens for HDFS/HBase STORM-346
STORM-347	BUG-20298	authentication should allow for groups not just users STORM-347
STORM-349	BUG-20297	(Security) ui actions should have nimbus like authroization - STORM-349
STORM-361	BUG-20296	Add JSON-P support to Storm UI API STORM-361
TEZ-1024 TEZ-1470	BUG-17939	Hive on tez job failed as it never finished
TEZ-1027	BUG-15413	Recovery Information does not get stored at tez.staging-dir with session on
TEZ-1032 TEZ-1205	BUG-17871	allow specifying tasks/vertices to be profiled
TEZ-1033	BUG-16512	[TEZ-1033] AM hangs during recovery with Awaiting init event
TEZ-1044	BUG-16609	Need to consider map/reduce java.opts values for container reuse

Apache JIRA	Key	Summary
TEZ-1080	BUG-19324	Configuration for Intermediate Inputs and Outputs
TEZ-1083	BUG-24993	Enable IFile RLE for DefaultSorter
TEZ-1083	BUG-25071	[TEZ-1083] Enable IFile RLE for DefaultSorter
TEZ-1141	BUG-25072	[TEZ-1141] DAGStatus.Progress should include number of failed attempts
TEZ-1157	BUG-22116	[TEZ-1157] Optimize broadcast :- Tasks pertaining to same job in same machine should not download multiple copies of broadcast data
TEZ-1160	BUG-18437	[TEZ-1160] Allow users to set non-tez configs on client side
TEZ-1209	BUG-19326	Add Javadoc for Tez APIs
TEZ-1210	BUG-25081	[TEZ-1210] TezClientUtils.localizeDagPlanAsText() needs to be fixed for session mode
TEZ-1223	BUG-19284	At large scale (10 TB), shuffle throws errors "Encoded byte size for String was -30, which is outside of 0..1000 range."
TEZ-1238	BUG-19880	Display more clear diagnostics info on client side if missing jar in LocalResource or Exception happen in Processor
TEZ-1240	BUG-17787	Tracking jira for Tez diagnostic message improvement.
TEZ-1387	BUG-18923	BUG-17787 Insufficient diagnostics/ error messages when jobs fail on a badly setup cluster
TEZ-1424	BUG-25073	[TEZ-1424] Fixes to DAG text representation in debug mode
TEZ-1494	BUG-22111	[TEZ-1494] DAG hangs waiting for ShuffleManager.getNextInput()
TEZ-15	RMP-1998	RMP-1988 AM Recovery Umbrella
TEZ-1536	BUG-22081	Tez - There's a typo in the code, it has 'configuration' spelled wrong in the error message.
TEZ-1540	BUG-22115	[TEZ-1540] [Umbrella] Handle parallelism updates and versioning w/ custom InputInitializerEvents
TEZ-1549	BUG-22312	Windows Tez Unit Test Failure
TEZ-1567	BUG-25095	Avoid blacklisting nodes when the disable blacklisting threshold is about to be hit
TEZ-1575	BUG-22891	[TEZ-1575]Add MR settings support to MRRSleep
TEZ-1585	BUG-22939	Investigate: Memory leak in session mode
TEZ-1590	BUG-25074	[TEZ-1590] Fetchers should not report failures after the Processor on the task completes
TEZ-1596	BUG-25079	[TEZ-1596] Secure Shuffle utils is extremely expensive for fast queries
TEZ-1624	BUG-23104	tez champlain linux unittest failure

Apache JIRA	Key	Summary
TEZ-1629	BUG-23918	[TEZ-1629] ContainerLauncherImpl's event handler thread should check for threadpool's status before submitting a task
TEZ-1634	BUG-22962	Tez job via fails with Caused by: java.io.IOException: Exceeded MAX_FAILED_UNIQUE_FETCHES
TEZ-1638	BUG-25078	[TEZ-1638] Missing type parametrization in UnorderedKVInputConfig?
TEZ-1643	BUG-24447	[TEZ-1643] DAGAppMaster kills DAG & shuts down, when RM is restarted
TEZ-1647	BUG-24515	[TEZ-1647] Issue with caching of events in VertexManager::onRootVertexInitialized
TEZ-1649	BUG-24692	Cross_5 hangs for more than 7 hours
TEZ-1666	BUG-24975	UserPayload should be null if the payload is not specified
TEZ-1666	BUG-25093	UserPayload should be null if the payload is not specified
TEZ-1667	BUG-24979	Add a system test for InitializerEvents
TEZ-1668	BUG-24980	InputInitializers should be able to register for Vertex state updates in the constructor itself
TEZ-1669	BUG-25089	yarn-swimlanes.sh throws error
TEZ-1689 TEZ-1703 TEZ-1267	BUG-25034	[TEZ-1267, TEZ-1689, TEZ-1703] Exception handling when Routing Events
TEZ-1698	BUG-26186	Use ResourceCalculatorPlugin instead of ResourceCalculatorProcessTree in Tez
TEZ-1712	BUG-25386	Union_9, Rank_10 fails with Vertex killed as other vertex failed.
TEZ-1713	BUG-26098	Tez complains about path not being fully qualified if the schema is not specified
TEZ-1731	BUG-26573	Checksum Errors when running queries in Tez
TEZ-1750	BUG-27233	Possible Out of Order scheduling can end up slowing down queries
TEZ-645 TEZ-646 TEZ-776	BUG-10941	GRIDMIX000000 got final Status KILLED with container reuse disabled - AM OOM
TEZ-657	BUG-18146	Tez should not consider tasks running in pre-empted containers to have failed
TEZ-792	BUG-25075	[TEZ-792] Default staging path should have user name
TEZ-857	BUG-19328	Separate user facing and framework facing interfaces on Tez runtime
TEZ-978	BUG-22113	[TEZ-978] Enhance auto parallelism tuning for queries having empty outputs or data skewness

Apache JIRA	Key	Summary
TEZ-978	BUG-20718	Enhance auto parallelism tuning for queries having empty outputs or data skewness
YARN-1104	BUG-17884	[YARN-1104] NMs to support rolling logs of stdout & stderr
YARN-1250	BUG-17034	[YARN-1250] Add ACLs manager of the generic history
YARN-1338	BUG-18018	[YARN-1338] [Review] Recover localized resource cache state upon nodemanager restart
YARN-1339	BUG-18662	[YARN-1339] Code review for recovering DeletionService state in NodeManager Restart
YARN-1445	BUG-11102	[YARN-1445] Separate FINISHING and FINISHED state in YarnApplicationState
YARN-1506	BUG-21222	[YARN-1506] Replace set resource change on RMNode/SchedulerNode directly with event notification.
YARN-1885	BUG-15613	[YARN-1885] yarn logs command did not provide logs for all containers across all attempts
YARN-1922	BUG-16862	Process group remains alive after container process is killed externally
YARN-1922	BUG-26552	Child processes aren't killed when container process is killed
YARN-1964	BUG-26423	[Docker] Review and commit docker patch from YARN-1964
YARN-1976	BUG-16829	[YARN-1976] Tracking url missing http protocol for FAILED application
YARN-2016	BUG-16566	Yarn getApplicationRequest start time range is not honored
YARN-2045	BUG-19884	[YARN-2045] Versioning NMStateStore for Rolling Upgrade
YARN-2102	BUG-19113	[YARN-2102] More generalized timeline ACLs
YARN-2167	BUG-19019	[YARN-2167] LevelDbIterator should get closed in NMLevelDbStateStoreService#loadLocalizationState() within finally block
YARN-2173	BUG-19111	[YARN-2173] Enabling HTTPS for the reader REST APIs of TimelineServer
YARN-2174	BUG-19112	[YARN-2174] Enabling HTTPS for the writer REST API of TimelineServer
YARN-2182	BUG-21671	Container Id format gets changed after work preserving rm restart
YARN-2208	BUG-19496	BUG-17882 [YARN-2208] AMRMTOKENManager need to have a way to roll over AMRMTOKEN
YARN-2211	BUG-19731	BUG-17882 [YARN-2211] RMStateStore needs to save AMRMTOKEN master key for recovery when RM restart/failover happens
YARN-2212	BUG-19732	BUG-17882 [YARN-2212] ApplicationMaster

