

Release Notes 3

HDP 3.1.4 Release Notes

Date of Publish: 2019-08-26



<https://docs.hortonworks.com>

Contents

| | |
|--|-----------|
| HDP 3.1.4 Release Notes..... | 4 |
| Component Versions..... | 4 |
| Descriptions of New Features..... | 5 |
| Deprecation Notices..... | 6 |
| Terminology..... | 6 |
| Removed Components and Product Capabilities..... | 6 |
| Testing Unsupported Features..... | 6 |
| Descriptions of the Latest Technical Preview Features..... | 7 |
| Upgrading to HDP 3.1.4..... | 7 |
| Behavioral Changes..... | 7 |
| Apache Patch Information..... | 11 |
| Accumulo..... | 11 |
| Atlas..... | 12 |
| Calcite..... | 16 |
| DataFu..... | 16 |
| Hadoop..... | 16 |
| HBase..... | 22 |
| Hive..... | 26 |
| Kafka..... | 33 |
| Knox..... | 34 |
| Livy..... | 34 |
| Oozie..... | 34 |
| Phoenix..... | 38 |
| Pig..... | 38 |
| Ranger..... | 39 |
| Spark..... | 41 |
| Sqoop..... | 42 |
| Storm..... | 42 |
| Tez..... | 44 |
| Zeppelin..... | 45 |
| ZooKeeper..... | 45 |
| Fixed Common Vulnerabilities and Exposures..... | 46 |

| | |
|-------------------------------|-----------|
| Fixed Issues..... | 48 |
| Known Issues..... | 64 |
| Legal Information..... | 76 |

HDP 3.1.4 Release Notes

This document provides you with the latest information about the Hortonworks Data Platform (HDP) 3.1.4 release.

Component Versions

List of the official Apache component versions for Hortonworks Data Platform (HDP) 3.1.4. To ensure that you are working with the most recent stable software available, you must be familiar with the latest Apache component versions in HDP 3.1.4. You should also be aware of the available Technical Preview components and use them only in a testing environment.

The Hortonworks approach is to provide patches only when necessary, to ensure the interoperability of components. Unless you are explicitly directed by Hortonworks Support to install a patch, each of the HDP components should remain at the following package version levels to ensure a certified and supported copy of HDP 3.1.4.

Official Apache component versions for HDP 3.1.4:

- Apache Accumulo 1.7.0
- Apache Atlas 1.1.0
- Apache Calcite 1.16.0
- Apache DataFu 1.3.0
- Apache Druid 0.12.1 (incubating)
- Apache Hadoop 3.1.1
- Apache HBase 2.0.2
- Apache Hive 3.1.0
- Apache Kafka 2.0.0
- Apache Knox 1.0.0
- Apache Livy 0.5.0
- Apache Oozie 4.3.1
- Apache Phoenix 5.0.0
- Apache Pig 0.16.0
- Apache Ranger 1.2.0
- Apache Spark 2.3.2
- Apache Sqoop 1.4.7
- Apache Storm 1.2.1
- Apache TEZ 0.9.1
- Apache Zeppelin 0.8.0
- Apache ZooKeeper 3.4.6

You can use the following [Technical Previews](#) components, but only in a testing environment:

- Apache Superset 0.23.3



Note: For information about open source software licensing and notices, refer to the Licenses and Notices files included with the software installation package.

Related reference

[Fixed Common Vulnerabilities and Exposures](#)

[Apache Patch Information](#)

[Fixed Issues](#)

Descriptions of New Features

Hortonworks Data Platform (HDP) 3.1.4 provides you with several new capabilities. You must learn about the new features and their benefits to use the new features in HDP 3.1.4.

| Apache component | Feature areas, benefits, and links to further information |
|------------------|---|
| Atlas | <p>Many fixes on top of Apache Atlas 1.1.0, including these new features:</p> <ul style="list-style-type: none"> • Support for process execution entities that represent separate execution instances for a single process. • Hive Metastore hook for collecting metadata. • Bulk API for retrieving entities by unique attributes. |
| Hive | <p>Core Capabilities</p> <ul style="list-style-type: none"> • Automatic partition management synchronizes changes in the metadata and on the file system. • You can now configure how long you can retain partition data and metadata. • Hive Warehouse Connector now validates mapping of columns against those in Hive to alert the user to input errors. • Writing a DataFrame to Hive supports specifying partitions. • A new MergeBuilder interface for HiveWarehouseSession API operations supports merging tables. <p>Upgrade</p> <ul style="list-style-type: none"> • HiveStrictManagedMigration has two new options: <p>Specify the number of threads for processing tables in parallel. Default: The number of CPU cores.</p> <p>Specify the the type of tables to process. For example, MANAGED_TABLE, EXTERNAL_TABLE. Default: All tables.</p> |
| Kakfa | <p>Core Capabilities</p> <ul style="list-style-type: none"> • New Kafka Streams configuration parameter upgrade allows rolling bounce upgrade from an older version. • Updated ProcessorStateManager APIs in Kafka Streams for registering state stores to the processor topology. For more details please read the Streams Upgrade Guide. • Dynamic broker configuration options can be stored in ZooKeeper using kafka-configs.sh before brokers are started. This option can be used to avoid storing clear passwords in server.properties as all password configs may be stored encrypted in ZooKeeper. • ZooKeeper hosts are now re-resolved if the connection attempt fails. But if your ZooKeeper host names resolve to multiple addresses, and some of them are not reachable, you may need to increase the connection timeout zookeeper.connection.timeout.ms property. • Ability to define ACLs on prefixed resources, for example, any topic starting with 'foo'. |
| Ranger | <p>Core Capabilities</p> <ul style="list-style-type: none"> • Support for time-bound classification or business catalog mapping. Ranger policy engine recognizes the start and end times for tags and enforces the policies based on the period of validity of the tags. • Support for Safenet KeySecure crypto platform for storing master keys of Ranger KMS. • Support to allow users to define policy level conditions in Ranger. |
| Spark | <p>Integration</p> <ul style="list-style-type: none"> • Spark Kafka source integration with schema registry. |