Apache JIRA	Key	Summary
		needs to find a way to update the AMRMTOKEN periodically
YARN-2228	BUG-19607	[YARN-2228] TimelineServer should load pseudo authentication filter when authentication = simple
YARN-2237	BUG-19733	BUG-17882 [YARN-2237]MRAppMaster changes for AMRMTOKEN roll-up
YARN-2247	BUG-19645	[YARN-2247] Add support for authentication for RM web services using delegation tokens
YARN-2257	RMP-2260	(2.1) - [YARN-2411] Support Capacity Scheduler server side mapping of user to a default queue
YARN-2269	BUG-17105	External links needs to be removed from YARN UI
YARN-2279	BUG-25576	[YARN-2279] Add UTs to cover timeline server authentication
YARN-2288	BUG-20065	[YARN-2288] LevelDbTimelineStore should be versioned for rolling upgrade
YARN-2298	BUG-20202	[YARN-2298] Move TimelineClient to yarn-common
YARN-2302	BUG-20246	[YARN-2302] Refactor TimelineWebServices
YARN-2347	BUG-20467	[YARN-2347]Consolidate RMStateVersion and NMDBSchemaVersion into StateVersion
YARN-2425	BUG-21377	[YARN-2425] When Application submitted by via Yarn RM WS, log aggregation does not happens
YARN-2426	BUG-21372	[HDFS-6904] ResourceManger is not able renew WebHDFS token when application submitted by Yarn WebService
YARN-2446	BUG-21692	[YARN-2446] Using TimelineNamespace to shield the entities of a user
YARN-2449	BUG-21693	[YARN-2449] Regression after YARN-2397 for timeline server
YARN-2459	BUG-22008	[YARN-2459] ResourceManager receives FATAL "org.apache.hadoop.yarn.server.resourcemanager.ResourceManager Received a org.apache.hadoop.yarn.server.resourcemanager.RMFatalEvent of type STATE_STORE_OP_FAILED."
YARN-2464	BUG-22635	[YARN-2464] Provide Hadoop as a local resource (on HDFS) which can be used by other projects
YARN-2468	BUG-21985	BUG-12006 [YARN-2468] Log handling for LRS
YARN-2505	BUG-20590	[YARN-2505] Support get/set hard node labels using RM REST API
YARN-2515	BUG-22297	Tez task fails to find correct container in HA configuration

Apache JIRA	Key	Summary
YARN-2529	BUG-22602	[YARN-2529] Generic history service RPC interface doesn't work when service authorization is enabled
YARN-2538	BUG-21674	BUG-17882 [YARN-2538]Add logs when new AMRMTOKEN arrives.
YARN-2557	BUG-23007	[YARN-2557]Need to add a parameter "attemptFailuresValidityInterval" in DistributedShell
YARN-2559	BUG-23035	[YARN-2559] ResourceManager sometime become un-responsive due to NPE in SystemMetricsPublisher
YARN-2563	BUG-22911	[YARN-2563] Oozie jobs are failing with authentication errors
YARN-2565	BUG-23130	[YARN-2565]RM is fails to start when AHS/GHS is enabled in secure mode without doing manual kinit as yarn
YARN-2569	BUG-23469	BUG-12006 [YARN-2569]Log Handling for LRS API Changes
YARN-2581	BUG-23470	BUG-12006 [YARN-2581]NMs need to find a way to get LogAggregationContext
YARN-2582	BUG-23472	BUG-12006 [YARN-2582]Log related CLI and Web UI changes for LRS
YARN-2583	BUG-23471	BUG-12006 [YARN-2583]Modify the LogDeletionService to support Log aggregation for LRS
YARN-2591	BUG-23741	[YARN-2591] AHSWebServices should return FORBIDDEN(403) if the request user doesn't have access to the history data
YARN-2594	BUG-23355	[YARN-2594] ResourceManager sometimes Hangs blocking tests to run
YARN-2598	BUG-23602	[YARN-2598] GHS should show N/A instead of null for the inaccessible information
YARN-2602	BUG-23622	[YARN-2602] Generic History Service of TimelineServer sometimes not able to handle NPE
YARN-2613	BUG-23512	[YARN-2613] NMClient doesn't have retries for supporting rolling-upgrades
YARN-2615	BUG-23761	[YARN-2615] ClientToAMTokenIdentifier and DelegationTokenIdentifier(s)
YARN-2621	BUG-23857	[YARN-2621] Simplify the output when the user doesn't have the access for getDomain(s)
YARN-2627	BUG-23813	[YARN-2627]RM log does not have value for "retry-count-window-ms"
YARN-2629	BUG-23962	[YARN-2629] Make distributed shell use the domain-based timeline ACLs
YARN-2631	BUG-23969	BUG-12006 [YARN-2631]DistributedShell lacks support for LRS log collection interval
YARN-2632	BUG-26912	[YARN-2632] Document NodeManager Restart Feature

Apache JIRA	Key	Summary
YARN-2637	BUG-24760	Hive queries run into AM resource hangs when running under xasecure
YARN-2651	BUG-24570	BUG-12006 [YARN-2651]Spin off the LogRollingInterval from LogAggregationContext
YARN-2668	BUG-24553	hadoop champlain build failure: registry linking to ZK
YARN-2689	BUG-24919	TestSecureRMRegistryOperations failing on windows
YARN-2701	BUG-24065	[YARN-2701]HA testing: Tasks failing with error "Can't create directory /grid/0/yarn/local/usercache/hrt_qa/appcache/application_1412158618225_0001 - File exists"
YARN-2703	BUG-25442	BUG-12006 [YARN-2703]Add logUploadedTime into LogValue for better display
YARN-2715	BUG-25313	[YARN-2715] RM get delegation token is broken for proxy users when using hadoop RPC
YARN-2717	BUG-25470	BUG-12006 [YARN-2717]containerLogNotFound log shows multiple time for the same container
YARN-2720	BUG-25323	Windows: Wildcard classpath variables not expanded against resources contained in archives
YARN-2723	BUG-25560	[YARN-2723] rmdadmin - replaceLabelsOnNode does not correctly parse port
YARN-2726	BUG-25585	[YARN-2726] CapacityScheduler should explicitly log when an accessible label has no capacity
YARN-2734	BUG-25732	BUG-12006 [YARN-2734] If a sub-folder is encountered by log aggregator it results in invalid aggregated file
YARN-2741	BUG-25525	[YARN-2741] HDP 2.1.3 - Customer getting "No Logs available for Container Container_<number>" when output to non-default disk
YARN-2743	BUG-25523	[YARN-2743] yarn apps failing JA009: Failed to run job : Failed to renew token: Kind: RM_DELEGATION_TOKEN on secure clusters
YARN-2743	BUG-25838	[YARN-2743] yarn applications fail with digest mismatch errors on un secure cluster
YARN-2747	BUG-25968	[YARN-2747]TestAggregatedLogFormat fails in trunk
YARN-2758	BUG-22442	BUG-22334 [YARN-2758] Unit test failures in TestApplicationHistoryService in Champlain builds
YARN-2769	BUG-26239	[YARN-2769] Timeline Acls test failing on windows

Apache JIRA	Key	Summary
YARN-2770	BUG-25924	[YARN-2770] Timeline Service delegation tokens need to be renewed by the RM
YARN-2779	BUG-26386	[YARN-2779] SystemMetricsPublisher can use Kerberos directly instead of timeline DT
YARN-2786	BUG-26604	System tests are disabled for Node Labels
YARN-2788	BUG-26251	[YARN-2788] YARN champlain JobHistoryServer cannot read logs written by 2.4.0 builds
YARN-2795	BUG-26526	[YARN-2795] Resource Manager fails startup with HDFS label storage and secure cluster
YARN-2798	BUG-26658	[YARN-2798] Oozie jobs failing on secure cluster with timeline server DT issues
YARN-2804	BUG-26666	[YARN-2804] Timeline server out log explodes with tons of repeated exceptions
YARN-2805	BUG-26831	[YARN-2805] RM failed to startup because of 'Failed to login' exception
YARN-2812	BUG-26892	BUG-22334 [YARN-2812] Unit test failures in TestApplicationHistoryServer in Champlain builds
YARN-2813	BUG-27014	[YARN-2813] NPE from MemoryTimelineStore.getDomains
YARN-2818	BUG-27084	[YARN-2818] Remove the logic to inject entity owner as the primary filter
YARN-2830	BUG-27396	Cherry-pick YARN-2830 in Champlain
YARN-2834	BUG-27359	[YARN-2834] RM crashed with NPE
YARN-2841	BUG-27442	[YARN-2841]RMProxy should retry EOFException
YARN-2843	BUG-27494	[YARN-2843] NodeLabels manager should trim all inputs for hosts and labels
YARN-2846	BUG-27182	[YARN-2846] Work Preserving NM Restart - Restarting NMs causes app attempt failure
YARN-2853	BUG-27526	[YARN-2853] One of our system test running Tez are hung, looks like RM is hung
YARN-2879	BUG-27784	[YARN-2879] Verify downlevel client compatibility behavior
YARN-611	BUG-19235	[YARN-611] Add an AM retry count reset window to YARN RM
YARN-614	BUG-19428	[YARN-614] Retry attempts automatically for hardware failures or YARN issues and set default app retries to 1
YARN-668	BUG-20258	[YARN-668] TokenIdentifier serialization should consider unknown fields for future extension
YARN-796	BUG-9137	Make nodes in a Yarn cluster be identifiable by its label