| Apache component | Feature areas, benefits, and links to further information |
|------------------|---|
| Zeppelin | <p>Security</p> <ul style="list-style-type: none"> Support for SPNEGO based user authentication for zeppelin. With SPNEGO based authentication, LDAP/Active Directory based user search for notebook ACLs will continue to work. <p>Integration</p> <ul style="list-style-type: none"> Unbundle Zeppelin Notebook from Zeppelin installation. |

Deprecation Notices

When upgrading from one Hortonworks Data Platform release to another, you must be aware of any technology that is deprecated, moved, or removed from the latest release. Use the deprecation notices as a guide for your implementation plans.

Terminology

Items in this section are designated as follows:

Deprecated

Technology that Hortonworks is removing in a future HDP release. Marking an item as deprecated gives you time to plan for removal in a future HDP release.

Moving

Technology that Hortonworks is moving from a future HDP release and is making available through an alternative Hortonworks offering or subscription. Marking an item as moving gives you time to plan for removal in a future HDP release and plan for the alternative Hortonworks offering or subscription for the technology.

Removed

Technology that Hortonworks has removed from HDP and is no longer available or supported as of this release. Take note of technology marked as removed since it can potentially affect your upgrade plans.

Removed Components and Product Capabilities

No components are deprecated or removed in this HDP release.

Please contact Hortonworks Support or your Hortonworks Account Team if you have any questions.

Testing Unsupported Features

Unsupported features are those features that are not officially supported in HDP 3.1.4. Although the features listed in the *Descriptions of the Latest Technical Preview Features* sections exist within HDP, you must not use these features in production because Hortonworks does not support them.

| Apache Component | Unsupported Feature |
|------------------|--|
| YARN | move-application operation: Moving jobs between queues is not supported. |

Descriptions of the Latest Technical Preview Features

If you want to experiment with the Technical Preview features included in Hortonworks Data Platform 3.1.4, you should be aware that these features are not ready for production deployment. Hortonworks encourages you to explore these technical preview features in non-production environments and provide feedback on your experiences through the [Hortonworks Community Forums](#).

| Apache Component | Technical Preview Feature |
|------------------|---|
| YARN | YARN can now handle NEC VE devices as a resource. Workloads can now use these devices for your use cases. |

Upgrading to HDP 3.1.4

You must use Ambari to either do a major (HDP 2.6.x to HDP 3.1.4) or a minor (HDP 3.0.x to HDP 3.1.4) upgrade.

If you already have HDP 2.6.x or HDP 3.0.x installed, upgrading your cluster to HDP 3.1.0 means:

- Keeping the same configuration files you used for HDP 2.6.x or HDP 3.0.x.
- Keeping the same data and metadata in the same location you used for HDP 2.6.x or HDP 3.0.x.
- Installing any new components (added for the first time in HDP 2.6.x or HDP 3.0.x) side-by-side with existing components.

Before you begin

- Ensure that you know whether you want to do a major or a minor upgrade.
- Ensure that you know which HDP components you need to upgrade at your installation.
- Read the Prepare to Upgrade section in the *Ambari Upgrade Guide*.
- See the *Ambari Upgrade Guide* for upgrade information and instructions.

Related Information

[Ambari Major Upgrade Guide](#)

[Ambari Minor Upgrade Guide](#)

[Ambari Major Upgrade Guide for IBM Power Systems](#)

Behavioral Changes

Behavioral changes denote a marked change in behavior from the previously released version to this version of HDP.

Table 1: Behavioral Changes

| Bug ID | Apache Component | Apache JIRA | Summary | Details |
|------------|------------------|-----------------------------|---|---|
| BUG-118583 | Ranger | N/A | Allow users to define policy level conditions | <p>Previous behavior:</p> <p>Before this feature, users had provision to define policy conditions at individual policy items.</p> <p>New behavior:</p> <p>This allows users to define policy wide conditions. End users has flexibility to use both the features.</p> |
| BUG-119330 | Ranger | RANGER-2408 | With introduction of Security zones in Ranger. Users of different roles in Ranger have different permissions based on they being Zone Admin or Zone Auditor for particular set of zones. | <p>Previous behavior:</p> <p>Before this change, the user's capabilities were not managed by their involvement in zones. So if user is not part of any zone, it works as usual.</p> <p>New behavior:</p> <p>This feature takes care of restricting access to different APIs based on individual's role and their part in zone. See Apache Ranger Cwiki for more information.</p> |
| N/A | Kafka | KIP-186 | This change increases the default offset retention time from 1 day to 7 days. This makes it less likely to "lose" offsets in an application that commits infrequently. It also increases the active set of offsets and therefore can increase memory usage on the broker. | <p>Previous behavior:</p> <p>Default offset retention time was one day.</p> <p>New behavior:</p> <p>The console consumer currently enables offset commit by default, and can be the source of a large number of offsets. Now, the offsets are preserved for 7 days instead of 1. You can retain the existing behavior by setting the broker configuration <code>offsets.retention.minutes</code> to 1440.</p> |
| N/A | Kafka | N/A | Support for Java 7 has been dropped, Java 8 is now the minimum version required. | <p>Previous behavior:</p> <p>Support for Java 7.</p> <p>New behavior:</p> <p>Java 8 is now the minimum version required.</p> |
| N/A | Kafka | N/A | The default value for <code>ssl.endpoint.identification.algorithm</code> was changed to HTTPS, which performs hostname verification (man-in-the-middle attacks are possible otherwise). | <p>New behavior:</p> <p>The default value for <code>ssl.endpoint.identification.algorithm</code> was changed to HTTPS.</p> <p>Expected customer action:</p> <p>Set <code>ssl.endpoint.identification.algorithm</code> to an empty string to restore the previous behavior.</p> |
| KAFKA-5674 | Kafka | N/A | Extends the lower interval of <code>max.connections.per.ip</code> minimum to zero and therefore allows IP-based filtering of inbound connections. | <p>New behavior:</p> <p>Extends the lower interval of <code>max.connections.per.ip</code> minimum value to zero, and therefore allows IP-based filtering of inbound connections.</p> <p>Expected customer action:</p> <p>None.</p> |