Apache JIRA	Key	Summary
YARN-90	BUG-12327	[Yarn-90] NM cannot detect when bad disks become healthy again.
YARN-913	RMP-2116	Add a registry for services in a YARN cluster
YARN-920	BUG-16908	YARN RM shows one running app but a node shows its running more than one
ZOOKEEPER-1897	BUG-26062	Zookeeper 3.4.6 does not execute via command line
ZOOKEEPER-2064	BUG-24379	Review Coverity scan results for Zookeeper

1.7. Apache Patch Information

The following sections list patches in each HDP 2.2.0 component beyond what was fixed in the base version of the Apache component.

1.7.1. DataFu 1.2.0

- [DATAFU-58](#) Unit tests against Hadoop 2
- [DATAFU-59](#) Fix unit tests under Window

1.7.2. Falcon 0.6.0

- [FALCON-403](#) Add support for Falcon to run on Windows
- [FALCON-434](#) Migrate HCatalog APIs to new namespace
- [FALCON-417](#) Upgrade Hive and HCatalog to latest stable version
- [FALCON-896](#) Upgrade Oozie to 4.1.0 and make it the default for Falcon

1.7.3. Flume 1.5.2

- [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-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)

1.7.4. Hadoop Common/HDFS 2.6.0

- [HADOOP-10840](#) Fix OutOfMemoryError caused by metrics system in Azure File
- [HADOOP-10839](#) Add unregisterSource() to MetricsSystem API.
- [HDFS-7153](#) Add storagePolicy to NN edit log during file creation.
- [HDFS-7093](#) Add config key to restrict setStoragePolicy.
- [HDFS-6581](#) Support for writing to single replica in RAM. Applied HDFS-6581.merge.15.patch to C4hamplain. This is a consolidated patch containing all feature changes and was applied to trunk
- [YARN-2254](#) TestRMWebServicesAppsModification should run against both
- [YARN-1979](#) TestDirectoryCollection fails when the umask is unusual.

1.7.5. HBase 0.98.4

- [HBASE-12478](#) [HBASE-10141](#) and MIN_VERSIONS are not compatible
- [HBASE-12445](#) hbase is removing all remaining cells immediately after the cell marked with marker = KeyValue.Type.DeleteColumn via PUT
- [HBASE-12432](#) RpcRetryingCaller should log after fixed number of retries like AsyncProcess
- [HBASE-12449](#) Use the max timestamp of current or old cell's timestamp in HRegion.append()
- [HBASE-12416](#) RegionServerCallable should report what host it was communicating with
- [HBASE-10844](#) Coprocessor failure during batchmutation leaves the memstore datastructs in an inconsistent state
- [HBASE-12548](#) Improve debuggability of IntegrationTestTimeBoundedRequestsWithRegionReplicas

- [HBASE-12407](#) HConnectionKey doesn't contain CUSTOM_CONTROLLER_CONF_KEY in CONNECTION_PROPERTIES
- [HBASE-12402](#) ZKPermissionWatcher race condition in refreshing the cache leaving stale ACLs and causing AccessDenied
- [HBASE-11972](#) The "doAs user" used in the update to hbase:acl table RPC is incorrect
- [HBASE-12549](#) Fix TestAssignmentManagerOnCluster#testAssignRacingWithSSH() flaky test
- [HBASE-12403](#) IntegrationTestMTTR flaky due to aggressive RS restart timeout
- [HBASE-12401](#) Add some timestamp signposts in IntegrationTestMTTR
- [HBASE-12562](#) Handling memory pressure for secondary region replicas
- [HBASE-12398](#) Region isn't assigned in an extreme race condition
- [HBASE-12377](#) HBaseAdmin#deleteTable fails when META region is moved around the same timeframe
- [HBASE-12372](#) [WINDOWS] Enable log4j configuration in hbase.cmd
- [HBASE-12560](#) [WINDOWS] Append the classpath from Hadoop to HBase classpath in bin/hbase.cmd
- [HBASE-11716](#) LoadTestDataGeneratorWithVisibilityLabels should handle Delete mutations
- [HBASE-11976](#) Server startcode is not checked for bulk region assignment
- [HBASE-12367](#) Integration tests should not restore the cluster if the CM is not destructive
- [HBASE-12485](#) Maintain SeqId monotonically increasing when Region Replica is on
- [HBASE-12356](#) Rpc with region replica does not propagate tracing spans
- [HBASE-12335](#) IntegrationTestRegionReplicaPerf is flaky
- [HBASE-12317](#) Run IntegrationTestRegionReplicaPerf w.o mapred
- [HBASE-10955](#) HBCK leaves the region in masters in-memory RegionStates if region hdfs dir is lost
- [HBASE-12319](#) Inconsistencies during region recovery due to close/open of a region during recovery
- [HBASE-12301](#) user_permission command does not show global permissions
- [HBASE-12189](#) Fix new issues found by coverity static analysis
- [HBASE-12277](#) Refactor bulkLoad methods in AccessController to its own interface

- [HBASE-12278](#) Race condition in TestSecureLoadIncrementalHFilesSplitRecovery
- [HBASE-11788](#) hbase is not deleting the cell when a Put with a KeyValue, KeyValue.Type.Delete is submitted
- [HBASE-11609](#) LoadIncrementalHFiles fails if the namespace is specified
- [HBASE-12274](#) Race between RegionScannerImpl#nextInternal() and RegionScannerImpl#close() may produce null pointer exception
- [HBASE-12098](#) User granted namespace table create permissions can't create a table
- [HBASE-11794](#) StripeStoreFlusher causes NullPointerException
- [HBASE-12078](#) Missing Data when scanning using PREFIX_TREE DATA-BLOCK-ENCODING
- [HBASE-11728](#) Data loss while scanning using PREFIX_TREE DATA-BLOCK-ENCODING
- [HBASE-12156](#) TableName cache isn't used for one of valueOf methods
- [HBASE-12164](#) Check for presence of user Id in SecureBulkLoadEndpoint#secureBulkLoadHFiles() is inaccurate
- [HBASE-12162](#) HBaseAdmin#getTableDescriptor() may fail in case master fails over
- [HBASE-12095](#) SecureWALCellCodec should handle the case where encryption is disabled
- [HBASE-12123](#) Failed assertion in BucketCache after 11331
- [HBASE-11982](#) Bootstrapping hbase:meta table creates a WAL file in region dir
- [HBASE-10141](#) instead of putting expired store files thru compaction, just archive them
- [HBASE-12099](#) TestScannerModel fails if using jackson 1.9.13
- [HBASE-12072](#) We are doing 35 x 35 retries for master operations (Not full patch. Instead master operations use a 10min timeout)
- [HBASE-12056](#) RPC logging too much in DEBUG mode
- [HBASE-11988](#) AC/VC system table create on postStartMaster fails too often in test
- [HBASE-12055](#) TestBucketWriterThread hangs flakily based on timing
- [HBASE-11798](#) TestBucketWriterThread may hang due to WriterThread stopping prematurely
- [HBASE-11974](#) When a disabled table is scanned, NotServingRegionException is thrown instead of TableNotEnabledException
- [HBASE-12012](#) Improve cancellation for the scan RPCs
- [HBASE-12008](#) Remove IntegrationTestImportTsv#testRunFromOutputCommitter
- [HBASE-11868](#) Data loss in hlog when the hdfs is unavailable

- [HBASE-11989](#) IntegrationTestLoadAndVerify cannot be configured anymore on distributed mode
- [HBASE-11963](#) Synchronize peer cluster replication connection attempts
- [HBASE-11887](#) Memory retention in branch-1; millions of instances of LiteralByteString for column qualifier and value
- [HBASE-11705](#) callQueueSize should be decremented in a fail-fast scenario
- [HBASE-11898](#) CoprocessorHost.Environment should cache class loader instance.
- [HBASE-11876](#) RegionScanner.nextRaw should not update metrics
- [HBASE-11865](#) Result implements CellScannable; rather it should BE a CellScanner
- [HBASE-7910](#) Dont use reflection for security [HBASE-12052](#) BulkLoad Failed due to no write permission on input files
- [HBASE-11973](#) The url of the token file location set by IntegrationTestImportTsv should point to the localfs
- [HBASE-11940](#) Add utility scripts for snapshotting / restoring all tables in cluster
- [HBASE-11847](#) HFile tool should be able to print block headers
- [HBASE-11845](#) HFile tool should implement Tool, disable blockcache by default
- [HBASE-11942](#) Fix TestHRegionBusyWait
- [HBASE-11331](#) [blockcache] lazy block decompression
- [HBASE-11678](#) BucketCache ramCache fills heap after running a few hours
- [HBASE-11810](#) Access SSL Passwords through Credential Provider API (Larry McCay)
- [HBASE-11908](#) Region replicas should be added to the meta table at the time of table creation
- [HBASE-12561](#) Replicas of regions can be cached from different instances of the table in MetaCache
- [HBASE-11863](#) WAL files are not archived and stays in the WAL directory after splitting
- [HBASE-11813](#) CellScanner#advance may overflow stack (Stack)
- [HBASE-11903](#) Directly invoking split & merge of replica regions should be disallowed
- [HBASE-11886](#) The creator of the table should have all permissions on the table
- [HBASE-11857](#) Restore ReaderBase.initAfterCompression() and WALCellCodec.create(Configuration, CompressionContext)
- [HBASE-11762](#) Record the class name of Codec in WAL header

- [HBASE-11580](#) Failover handling for secondary region replicas
- [HBASE-11842](#) Integration test for async wal replication to secondary regions
- [HBASE-11564](#) Improve cancellation management in the rpc layer
- [HBASE-11297](#) Remove some synchronos in the rpcServer responder
- [HBASE-11719](#) Remove some unused paths in AsyncClient
- [HBASE-11718](#) Remove some logs in RpcClient.java
- [HBASE-10814](#) RpcClient: some calls can get stuck when connection is closing
- [HBASE-10942](#) support parallel request cancellation for multi-get
- [HBASE-11851](#) RpcClient can try to close a connection not ready to close
- [HBASE-11835](#) Wrong management of non expected calls in the client
- [HBASE-11836](#) IntegrationTestTimeBoundedMultiGetRequestsWithRegionReplicas tests simple get by default
- [HBASE-11654](#) WAL Splitting dirs are not deleted after replay
- [HBASE-11837](#) Scanner from RegionCoprocesorEnvironment.getTable(TableName) returns no local data
- [HBASE-11766](#) Backdoor CoprocessorHConnection is no longer being used for local writes
- [HBASE-11574](#) hbase:meta's regions can be replicated
- [HBASE-11571](#) Bulk load handling from secondary region replicas
- [HBASE-11569](#) Flush / Compaction handling from secondary region replicas
- [HBASE-11572](#) Add support for doing get/scans against a particular replica_id
- [HBASE-11567](#) Write bulk load COMMIT events to WAL
- [HBASE-11568](#) Async WAL replication for region replicas
- [HBASE-11820](#) ReplicationSource : Set replication codec class as RPC codec class on a cloned Configuration
- [HBASE-11367](#) Pluggable replication endpoint
- [HBASE-11512](#) Write region open/close events to WAL
- [HBASE-11511](#) Write flush events to WAL
- [HBASE-11594](#) Unhandled NoNodeException in distributed log replay mode
- [HBASE-11550](#) Custom value for BUCKET_CACHE_BUCKETS_KEY should be sorted

- [HBASE-11551](#) BucketCache.run() doesn't handle exceptions correctly
- [HBASE-11620](#) Record the class name of Writer in WAL header so that only proper Reader can open the WAL file
- [HBASE-11557](#) Counter.class should always be added as dependencyJar
- [HBASE-10674](#) HCK should be updated to do replica related checks
- [HBASE-11261](#) Handle splitting/merging of regions that have region_replication greater than one
- [HBASE-10818](#) Add integration test for bulkload with replicas
- [HBASE-11214](#) Fixes for scans on a replicated table
- [HBASE-10357](#) Failover RPC's for scans
- [HBASE-11344](#) Hide row keys and such from the web UIs
- [HBASE-11293](#) Master and Region servers fail to start when hbase.master.ipc.address=0.0.0.0, hbase.regionserver.ipc.address=0.0.0.0 and Kerberos is enabled
- [HBASE-10957](#) [HBASE-10070](#): HMaster can abort with NPE in #rebuildUserRegions
- [HBASE-10859](#) HStore.openStoreFiles() should pass the StoreFileInfo object to createStoreFileAndReader().
- [HBASE-10858](#) TestRegionRebalancing is failing
- [HBASE-10785](#) Metas own location should be cached
- [HBASE-10814](#) RpcClient: some calls can get stuck when connection is closing
- [HBASE-10791](#) Add integration test to demonstrate performance improvement
- [HBASE-10810](#) LoadTestTool should share the connection and connection pool
- [HBASE-10794](#) multi-get should handle replica location missing from cache
- [HBASE-10634](#) Multiget doesn't fully work.
- [HBASE-10701](#) Cache invalidation improvements from client side
- [HBASE-10778](#) Unique keys accounting in MultiThreadedReader is incorrect
- [HBASE-10743](#) Replica map update is problematic in RegionStates
- [HBASE-10616](#) Integration test for multi-get calls
- [HBASE-10734](#) Fix RegionStates.getRegionAssignments to not add duplicate regions
- [HBASE-10729](#) Enable table doesn't balance out replicas evenly if the replicas were unassigned earlier

- [HBASE-10720](#) rpcClient: Wrong log level when closing the connection
- [HBASE-10688](#) Add a draining_node script to manage nodes in draining mode
- [HBASE-10704](#) BaseLoadBalancer#roundRobinAssignment() may add same region to assignment plan multiple times
- [HBASE-10633](#) StoreFileRefresherChore throws ConcurrentModificationException sometimes
- [HBASE-10572](#) Create an IntegrationTest for region replicas.
- [HBASE-10703](#) TestAsyncProcess does not pass on
- [HBASE-10070](#) [HBASE-10637](#) rpcClient: Setup the iostreams when writing
- [HBASE-10635](#) thrift#TestThriftServer fails due to TTL validity check
- [HBASE-10620](#) LoadBalancer.needsBalance() should check for co-located region replicas as well
- [HBASE-10672](#) Table snapshot should handle tables whose REGION_REPLICATION is greater than one.
- [HBASE-10591](#) Sanity check table configuration in createTable
- [HBASE-10661](#)
TestStochasticLoadBalancer.testRegionReplicationOnMidClusterWithRacks() is flaky
- [HBASE-10351](#) LoadBalancer changes for supporting region replicas
- [HBASE-10630](#) NullPointerException in ConnectionManager.locateRegionInMeta() due to missing region info
- [HBASE-10356](#) Failover RPC's for multi-get.
- [HBASE-10525](#) Allow the client to use a different thread for writing to ease interrupt.
- [HBASE-10355](#) Failover RPC's from client using region replicas.
- [HBASE-10352](#) Region and RegionServer changes for opening region replicas, and refreshing store files
- [HBASE-10351](#) LoadBalancer changes for supporting region replicas.
- [HBASE-10359](#) Master/RS WebUI changes for region replicas.
- [HBASE-10362](#) HBACK changes for supporting region replicas.
- [HBASE-10361](#) Enable/AlterTable support for region replicas.
- [HBASE-10350](#) Master/AM/RegionStates changes to create and assign region replicas.
- [HBASE-10490](#) Simplify RpcClient code

- [HBASE-10517](#) NPE in MetaCache.clearCache()
- [HBASE-10479](#) HConnection interface is public but is used internally, and contains a bunch of methods
- [HBASE-10348](#) HTableDescriptor changes for region replicas
- [HBASE-10354](#) Add an API for defining consistency per request
- [HBASE-10347](#) HRegionInfo changes for adding replicaId and MetaEditor/MetaReader changes for region replicas
- [HBASE-10277](#) refactor AsyncProcess
- [HBASE-10427](#) clean up HRegionLocation/ServerName usage
- [HBASE-10472](#) Manage the interruption in ZKUtil#getData