| Bug ID | Apache Component | Apache JIRA | Summary | Details |
|--------|------------------|-------------|---|--|
| N/A | Kafka | KIP-272 | Added API version tag to the metric <code>kafka.network:type=RequestMetrics,name=RequestsPerSec,request={Produce FetchConsumer FetchFollower ...}</code> . This metric now becomes <code>kafka.network:type=RequestMetrics,name=RequestsPerSec,request={Produce FetchConsumer FetchFollower ...},version={0 1 2 3 ...}</code> . | <p>New behavior:</p> <p>This change will impact JMX monitoring tools that do not automatically aggregate.</p> <p>Expected customer action:</p> <p>To get the total count for a specific request type, the tool needs to be updated to aggregate across different versions.</p> |
| N/A | Kafka | KIP-225 | Changed the metric "records.lag" to use tags for topic and partition. | <p>Previous behavior:</p> <p>The original version of the name format was "{topic}-{partition}.records-lag".</p> <p>New behavior:</p> <p>The original version with the name format "{topic}-{partition}.records-lag" has now been removed.</p> |
| N/A | Kafka | N/A | The Scala consumers (deprecated since 0.11.0.0) are now removed. | <p>Previous behavior:</p> <p>The Scala consumers has been deprecated since version 0.11.0.0.</p> <p>New behavior:</p> <p>This is now removed.</p> <p>Expected customer action:</p> <p>The Scala consumers in 1.1.0 (and older) will continue to work even if the brokers are upgraded to 2.0.0.</p> |
| N/A | Kafka | N/A | The Scala producers, which have been deprecated since Kafka version 0.10.0.0, have now been removed. The Java producer has been the recommended option since 0.9.0.0. | <p>Previous behavior:</p> <p>The Scala producers has been deprecated since 0.10.0.0.</p> <p>New behavior:</p> <p>This is now removed. Note that the behavior of the default partitioner in the Java producer differs from the default partitioner in the Scala producers.</p> <p>Expected customer action:</p> <p>Users migrating should consider configuring a custom partitioner that retains the previous behavior. The Scala producers in 1.1.0 (and older) will continue to work even if the brokers are upgraded to 2.0.0.</p> |
| N/A | Kakfa | N/A | MirrorMaker and ConsoleConsumer no longer support the Scala consumer, they always use the Java consumer. | <p>New behavior:</p> <p>MirrorMaker and ConsoleConsumer always use the Java consumer.</p> <p>Expected customer action:</p> <p>None.</p> |
| N/A | Kafka | N/A | The ConsoleProducer no longer supports the Scala producer, it always uses the Java producer. | <p>New behavior:</p> <p>The ConsoleProducer always uses the Java producer.</p> <p>Expected customer action:</p> <p>None.</p> |