1.7.6. Hive 0.14.0

- [HIVE-8947](#) HIVE-8876 also affects Postgres < 9.2
- [HIVE-8901](#) increase retry attempt, interval on metastore database errors
- [HIVE-8890](#) HiveServer2 dynamic service discovery: use persistent ephemeral nodes curator recipe
- [HIVE-8880](#) non-synchronized access to split list in OrInputFormat
- [HIVE-8876](#) incorrect upgrade script for Oracle (13->14)
- [HIVE-8875](#) hive.optimize.sort.dynamic.partition should be turned off for ACID
- [HIVE-8813](#) Allow tests to be excluded based on pattern/regex
- [HIVE-8880](#) non-synchronized access to split list in OrInputFormat
- [HIVE-8830](#) hcatalog process don't exit because of non daemon thread
- [HIVE-8812](#) TestMinimrCliDriver failure if run in the same command as TestHBaseNegativeCliDriver
- [HIVE-8811](#) Dynamic partition pruning can result in NPE during query compilation
- [HIVE-8805](#) CBO skipped due to SemanticException: Line 0:-1 Both left and right aliases encountered in JOIN 'avg_cs_ext_discount_amt'
- [HIVE-8801](#) Make orc_merge_incompat1.q deterministic across platforms
- [HIVE-8797](#) Simultaneous dynamic inserts can result in "partition already exists" error
- [HIVE-8715](#) Hive 14 upgrade scripts can fail for statistics if database was created using auto-create
- [HIVE-8374](#) schematool fails on Postgres versions < 9.2

- [HIVE-8317](#) WebHCat pom should explicitly depend on jersey-core
- [HIVE-8266](#) create function using <resource> statement compilation should include resource URI entity
- [HIVE-8182](#) beeline fails when executing multiple-line queries with trailing spaces
- [HIVE-8130](#) Support Date in Avro
- [HIVE-8113](#) Derby server fails to start on windows
- [HIVE-6148](#) Support arbitrary structs stored in HBase

1.7.7. Knox 0.6.0

- [KNOX-470](#) Add README to Samples and Documentation
- [KNOX-468](#) Add default config to optimize LDAP group lookup
- [KNOX-467](#) Unit tests failing on windows
- [KNOX-466](#) Log exception stack traces at INFO level when they reach gateway servlet
- [KNOX-465](#) Initial audit record can contain leftover principal name
- [KNOX-464](#) Location headers have wrong hostname when used behind load balancer
- [KNOX-459](#) Apache Knox appears to be leaking sockets

1.7.8. Mahout 0.9.0

- [MAHOUT-1464](#) Cooccurrence Analysis on Spark
- [MAHOUT-1578](#) Optimizations in matrix serialization
- [MAHOUT-1572](#) blockify() to detect (naively) the data sparsity in the loaded data
- [MAHOUT-1571](#) Functional Views are not serialized as dense/sparse correctly
- [MAHOUT-1566](#) (Experimental) Regular ALS factorizer with conversion tests, optimizer enhancements and bug fixes
- [MAHOUT-1537](#) Minor fixes to spark-shell
- [MAHOUT-1529](#) Finalize abstraction of distributed logical plans from backend operations
- [MAHOUT-1489](#) Interactive Scala & Spark Bindings Shell & Script processor
- [MAHOUT-1346](#) Spark Bindings
- [MAHOUT-1555](#) Exception thrown when a test example has the label not present in training examples
- [MAHOUT-1446](#) Create an intro for matrix factorization

- [MAHOUT-1480](#) Clean up website on 20 newsgroups
- [MAHOUT-1561](#) cluster-syntheticcontrol.sh not running locally with MAHOUT_LOCAL=true
- [MAHOUT-1558](#) Clean up classify-wiki.sh and add in a binary classification problem
- [MAHOUT-1560](#) Last batch is not filled correctly in MultithreadedBatchItemSimilarities
- [MAHOUT-1554](#) Provide more comprehensive classification statistics
- [MAHOUT-1548](#) Fix broken links in quickstart webpage
- [MAHOUT-1542](#) Tutorial for playing with Mahout's Spark shell
- [MAHOUT-1533](#) Remove Frequent Pattern Mining
- [MAHOUT-1532](#) Add solve() function to the Scala DSL
- [MAHOUT-1530](#) Custom prompt and welcome message for the Spark Shell
- [MAHOUT-1527](#) Fix wikipedia classifier example
- [MAHOUT-1526](#) Ant file in examples
- [MAHOUT-1523](#) Remove @author tags in sparkbindings
- [MAHOUT-1521](#) lucene2seq - Error trying to load data from stored field (when non-indexed)
- [MAHOUT-1520](#) Fix links in Mahout website documentation
- [MAHOUT-1519](#) Remove StandardThetaTrainer
- [MAHOUT-1517](#) Remove casts to int in ALSWRFactorizer
- [MAHOUT-1513](#) Deprecate Canopy Clustering
- [MAHOUT-1511](#) Renaming core to mrlegacy
- [MAHOUT-1510](#) Goodbye MapReduce
- [MAHOUT-1509](#) Invalid URL in link from "quick start/basics" page
- [MAHOUT-1508](#) Performance problems with sparse matrices
- [MAHOUT-1505](#) structure of clusterdump's JSON output
- [MAHOUT-1504](#) Enable/fix thetaSummer job in TrainNaiveBayesJob
- [MAHOUT-1503](#) TestNaiveBayesDriver fails in sequential mode
- [MAHOUT-1502](#) Update Naive Bayes Webpage to Current Implementation
- [MAHOUT-1501](#) ClusterOutputPostProcessorDriver has private default constructor

- [MAHOUT-1498](#) DistributedCache.setCacheFiles in DictionaryVectorizer overwrites jars pushed using oozie
- [MAHOUT-1497](#) mahout resplit not producing splited files
- [MAHOUT-1496](#) Create a website describing the distributed ALS recommender
- [MAHOUT-1491](#) Spectral KMeans Clustering doesn't clean its /tmp dir and fails when seeing it again
- [MAHOUT-1488](#) DisplaySpectralKMeans fails: examples/output/clusteredPoints/part-m-00000 does not exist
- [MAHOUT-1483](#) Organize links in web site navigation bar
- [MAHOUT-1482](#) Rework quickstart website
- [MAHOUT-1476](#) Cleanup website on Hidden Markov Models
- [MAHOUT-1475](#) Cleanup website on Naive Bayes
- [MAHOUT-1472](#) Cleanup website on fuzzy kmeans
- [MAHOUT-1471](#) Cleanup website for Canopy clustering
- [MAHOUT-1468](#) Creating a new page for StreamingKMeans documentation on mahout website
- [MAHOUT-1467](#) ClusterClassifier readPolicy leaks file handles
- [MAHOUT-1466](#) Cluster visualization fails to execute
- [MAHOUT-1465](#) Clean up README
- [MAHOUT-1463](#) Modify OnlineSummarizers to use the TDigest dependency from Maven Central
- [MAHOUT-1460](#) Remove reference to Dirichlet in ClusterIterator
- [MAHOUT-1459](#) Move Hadoop related code out of CanopyClusterer
- [MAHOUT-1458](#) Remove KMeansConfigKeys and FuzzyKMeansConfigKeys
- [MAHOUT-1457](#) Move EigenSeedGenerator into spectral kmeans package
- [MAHOUT-1455](#) Forkcount config causes JVM crashes during build
- [MAHOUT-1451](#) Cleaning up the examples for clustering on the website
- [MAHOUT-1450](#) Cleaning up clustering documentation on mahout website
- [MAHOUT-1449](#) Update the Known Issues in Random Forests Page
- [MAHOUT-1448](#) In Random Forest, the training does not support multiple input files. The input dataset must be one single file.

- [MAHOUT-1447](#) ImplicitFeedbackAlternatingLeastSquaresSolver tests and features
- [MAHOUT-1445](#) Create an intro for item based recommender
- [MAHOUT-1440](#) Add option to set the RNG seed for initial cluster generation in Kmeans/fKmeans
- [MAHOUT-1438](#) "quickstart" tutorial for building a simple recommender
- [MAHOUT-1434](#) Dead links on the web site
- [MAHOUT-1433](#) Make SVDRecommender look at all unknown items of a user per default
- [MAHOUT-1429](#) Parallelize YtransposeY in ImplicitFeedbackAlternatingLeastSquaresSolver
- [MAHOUT-1428](#) Recommending already consumed items
- [MAHOUT-1425](#) SGD classifier example with bank marketing dataset.
- [MAHOUT-1420](#) Add solr-recommender to examples
- [MAHOUT-1419](#) Random decision forest is excessively slow on numeric features
- [MAHOUT-1417](#) Random decision forest implementation fails in Hadoop 2
- [MAHOUT-1416](#) Make access of DecisionForest.read(dataInput) less restricted
- [MAHOUT-1415](#) Clone method on sparse matrices fails if there is an empty row which has not been set explicitly
- [MAHOUT-1413](#) Rework Algorithms page
- [MAHOUT-1388](#) Add command line support and logging for MLP
- [MAHOUT-1385](#) Caching Encoders don't cache
- [MAHOUT-1356](#) Ensure unit tests fail fast when writing outside mvn target directory
- [MAHOUT-1329](#) Mahout for hadoop 2
- [MAHOUT-1310](#) Mahout support windows

1.7.9. Oozie 4.1.0

- [OOZIE-2051](#) Allows multiple default action configuration files per action
- [OOZIE-2013](#) Bundle job with children status of "KILLED" and "FAILED" should transition to "KILLED" instead of "DNEWITHERROR"
- [OOZIE-2012](#) coordinator with an invalid cron frequency throws NPE after validation
- [OOZIE-2011](#) In non-secure environment, oozie fails "with E0508: User [?] not authorized for WF job" when killing a job

- [OOZIE-1814](#) Oozie should mask "oozie.https.keystore.pass" in instrumentation log
- [OOZIE-1726](#) Oozie does not support _HOST when configuring kerberos security
- [OOZIE-1561](#) When using Oozie HA, the logs should also be HA

1.7.10. Phoenix 4.2.0

- [PHOENIX-1442](#) Alter Index double normalize Index Table Name
- [PHOENIX-1434](#) Stats are not maintained correctly through splits
- [PHOENIX-1416](#) Given a schema name, DatabaseMetadata.getTableNames and getColumnNames calls erroneously match tables without schema
- [PHOENIX-1406](#) Use preCompact hook instead of preCompactScannerOpen for stats collection
- [PHOENIX-1405](#) Problem referencing lower-case column names with Phoenix / Pig / Spark
- [PHOENIX-1404](#) Ignore call to ResultSet#setFetchSize
- [PHOENIX-1403](#) Correct synchronization in ConnectionQueryServicesImpl.init
- [PHOENIX-1402](#) Don't recalculate stats on split
- [PHOENIX-1401](#) SYSTEM.SEQUENCE table is not pre-split as expected
- [PHOENIX-1400](#) Copy small attribute in ScanUtil
- [PHOENIX-1397](#) RVC combined with OR on first row key column results in NPE
- [PHOENIX-1396](#) Avoid unnecessary ResultSpooler.bin files creation/deletion every time SpoolingResultIterator
- [PHOENIX-1395](#) ResultSpooler spill files are left behind in /tmp folder
- [PHOENIX-1346](#) Base exception gets wrapped as incompatible jar exception

1.7.11. Pig 0.14.0

- [PIG-4138](#) Upgrade hbase to hbase-0.98.4

1.7.12. Ranger 0.4.0 (formerly XASecure)

- [RANGER-186](#) Improve failure handling in usersync service
- [RANGER-185](#) Optimize database transaction usage in admin web application
- [RANGER-183](#) Search policies public API call failing in Oracle DB
- [RANGER-182](#) Date format for Audit log is incorrect.

- [RANGER-180](#) Ranger lookup for hadoop repository fails when wire encryption is on
- [RANGER-158](#) SSL service should not accept SSLv3 protocol (due to CVE-2014-3566)

1.7.13. Sqoop 1.4.5

- [SQOOP-1759](#) TestIncrementalImport fails with NPE on Windows
- [SQOOP-1540](#) Accumulo unit tests fail with Accumulo 1.6.1 because of conflicts in libthrift libraries
- [SQOOP-1539](#) Empty partition keys and values in multiple partition key hcatalog usage should be validated
- [SQOOP-1519](#) Enable HCat/HBase/Accumulo operations with OraOop connection manager
- [SQOOP-1474](#) Fix usage of StringUtils.repeat method
- [SQOOP-1471](#) Use Hadoop CredentialProvider API to encrypt passwords at rest
- [SQOOP-1400](#) Failed to import data using mysql-connector-java-5.1.17.jar

1.7.14. Storm 0.9.3

- [STORM-216](#) Secure Multi-tenant Storm

1.7.15. Tez 0.5.2

- [TEZ-1774](#) AppLaunched event for Timeline does not have start time set.
- [TEZ-1770](#) Handle ConnectExceptions correctly when establishing connections to an NM which may be down.
- [TEZ-1761](#) TestRecoveryParser::testGetLastInProgressDAG fails in similar manner to TEZ-1686.
- [TEZ-1758](#) TezClient should provide YARN diagnostics when the AM crashes
- [TEZ-1750](#) Add a DAGScheduler which schedules tasks only when sources have been scheduled.
- [TEZ-1749](#) Increase test timeout for TestLocalMode.testMultipleClientsWithSession
- [TEZ-1745](#) TestATSHistoryLoggingService::testATSHistoryLoggingServiceShutdown can be flaky.
- [TEZ-1747](#) Increase test timeout for TestSecureShuffle.
- [TEZ-1746](#) Flaky test in TestVertexImpl and TestExceptionPropagation.
- [TEZ-1742](#) Improve response time of internal preemption.

1.7.16. Zookeeper 3.4.6

- [ZOOKEEPER-2064](#) Prevent resource leak in various classes
- [ZOOKEEPER-2053](#) Make scripts use ZOOKEEPER_HOME if set.
- [ZOOKEEPER-1926](#) Unit tests should only use build/test/data for data
- [ZOOKEEPER-1897](#) ZK Shell/Cli not processing commands
- [ZOOKEEPER-1876](#) Add support for installing windows services in .cmd scripts
- [ZOOKEEPER-1888](#) ZkCli.cmd commands fail with "'java' is not recognized as an internal or external command"
- [ZOOKEEPER-1848](#) [WINDOWS] Java NIO socket channels does not work with Windows ipv6 on JDK6
- [ZOOKEEPER-1791](#) ZooKeeper package includes unnecessary jars that are part of the package.
- [ZOOKEEPER-1575](#) adding .gitattributes to prevent CRLF and LF mismatches for source and text files
- [ZOOKEEPER-1574](#) Set svn:eol-style property to native.

1.8. Known Issues for HDP 2.2

HDP 2.2.0 has the following known issues, scheduled for resolution in a future release. Where available, a workaround has been provided.

1.8.1. Falcon Known Issues

- [BUG-26944: E0803: IO error, null during scheduling feed on a cluster with wire encryption on.](#)

Problem: Falcon has difficulty writing to target clusters with wire encryption on. The ssl-client on the CP must be valid for doing SSL negotiation with the target cluster.

Workaround: Have Falcon write to target clusters with wire encryption disabled.

- [BUG-25179: Falcon's default port is 15443 if TLS is enabled, 15000 otherwise.](#)

Problem: Falcon server is on port 15443 or 15000 based on setting for TLS. In Falcon startup.properties:

- If falcon.enableTLS is set to true explicitly or not set at all, falcon starts at port 15443 on https:// by default.
- If falcon.enableTLS is set to false, falcon starts at port 15000 on http://.

Workaround: To change the port, use -port option. If falcon.enableTLS is not set explicitly but port is set explicitly, port that ends with 443 will automatically put falcon on https://. Any other port will put falcon on http://.

- **BUG-25179:** You must add your own keystore to `/etc/falcon/conf/` directory as `prism.keystore`, when `falcon.enableTLS` is set to `true`.
- **BUG-28055:** Ambari currently configures the config store to local fs in `"*.config.store.uri"` (startup properties) and `"Falcon store URI"` (falcon server).

```
file:///hadoop/falcon/store
```

It is best to write to HDFS. In Ambari, modify your Falcon config to store to HDFS.

Manage the Falcon service > Select Config > in Falcon Store URI, specify a link to HDFS.
For example:

```
*.config.store.uri=hdfs://ip-172-31-13-179.ec2.internal:8020/apps/falcon
```

1.8.2. Hadoop Common/HDFS Known Issues

- **BUG-30030:** Incomplete installation of `hadoop-httpfs`.