| Bug ID | Apache Component | Apache JIRA | Summary | Details |
|--------|------------------|-------------|---|---|
| N/A | Kafka | N/A | A number of deprecated tools that rely on the Scala clients are now removed. | <p>New behavior:</p> <p>Removed Scala clients: <code>ReplayLogProducer</code>, <code>SimpleConsumerPerformance</code>, <code>SimpleConsumerShell</code>, <code>ExportZkOffsets</code>, <code>ImportZkOffsets</code>, <code>UpdateOffsetsInZK</code>, <code>VerifyConsumerRebalance</code>.</p> <p>Expected customer action:</p> <p>None.</p> |
| N/A | Kafka | N/A | The deprecated <code>kafka.tools.ProducerPerformance</code> parameter is removed. | <p>New behavior:</p> <p><code>kafka.tools.ProducerPerformance</code> parameter is removed.</p> <p>Expected customer action:</p> <p>Use <code>org.apache.kafka.tools.ProducerPerformance</code></p> |
| N/A | Kafka | KIP-284 | Changed the retention time for Kafka Streams repartition topics. | <p>New behavior:</p> <p>Changed the retention time for Kafka Streams repartition topics. by setting its default value to <code>Long.MAX_VALUE</code>.</p> <p>Expected customer action:</p> <p>Use <code>org.apache.kafka.tools.ProducerPerformance</code></p> |
| N/A | Kafka | KIP-266 | Adds a new consumer configuration <code>default.api.timeout.ms</code> to specify the default timeout to use for <code>KafkaConsumer</code> APIs that could block. This change also adds overloads for such blocking APIs to support specifying a specific timeout to use for each of them instead of using the default timeout set by <code>default.api.timeout.ms</code> . | <p>Previous behavior:</p> <p>The value of <code>request.timeout.ms</code> parameter was a little higher than 5 minutes to account for maximum time that a rebalance would take.</p> <p>New behavior:</p> <p>The default value of <code>request.timeout.ms</code> has been changed to 30 seconds. Now we treat the <code>JoinGroup</code> request in the rebalance as a special case and use a value derived from <code>max.poll.interval.ms</code> for the request timeout. All other request types use the timeout defined by <code>request.timeout.ms</code>. A new <code>poll(Duration)</code> API has been added which does not block for dynamic partition assignment. The old <code>poll(long)</code> API has been deprecated and will be removed in a future version. Overloads have also been added for other <code>KafkaConsumer</code> methods such as <code>partitionsFor</code>, <code>listTopics</code>, <code>offsetsForTimes</code>, <code>beginningOffsets</code>, <code>endOffsets</code>.</p> <p>Expected customer action:</p> <p>None.</p> |
| N/A | Kafka | N/A | The internal method <code>kafka.admin.AdminClient.deleteRecordsBefore</code> is removed. | <p>New behavior:</p> <p><code>deleteRecordsBefore</code></p> <p>The internal method <code>kafka.admin.AdminClient.deleteRecordsBefore</code> has been removed.</p> <p>Expected customer action:</p> <p>None.</p> |
| N/A | Kafka | KIP-277 | The <code>AclCommand</code> tool, producer convenience option uses the finer grained ACL on the given topic. | <p>New behavior:</p> <p>The <code>AclCommand</code> tool, producer convenience option uses the finer grained ACL on the given topic.</p> <p>Expected customer action:</p> <p>None.</p> |

| Bug ID | Apache Component | Apache JIRA | Summary | Details |
|--------|------------------|----------------------------------|--|---|
| N/A | Kafka | KIP-176 | Removes the --new-consumer option for all consumer based tools. | <p>New behavior:</p> <p>The --new-consumeroption is removed because it is redundant since the new consumer is automatically used if --bootstrap-server is defined.</p> <p>Expected customer action:</p> <p>None.</p> |
| N/A | Kafka | KIP-283 | <p>Improved message down-conversion handling on Kafka broker, which has typically been a memory-intensive operation.</p> <p>New topic and broker configurations <code>message.downconversion.enable</code> and <code>log.message.downconversion.enable</code> are added to control whether down-conversion is enabled. When disabled, broker does not perform any down-conversion and instead sends an <code>UNSUPPORTED_VERSION</code> error to the client.</p> | <p>New behavior:</p> <p>Adds a mechanism by which the operation becomes less memory-intensive by down-converting chunks of partition data at a time which helps put an upper bound on memory consumption. With this improvement, there is a change in FetchResponse protocol behavior where the broker could send an oversized message batch towards the end of the response with an invalid offset. Such oversized messages must be ignored by consumer clients, as is done by <code>KafkaConsumer</code>.</p> <p>Expected customer action:</p> <p>None.</p> |
| N/A | ZooKeeper | ZOOKEEPER-2046 ZOOKEEPER-2046 | Authentication/authorization is enforced when a server attempts to join a quorum in Apache ZooKeeper | <p>Previous behavior:</p> <p>The command doesn't check any permission and returns all information contained in the ACL Id field as plaintext string.</p> <p>New behavior:</p> <p>The command now does check any permission and returns all information contained in the ACL Id field as plaintext string.</p> <p>Expected customer action:</p> <p>None.</p> |

Apache Patch Information

The following sections list patches in each HDP 3.1.4 component beyond what was fixed in the base version of the Apache component.

[Related Concepts](#)

[Component Versions](#)

Accumulo

This release provides Accumulo 1.7.0 and no additional Apache patches.

HDP 3.0.0, HDP 3.0.1, and HDP 3.1.0 provided Accumulo 1.7.0 and no additional Apache patches.

Atlas

This release provides Atlas 1.1 with most of the Apache patches that will eventually be included in Atlas 2.0.

HDP 3.1.4 provided Atlas 1.1.0 and the following Apache patches:

- [ATLAS-3377](#): Update AtlasPatchRegistry to use graph query instead of index query
- [ATLAS-3362](#): Import Service: Table-level Imports: Updating Replication Info Improvements
- [ATLAS-3348](#): Update jackson.databind.version to 2.9.9.1
- [ATLAS-3345](#): AggregationMetrics on quick search is empty
- [ATLAS-3338](#): Avoid multiple calls to get name of index-field for an attribute
- [ATLAS-3333](#): notification processing to provide an option to ignore Hive dummy databases and tables
- [ATLAS-3332](#): UI, Regression : In Relationship tab , empty valued attributes are shown even though "Show empty values" is unchecked
- [ATLAS-3331](#): Atlas Metrics UI should display details of multiple topics consumed by Atlas
- [ATLAS-3329](#): Update usage of Servlets.getUserName method to AtlasAuthorizationUtils.getCurrentUserName()
- [ATLAS-3324](#): Export API - incremental export with hive_table
- [ATLAS-3323](#): Atlas log4j configurations for file-size and file-backup are wrongly placed.
- [ATLAS-3322](#): Search Weight used for quick search is not getting applied properly.
- [ATLAS-3318](#): Throw exception when getGraphInstance() is unable to open janusgraph instance
- [ATLAS-3317](#): java.lang.Exception: Timeout scanning annotations while running Integration Tests
- [ATLAS-3316](#): getResolvedEntityVertex returns null if AtlasRelatedObjectId is referenced with AtlasObjectId
- [ATLAS-3313](#): Enhance Atlas server to process notifications from multiple Kafka topics
- [ATLAS-3308](#): Enhance Quicksearch API to support searchParameters using Post method.
- [ATLAS-3307](#): UT fix to add Tag propagation property which is not configured for all tags
- [ATLAS-3298](#): set hive_process startTime to System.currentTimeMillis() if startTime is null
- [ATLAS-3297](#): Cache created/updated entities in request context for use in entity-change listeners
- [ATLAS-3294](#): UI: Quick-search UI visual appearance enhancement
- [ATLAS-3293](#): LARGE_MESSAGES log file goes to {{log_dir}} instead of \${atlas.log.dir}
- [ATLAS-3292](#): Remove legacy Attributes
- [ATLAS-3291](#): UI: Don't allow user to sort on multiple value columns
- [ATLAS-3283](#): Export-import UTs are getting skipped
- [ATLAS-3282](#): UI : use search suggestions API for quick-search
- [ATLAS-3279](#): Performance regression after recent updates to disable full-text index
- [ATLAS-3278](#): Suggestions is not working for existing Atlas data instances.
- [ATLAS-3277](#): Add default atlas conf path in the import-* scripts
- [ATLAS-3276](#): Fix stale transactions in atlas due to ATLAS-3246 (Free-text search)
- [ATLAS-3275](#): Basic Search with "query" string throws "Problem accessing /solr/vertex_index/freetext" when SOLR version is 7.4.
- [ATLAS-3274](#): Updating relationship instance with a wrong guid throws 500 , expected is 400
- [ATLAS-3273](#): UI : Lineage Increase double click time
- [ATLAS-3270](#): Enhance Atlas to have a new rest end point to get suggestions given a prefix string.
- [ATLAS-3269](#): basic-search for entities with no-classifications and a search string returns empty result
- [ATLAS-3263](#): UI: Assign button (Create Term) should be blocked while an API call is in progress
- [ATLAS-3262](#): disable full-text index population when free-text is enabled
- [ATLAS-3256](#): Modify export API to process with relationshipAttributes
- [ATLAS-3251](#): Implement Patch to populate classification text for legacy data.
- [ATLAS-3246](#): Improve the free text search functionality to not depend on solrconfig.xml and schema.xml files.
- [ATLAS-3244](#): UI : Allows user to search the entities from any page (quick-search)
- [ATLAS-3231](#): Service specific icon should be shown for "process_execution"
- [ATLAS-3222](#): Update httpClient jar version to 4.5.3 to fix NoSuchMethodError error

- [ATLAS-3218](#): UI : entity detail page columns sorting empty value check
- [ATLAS-3212](#): Import Service: Importing to Existing Database with Different Name
- [ATLAS-3202](#): Hive hook: getStorageDescEntity() throws NPE if bucketCols is null
- [ATLAS-3200](#): add missing memory requirements for dev-support/atlas-docker/Dockerfile
- [ATLAS-3198](#): update/delete/truncate operations ATLAS_HOOK and ATLAS_ENTITIES improvements
- [ATLAS-3197](#): Add QueryText for hive_table and hive_db
- [ATLAS-3196](#): UI : Lineage support for IE.
- [ATLAS-3193](#): Import Service: Importing to Existing Database Fails
- [ATLAS-3190](#): UI : Allow user to refresh metric data inside popup
- [ATLAS-3188](#): Regression: Simple insert queries are being captured in ATLAS_HOOK topic
- [ATLAS-3187](#): Regression, Relationship updates : In GET entity definition of a deleted table, columns, sd etc., are empty
- [ATLAS-3182](#): Import Transforms: Hive Table Qualified Name is Modified Incorrectly
- [ATLAS-3181](#): The URL produced by the swagger interactive API is missing "/api/atlas/"
- [ATLAS-3176](#): Regression , Export : Export zip doesn't contain all entities which are listed in atlas-export-order.json
- [ATLAS-3174](#): UI : Exclude process or deleted entity from search drop-down if respective options are chosen
- [ATLAS-3168](#): PatchFx: Support for HA Mode
- [ATLAS-3167](#): Exception while starting Atlas with Cassandra
- [ATLAS-3160](#): AtlasRelationshipDef throws useless Atlas exceptions in constructors
- [ATLAS-3159](#): UI : Add reset button for Lineage
- [ATLAS-3158](#): UI : Render schema tab using relationshipAttributes (Columns)
- [ATLAS-3155](#): UI - IE-11 - Name entry not clickable
- [ATLAS-3154](#): UI - Lineage PNG export not working on Firefox
- [ATLAS-3151](#): incorrect return type for AtlasClientV2.getRelationshipByGuid()
- [ATLAS-3148](#): Implement Hive Metastore hook for Atlas
- [ATLAS-3147](#): Atlas server should log stats/metrics periodically in log file
- [ATLAS-3146](#): UI : Toggle N/A values in relationship and property tables.
- [ATLAS-3145](#): UI :Entity showing Active even if the relationshipStatus is Deleted
- [ATLAS-3143](#): PatchFx: Improve Memory Usage
- [ATLAS-3139](#): basic search enhancement to include relationship attributes as well in the result
- [ATLAS-3138](#): UI : Stats popup table layout improve for better readability
- [ATLAS-3137](#): UI : Create entity add element validation and design improvement
- [ATLAS-3136](#): assigning a term to an invalid entity guid throws error 500, instead of 400
- [ATLAS-3135](#): delete null data from typedefs results in messy data
- [ATLAS-3134](#): Change Date.getTime() to System.currentTimeMillis() to improve performance
- [ATLAS-3133](#): Capture metadata for different executions of the same process in Atlas
- [ATLAS-3132](#): Data Patch Fx: Improve Data Patching Performance
- [ATLAS-3131](#): UI : Stats Popup Improvements
- [ATLAS-3129](#): Truststore reloader thread from AtlasBaseClient.
- [ATLAS-3128](#): UI: Create entity is not working after new relationship attribute introduce
- [ATLAS-3127](#): UI : On Load Graph zoom buttons (+/-) is visible on top of relationship Table.
- [ATLAS-3116](#): Fix LDAP login, once a sign-on is made through Knox SSO.
- [ATLAS-3113](#): Use index query to search for active entities and better logging in java patch framework
- [ATLAS-3110](#): Add a bulk api to get entities by unique attributes
- [ATLAS-3109](#): Allow ignoring relationshipAttributes for AtlasEntityStore.getByUniqueAttributes
- [ATLAS-3107](#): update AtlasEntityDef with a field to capture attributes injected by relationship-defs
- [ATLAS-3105](#): UI : Distorted image is downloaded, for certain conditions
- [ATLAS-3104](#): Fix stale transaction alerts in atlas logs
- [ATLAS-3103](#): UI : Render Schema Table from relationship attributes node in the entity API

- [ATLAS-3102](#): Missing `setServiceType` in `AtlasBaseTypeDef` copy constructor
- [ATLAS-3101](#): UI Regression : Unable to search by `_CLASSIFIED`
- [ATLAS-3100](#): Enhance Full Text Search backend implementation to use SOLR request handler.
- [ATLAS-3099](#): Update graph storage and index backend property to be configurable
- [ATLAS-3097](#): Missing `AtlasRelationshipDef` constructors for `serviceType`
- [ATLAS-3095](#): Update `QuickStartV2` to use relationships
- [ATLAS-3093](#): UI : Add missing tooltips to lineage page.
- [ATLAS-3092](#): Atlas Plugin `ClassLoader` Doesn't Restore Thread `ClassLoader`
- [ATLAS-3091](#): `AtlasGraphProvider`: Allow for Creation of `JanusGraph` DB Connection Initialized with Batch Loading Enabled
- [ATLAS-3090](#): PC Framework: Improve Existing Producer Consumer Framework
- [ATLAS-3089](#): Refactor: Move Producer Consumer Framework to intg Project for Use Across Modules
- [ATLAS-3084](#): Adding default value to the relationship
- [ATLAS-3077](#): Handle java patches in patch framework
- [ATLAS-3076](#): UI : Relationship Tab improvements.
- [ATLAS-3075](#): Add tag propagation details in relationship notification message
- [ATLAS-3073](#): Investigate and fix IT failures in `EntityJerseyResourceIT`
- [ATLAS-3072](#): UI : Increasing the width of the Lineage slider.
- [ATLAS-3071](#): Add Functionalities to Collect Notification Metrics/Entity Lifecycle
- [ATLAS-3069](#): Avoid repeated printing of missing relationship def warning messages
- [ATLAS-3067](#): update hive types to remove use of `ownedRef/inverseRef` constraints for relationships
- [ATLAS-3066](#): UI : Various table layout improvements
- [ATLAS-3065](#): Enable relationship typedef update to remove legacy attributes
- [ATLAS-3063](#): Update entity create/update to specify relationship-type for attributes
- [ATLAS-3062](#): UI : Add relationship table inside relationships tab
- [ATLAS-3061](#): UI: Search Term is not working
- [ATLAS-3060](#): UI: Allow to drag node and place it accordingly
- [ATLAS-3059](#): UI : Lineage data structure change to support hide deleted and hide process entities to gather
- [ATLAS-3058](#): Saving of favorite-search fails
- [ATLAS-3057](#): Index Repair Tool: Add `JanusGraph`-Specific Index Repair Tool
- [ATLAS-3056](#): update `rbms` types to remove use of `ownedRef/inverseRef` constraints for relationships
- [ATLAS-3055](#): fix entity create/update to handle relationship attributes consistently
- [ATLAS-3054](#): notification processing to support batch-commits for faster processing of large messages
- [ATLAS-3053](#): Update notification processing to propagate classifications only when specified
- [ATLAS-3052](#): metrics query results caching - to be disabled
- [ATLAS-3051](#): On container entity-delete, contained entities in composition relationship should be deleted
- [ATLAS-3044](#): Notifications are not sent for entities deleted during import
- [ATLAS-3043](#): Unique attributes constraint at store-level
- [ATLAS-3041](#): Cannot delete relationship types using the `types/typedef/name` REST API
- [ATLAS-3036](#): Improve `FullTextMapper` performance during entity retrieval
- [ATLAS-3031](#): UI : Allow user to export the lineage in PNG or JPEG format
- [ATLAS-3030](#): UI : Allow to search the node in the lineage graph
- [ATLAS-3017](#): Add Atlas server statistics rest endpoint
- [ATLAS-3008](#): UI: Add local sorting for the Search table
- [ATLAS-3001](#): `AtlasClientV2.updateClassifications()` throws a serialization error
- [ATLAS-2987](#): Update component versions of Atlas to use Hadoop3, HBase2 and Solr7
- [ATLAS-2949](#): Remove duplicate attribute from the `rbms` entity
- [ATLAS-2945](#): UI - If a Term is assigned to a Category then the user is able to assign the same term again
- [ATLAS-2933](#): Empty array attributes are returned as null instead of an empty list
- [ATLAS-2885](#): Skip handling of `INSERT INTO...VALUES` hive query messages