Problem: `hadoop-httpfs` RPM scripts are incomplete for HDP 2.2. This affects proxy setup that allows Hue to interact with HA clusters.

Workaround: Have Falcon write to target clusters with wire encryption disabled.

Workaround: The following steps should be executed on the node that will be hosting the HttpFS service. This assumes that HDP 2.2 has already been deployed and configured.

1. Install the `hadoop_2_2_0_0_2041-httpfs` package:

```
yum install hadoop-httpfs
```

2. Create a symlink in `/usr/lib` to the location where `httpfs` is installed:

```
ln -s /usr/hdp/current/hadoop-client/./hadoop-httpfs /usr/lib/hadoop-httpfs
```

3. Create a symlink to `libexec`:

```
ln -s /usr/hdp/current/hadoop-client/./hadoop/libexec /usr/lib/hadoop-httpfs/libexec
```

4. Create configuration directories and copy the default set of configurations:

```
mkdir -p /etc/hadoop-httpfs/conf
mkdir -p /etc/hadoop-httpfs/tomcat-deployment/conf
cp -r /usr/hdp/2.2.0.0-2041/etc/hadoop-httpfs/conf.empty/* /etc/hadoop-httpfs/conf/
cp -r /usr/hdp/2.2.0.0-2041/etc/hadoop-httpfs/tomcat-deployment.dist/conf/* /etc/hadoop-httpfs/tomcat-deployment/conf
```

5. Enable the `init.d` script for HttpFS by creating the following symlink:

```
ln -s /usr/hdp/2.2.0.0-2041/etc/rc.d/init.d/hadoop-httpfs /etc/init.d/hadoop-httpfs
```

6. Create the following symlinks to point to the Tomcat libraries and binary scripts that HttpFS will use:

```
ln -s /usr/lib/bigtop-tomcat/lib /etc/hadoop-httpfs/tomcat-deployment/  
libln -s /usr/lib/bigtop-tomcat/bin /etc/hadoop-httpfs/tomcat-deployment/  
bin
```

7. Start the HttpFs server:

```
service hadoop-httpfs start
```

1.8.3. HBase Known Issues

- BUG-26592: Multiple tests are failing because of System time goes back frequently on SLES

Problem: SUSE 11 SP3 clusters running on Amazon EC2 have had the system clock frequently going back (hundreds of ms), which causes multiple issues in HBase. The HBase data model assumes that timestamps are monotonically increasing, and failing that may result in updates being eclipsed by previous mutations.

Workaround: Setup NTP and ensure that system clock time does not drift when running HBase. See <https://hbase.apache.org/book.html#basic.prerequisites> for more details.

- BUG-23981 [Windows]: IOException: Invalid HFile block magic in Windows unit tests

If you use HBase on a local filesystem (not HDFS) on Windows for single node non-distributed clusters, concurrent hfile readers might return data belonging to a different offset than what was requested, resulting in read failures.

1.8.4. HDFS Known Issues

- BUG-26890: Commands in HA mode behave differently compared to non-HA

Problem: Commands in HA mode behave differently compared to non-HA. For example:

```
hdfs dfsadmin -safemode enter - if nn2 is active and nn1 is standby, command  
will run with no errors, but it will set standby to safemode and not active. If nn1 is  
actually off, the command will try nn2 and the command in effect will do what it needs  
to do.
```

The same behavior is true for `hdfs dfsadmin -saveNamespace`.

Workaround: You can add the `fs` and specify `hdfs://hostname:8020` of the active namenode, for example: `hdfs dfsadmin -fs hdfs://<namenodehostactive>:8020 -saveNamespace`.

1.8.5. Hive Known Issues

- BUG-28161, BUG-28234: Hive Server2 Long running test failed after 44 hours with OutOfMemory

Problem: Hive Server 2 out of memory due to FileSystem.CACHE leak

Workaround: To prevent memory leaks in unsecure mode, disable file system caches by setting the following parameters to true in hive configuration (hive-site.xml).

`fs.hdfs.impl.disable.cache` – Disable HDFS filesystem cache, default false.

`fs.file.impl.disable.cache` – Disable local filesystem cache, default false.

- **BUG-27913 Error while upgrading Hive Metastore schema for Oracle**

Problem: The pre-upgrade script causes an error when upgrading Hive Metastore using an Oracle DB. For example, after executing the command `-upgradeSchema`, the following error message occurs:

```
root@ip-172-31-42-3 ~# /usr/hdp/2.2.0.0-2036/hive/bin/schematool -
upgradeSchema -dbType oracle
14/11/17 14:11:38 WARN conf.HiveConf: HiveConf of name
hive.optimize.mapjoin.mapreduce does not exist
14/11/17 14:11:38 WARN conf.HiveConf: HiveConf of name hive.heapsize
does not exist
14/11/17 14:11:38 WARN conf.HiveConf: HiveConf of name
hive.server2.enable.impersonation does not exist
14/11/17 14:11:38 WARN conf.HiveConf: HiveConf of name
hive.semantic.analyzer.factory.impl does not exist
14/11/17 14:11:38 WARN conf.HiveConf: HiveConf of name
hive.auto.convert.sortmerge.join.noconditionaltask does not exist
Metastore connection URL: jdbc:oracle:thin:@//ip-172-31-42-
1.ec2.internal:1521/XE
Metastore Connection Driver : oracle.jdbc.driver.OracleDriver
Metastore connection User: hiveuser
Starting upgrade metastore schema from version 0.13.0 to 0.14.0
Upgrade script upgrade-0.13.0-to-0.14.0.oracle.sql
Error: ORA-00955: name is already used by an existing object
(state=42000,code=955)
Warning in pre-upgrade script pre-0-upgrade-0.13.0-to-
0.14.0.oracle.sql: Schema script failed, errorCode 2
Completed upgrade-0.13.0-to-0.14.0.oracle.sql
schemaTool completed
```

Workaround: Ignore the error message. The schema tool has upgraded.

- **BUG-27637: Postgres Metastore has intermittent test failure due to "ERROR: could not serialize access due to read/write dependencies among transactions"**

Problem: Sometimes the `acid_concurrency` test on Postgres Metastore test fails with this error:

```
ERROR: could not serialize access due to read/write dependencies among
transactions
Detail: Reason code: Canceled on identification as a pivot, during write.
Hint: The transaction might succeed if retried.
at org.postgresql.core.v3.QueryExecutorImpl.
receiveErrorResponse(QueryExecutorImpl.java:2198)
at org.postgresql.core.v3.QueryExecutorImpl.
processResults(QueryExecutorImpl.java:1927)
[...]
ERROR: could not serialize access due to read/write dependencies among
transactions
Detail: Reason code: Canceled on identification as a pivot, during write.
Hint: The transaction might succeed if retried.
org.postgresql.core.v3.QueryExecutorImpl.
receiveErrorResponse(QueryExecutorImpl.java:2198)
```

```
org.postgresql.core.v3.QueryExecutorImpl.processResults(QueryExecutorImpl.
java:1927)
[...]
ERROR [HiveServer2-Background-Pool: Thread-392]: SessionState (SessionState.
java:printError(833)) - Execution has failed.
2014-11-12 08:40:57,813
INFO [HiveServer2-Background-Pool: Thread-392]: log.PerfLogger (PerfLogger.
java:PerfLogEnd(135))
[...]
ERROR [HiveServer2-Background-Pool: Thread-392]: ql.Driver (SessionState.
java:printError(833)) - FAILED: Execution Error, return code 1 from org.
apache.hadoop.hive.ql.exec.tez.TezTask
2014-11-12 08:40:57,869
INFO [HiveServer2-Background-Pool: Thread-392]: log.PerfLogger (PerfLogger.
java:PerfLogEnd(135))
[...]
```

- **BUG-27636 Oracle Metastore: intermittent acid_concurrency test failures due to NoSuchLockException**

```
ERROR [HiveServer2-Background-Pool: Thread-1229]:
metastore.RetryingHMSHandler (RetryinHMSHandler.java:invoke(141)) -
NoSuchLockException(message:No such lock: 59) at org.apache.hadoop.hive.
metastore.txn.
TxnHandler.heartbeatLock(TxnHandler.java:1532) at org.apache.hadoop.hive.
metastore.txn.
TxnHandler.heartbeat(TxnHandler.java:608)
[...]
```

- **BUG-27197: MatchPath not available in HDP 2.2.**

Intellectual property issues prevent Hortonworks from shipping MatchPath in Hive 0.14.x.

- **BUG-26418: SSLv3 not supported in Hive 0.14.x**

Problem: SSLv3 is vulnerable to the Poodle Attack and is therefore disabled in Hive 0.14.x. Do not enable SSLv3 in Hive.

1.8.6. Knox Known Issues

- **BUG-24945: [KNOX-443] Support for active HS2 clients (through Knox) during rolling upgrades will be addressed in a later release.**

1.8.7. Oozie Known Issues

- **BUG-27982 Oozie client package installation from public HDP repo fails because of hdp-select**

Problem: Oozie-client is missing a dependency on hdp-select. If you install oozie-client prior to any other stack applications, installation fails with with an error message similar to the following:

```
Installing: oozie_2_2_0_0_1947-client-4.1.0.2.2.0.0-1947 [error]
Installation of oozie_2_2_0_0_1947-client-4.1.0.2.2.0.0-1947 failed:
(with --nodeps --force) Error: Subprocess failed. Error: RPM failed:
warning: group jenkins does not exist - using root
```

```
warning: group jenkins does not exist - using root
warning: group jenkins does not exist - using root
warning: group jenkins does not exist - using root
warning: group jenkins does not exist - using root
warning: group jenkins does not exist - using root
warning: group jenkins does not exist - using root
warning: group jenkins does not exist - using root
warning: group jenkins does not exist - using root
warning: group jenkins does not exist - using root
warning: group jenkins does not exist - using root
warning: group jenkins does not exist - using root
warning: group jenkins does not exist - using root
warning: group jenkins does not exist - using root
warning: group jenkins does not exist - using root
/var/tmp/rpm-tmp.48695: line 6: /usr/bin/hdp-select: No such file or
directory
error: %post(oozie_2_2_0_0_1947-client-4.1.0.2.2.0.0-1947.noarch)
scriptlet failed, exit status 127
Abort, retry, ignore? [a/r/i] (a):
```

Workaround: If you only intend to install oozie-client and no other HDP 2.2 components, install and run hdp-select first.

1.8.8. Phoenix Known Issues

- BUG-26661: Windows runs failed with "region isn't co-located with its data" exception

Problem: The Local Index feature is considered experimental and should not be enabled at this time, except for preview purposes.

- BUG-28226: Error 504 after upgrading from HDP 2.1 to HDP 2.2.

Problem: After upgrading from HDP 2.1 to HDP 2.2, queries such as the following generate an ERROR 504: `./psql.py ip-172-31-43-221.ec2.internal:2181 /grid/0/hdp/2.2.0.0- 2041/ phoenix/doc/examples/WEB_STAT.sql /grid/0/hdp/2.2.0.0- 2041/ phoenix/doc/examples/WEB_STAT.csv /grid/0/hdp/2.2.0.0- 2041/ phoenix/doc/examples/WEB_STAT_QUERIES.sql.`

Error detail: `org.apache.phoenix.schema.ColumnNotFoundException: ERROR 504 (42703): Undefined column. columnName=INDEX_TYPE`

Workaround: Follow instructions in the "Manual Upgrade to HDP 2.2" document, "Configure Phoenix" section, step 4: replace the client jar file before connecting the new Phoenix client to the Phoenix cluster.

1.8.9. Pig Known Issues

- Nested fields will experience slower performance because they are not currently supported by predicate pushdown in Pig.

1.8.10. Ranger (formerly XA Secure) Known Issues

- BUG-28102: Update to headers needed for `/etc/init.d/ranger-admin` and `/etc/init.d/ranger-usersyncscripts` after Ranger install on Debian 6.

Problem: The init scripts installed by Ranger for admin and usersync are missing a comment block at the beginning (details <https://wiki.debian.org/LSBInitScripts/>), resulting in error messages.

Workaround: Add the following header to these scripts:

```
### BEGIN INIT INFO
# Provides:          scriptname
# Required-Start:    $remote_fs $syslog
# Required-Stop:     $remote_fs $syslog
# Default-Start:     2 3 4 5
# Default-Stop:      0 1 6
# Short-Description: Start daemon at boot time
# Description:       Enable service provided by daemon.
### END INIT INFO
```

1.8.11. Sqoop Known Issues

- BUG-25282: `j.l.ArrayIndexOutOfBoundsException` using `--map-column-hive`

Problem: Running this command with the `--map-column-hive` statement:

```
sqoop import --verbose --connect
'jdbc:mysql://sandbox.hortonworks.com/demo' --table customer_address --
username carter --password password --create-hcatalog-table --hcatalog-
table customer_address --hcatalog-storage-stanza "stored as orc" --map-
column-hive fdsa -m 1
```

Returns:

```
2014-10-17 21:49:27,597 DEBUG sqoop.Sqoop (Sqoop.java:run(137)) - 1
java.lang.ArrayIndexOutOfBoundsException: 1
at org.apache.sqoop.SqoopOptions.parseColumnMapping(SqoopOptions.java:1196)
at org.apache.sqoop.SqoopOptions.setMapColumnHive(SqoopOptions.java:1201)
at org.apache.sqoop.tool.BaseSqoopTool.applyHiveOptions(BaseSqoopTool.
java:1111)
at org.apache.sqoop.tool.ImportTool.applyOptions(ImportTool.java:961)
at org.apache.sqoop.tool.SqoopTool.parseArguments(SqoopTool.java:433)
at org.apache.sqoop.Sqoop.run(Sqoop.java:131)
at org.apache.hadoop.util.ToolRunner.run(ToolRunner.java:70)
at org.apache.sqoop.Sqoop.runSqoop(Sqoop.java:179)
at org.apache.sqoop.Sqoop.runTool(Sqoop.java:218)
at org.apache.sqoop.Sqoop.runTool(Sqoop.java:227)
at org.apache.sqoop.Sqoop.main(Sqoop.java:236)
```

There is no `fdsa` column. The command works without the `--map-column-hive` statement.

Workaround: Sqoop is expecting comma separated list of mapping in form `<name of column>=<new type>`. For example: `$ sqoop import ... --map-column-java id=String,value=Integer`

1.8.12. YARN Known Issues

- BUG-27604: After upgrade from 2.1 of Highly Available system, ResourceManager fails to elect a master.

Problem: After upgrading an HA system from HDP 2.2 from HDP 2.1, you must manually clear the ResourceManager state store as part of the upgrade or the following error message will occur:

```
Exception transitioning to active PERMISSIONS=All users are allowed
2014-11-12 05:36:04,941 WARN ha.ActiveStandbyElector
(ActiveStandbyElector.java:becomeActive(809)) - Exception handling the
winning of election
org.apache.hadoop.ha.ServiceFailedException: RM could not transition to
Active at org.apache.hadoop.yarn.server.resourcemanager.
EmbeddedElectorService.
becomeActive(EmbeddedElectorService.java:128)
[...]
```

Workaround: Manually clear the ResourceManager state store:

1. Run:

```
yarn resourcemanager -format-state-store
```

2. Start the ResourceManager:

```
su -l yarn -c "/usr/hdp/current/hadoop-yarn-
resourcemanager/sbin/yarn-daemon.sh start resourcemanager"
```

1.8.13. ZooKeeper Known Issues

- **BUG-19351:** Zookeeper does not have shell access on Ubuntu and Debian Hosts.

Description: The /etc/passwd file has no shell access for Zookeeper:

```
zookeeper:x:108:115:ZooKeeper User,,,:/var/lib/zookeeper:/bin/false
```

Workaround: Modify the zookeeper user to have shell access. For example:

```
usermod zookeeper -s /bin/bash
```

- **BUG-28102:** Update to headers needed for /etc/init.d/ranger-admin and /etc/init.d/ranger-usersyncscripts after Ranger install on Debian 6.

Problem: The init scripts installed by Ranger for admin and usersync are missing a comment block at the beginning (details <https://wiki.debian.org/LSBInitScripts/>), resulting in error messages.

Workaround: Add the following header to these scripts:

```
### BEGIN INIT INFO
# Provides:          scriptname
# Required-Start:    $remote_fs $syslog
# Required-Stop:     $remote_fs $syslog
# Default-Start:     2 3 4 5
# Default-Stop:      0 1 6
# Short-Description: Start daemon at boot time
# Description:       Enable service provided by daemon.
### END INIT INFO
```