- [ATLAS-2842](#): Entities with different GUIDs do not get referenced in array with SET cardinality
- [ATLAS-2799](#): Import Process: Additional Transforms
- [ATLAS-2458](#): RelationshipApi Delete doesn't remove relationshipAttributes from entity
- [ATLAS-2183](#): Export-API documentation curl example JSON syntax and content wrong

HDP 3.0.0, HDP 3.0.1, HDP 3.1.0 provided Atlas 1.1.0 and the following Apache patches:

- [ATLAS-2961](#): updated Hive hook to recognize configuration 'hdfs_path.convert_to_lowercase'
- [ATLAS-2950](#): UI - In newer versions of chrome such as 70 onwards, if a new column is added the instead of right most it appears at the left most
- [ATLAS-2943](#): Export options null check added.
- [ATLAS-2900](#): Export connected addressed case where imported file is not usable.
- [BUG-113089](#): fix build issue for npm mime-db compatible version issue
- [ATLAS-2927](#): Update lineage query for Process entities
- [ATLAS-2926](#): ZipSink OOM
- [ATLAS-2921](#): Update import scripts exit status in case of java error
- [ATLAS-2924](#): Use encoded property name in getEncodedProperty() and setEncodedProperty() methods
- [ATLAS-1950](#): Import API: Improvement: Specify Supertypes in Import Transforms
- [ATLAS-2797](#): Atlas Cluster.
- [ATLAS-2798](#): Export & Import Audits.
- [ATLAS-2802](#): Atlas Client Update for Export and Import.
- [ATLAS-2804](#): Export & Import Detailed Audits.
- [ATLAS-2805](#): Applying classification updates modificationTimestamp
- [ATLAS-2806](#): Using replication attributes during export and import process.
- [ATLAS-2811](#): Skip Lineage Export option.
- [ATLAS-2738](#): Export Process: Support for incremental export.
- [ATLAS-2820](#): UI : Add replication audits tab for AtlasCluster entity to show Export & Import audit
- [ATLAS-2814](#): Cluster stores replication details.
- [ATLAS-2818](#): Entity tagging after import.
- [ATLAS-2843](#): AtlasClient updates for exportData and importData.
- [ATLAS-2845](#): Rename AtlasCluster to AtlasServer.
- [ATLAS-2856](#): added utility methods to RequestContext, to find number of active requests and earliest active request-time
- [ATLAS-2862](#): Incremental Export now uses request context to determine change marker.
- [ATLAS-2870](#): Improvement to AddClassification transform to use filters.
- [ATLAS-2869](#): Hdfs_path if requested are created and then proceeds with export.
- [ATLAS-2864](#): Improved incremental export queries.
- [ATLAS-2873](#): Atlas Import Transform Handler Implementation
- [ATLAS-2874](#): Include handling of Atlas Entity Transformers in current Import logic
- [ATLAS-2875](#): Implement clear attribute value transformer for Atlas Entity Transformer
- [ATLAS-2882](#): AddClassification transform for new transforms
- [ATLAS-2886](#): Support for fully qualified server name
- [ATLAS-2892](#): Delete by name REST endpoint.
- [ATLAS-2895](#): Server full name processing
- [ATLAS-2888](#): Change marker fix for server name. Unit test fix.
- [ATLAS-2906](#): Allow transforms to be applied when entity-level transforms are present.
- [ATLAS-2897](#): Better handling of empty zip files. Unit test fix.
- [ATLAS-2909](#): ChangeMarker updated during initialization.
- [ATLAS-2923](#): Options copied to AttributeDefinition
- [ATLAS-2922](#): Multiplicity computation updated for SET/optional
- [ATLAS-2724](#): UI enhancement for Avro schemas and other JSON-valued attributes

- [ATLAS-2786](#): updated HBase client to use all zookeeper entries
- [ATLAS-2827](#): Fix to handle failure in saving indexable string property of large size
- [ATLAS-2844](#): added eventTime field in EntityNotificationV2
- [ATLAS-2847](#): UI - Add grunt scripts in package.json
- [ATLAS-2848](#): Fix property name in ApplicationProperties
- [ATLAS-2849](#): Basic search result list for classification= '_NOT_CLASSIFIED' returns type vertices in addition to matching entity vertices
- [ATLAS-2852](#): Export Import simultaneous requests updated error code.
- [ATLAS-2853](#): updated to send entity-notifications after successful graph transaction commit
- [ATLAS-2857](#): Fix Atlas release version in /api/atlas/version API.
- [ATLAS-2860](#): Remove dependency on atlas-application.properties for AtlasClient.
- [ATLAS-2872](#): updated HiveHook to purge name cache periodically
- [ATLAS-2877](#): updated notification processing to wait only before retry
- [ATLAS-2878](#): avoid retrieval of entiyWithExtInfo when extInfo is not needed
- [ATLAS-2891](#): updated hook notification processing with option to ignore potentially incorrect hive_column_lineage
- [ATLAS-2905](#): Generate lineage information for process entities
- [ATLAS-2907](#): UI - show lineage information for process entities
- [ATLAS-2917](#): Filter added for Lineage to hide Process entity.
- [ATLAS-2918](#): Update JanusGraph version to 0.3.1
- [ATLAS-2920](#): Update JanusGraph Solr clients to use all zookeeper entries

HDP 3.0.0 and HDP 3.0.1 provided Atlas 1.0.0 and the following Apache patches:

- [ATLAS-2740](#): Updated import-hbase.sh script to populate values for missed HBase column-family attributes.
- [ATLAS-2751](#): Atlas is not consuming messages from ATLAS_HOOK topic after recovering from zookeeper connection timeout.

Calcite

This release provides Calcite 1.16.0 and no additional Apache patches.

HDP 3.0.0, HDP 3.0.1, and HDP 3.1.0 provided Calcite 1.16.0 and no additional Apache patches.

DataFu

This release provides DataFu 1.3.0 and no additional Apache patches.

HDP 3.0.0, HDP 3.0.1, and HDP 3.1.0 provided DataFu 1.3.0 and no additional Apache patches.

Hadoop

This release provides Hadoop Common 3.1.1 and no additional Apache patches.

HDP 3.1.0 provided Hadoop Common 3.1.1 and no additional Apache patches.

HDP 3.0.1 provided Hadoop Common 3.1.1 and no additional Apache patches.

HDP 3.0.0 provided Hadoop Common 3.1.0 and the following Apache patches:

- [HADOOP-15571](#): Multiple FileContexts created with the same configuration object should be allowed to have different umask. Contributed by Vinod Kumar Vavilapalli.
- [HDFS-13712](#): BlockReaderRemote.read() logging improvement. Contributed by Gergo Repas.
- [HDFS-13715](#): diskbalancer does not work if one of the blockpools are empty on a Federated cluster. Contributed by Bharat Viswanadham

- [HDFS-13715](#): Fix TestNameNodePrunesMissingStorages test failures on Windows. Contributed by Lukas Majercak.
- [HDFS-13635](#): Incorrect message when block is not found. Contributed by Gabor Bota.
- [HADOOP-15554](#): Improve JIT performance for Configuration parsing. Contributed by Todd Lipcon.
- [HDFS-13702](#): Remove HTrace hooks from DFSClient to reduce CPU usage. Contributed by Todd Lipcon.
- [HDFS-13703](#): Avoid allocation of CorruptedBlocks hashmap when no corrupted blocks are hit. Contributed by Todd Lipcon.
- [HADOOP-15574](#): Suppress build error if there are no docs after excluding private annotations. Contributed by Takanobu Asanuma.
- [HADOOP-15548](#): Randomize local dirs. Contributed by Jim Brennan.
- [HDFS-13705](#): The native ISA-L library loading failure should be made warning rather than an error message. Contributed by Shashikant Banerjee.
- [HADOOP-15458](#): TestLocalFileSystem#testFSOutputStreamBuilder fails on Windows. Contributed by Xiao Liang.
- [HDFS-13692](#): StorageInfoDefragmenter floods log when compacting StorageInfo TreeSet. Contributed by Bharat Viswanadham.
- [HDFS-13682](#): Cannot create encryption zone after KMS auth token expires.
- [HADOOP-15527](#): Improve delay check for stopping processes. Fixed script location. Contributed by Vinod Kumar Vavilapalli
- [HADOOP-15527](#): Improve delay check for stopping processes. Contributed by Vinod Kumar Vavilapalli
- [HADOOP-15533](#): Make WASB listStatus messages consistent. Contributed by Esfandiar Manii
- [HADOOP-15523](#): Shell command timeout given is in seconds whereas it is taken as millisec while scheduling. Contributed by Bilwa S T.
- [HDFS-13681](#): Fix TestStartup.testNNFailToStartOnReadOnlyNNDir test failure on Windows. Contributed by Xiao Liang.
- [HDFS-13686](#): Add overall metrics for FSNamesystemLock. Contributed by Lukas Majercak.
- [HDFS-13174](#): hdfs mover -p /path times out after 20 min. Contributed by Istvan Fajth.
- [HDFS-13676](#): TestEditLogRace fails on Windows. Contributed by Zuoming Zhang.
- [HDFS-13673](#): TestNameNodeMetrics fails on Windows. Contributed by Zuoming Zhang.
- [HDFS-13675](#): Speed up TestDFSAdminWithHA. Contributed by Lukas Majercak.
- [HDFS-13563](#): TestDFSAdminWithHA times out on Windows. Contributed by Lukas Majercak.
- [HDFS-13641](#): Add metrics for edit log tailing. Contributed by Chao Sun.
- [BUG-103209](#): Dependency upgrades related to blackduck scan.
- [HADOOP-15527](#): Improve delay check for stopping processes. Contributed by Vinod Kumar Vavilapalli
- [HADOOP-15532](#): TestBasicDiskValidator fails with NoSuchFileException. Contributed by Giovanni Matteo Fumarola.
- [HDFS-13511](#): Provide specialized exception when block length cannot be obtained. Contributed by Gabor Bota.
- [HADOOP-15529](#): ContainerLaunch#testInvalidEnvVariableSubstitutionType is not supported in Windows. Contributed by Giovanni Matteo Fumarola.
- [HADOOP-15506](#): Upgrade Azure Storage Sdk version to 7.0.0 and update corresponding code blocks. Contributed by Esfandiar Manii.
- [HDFS-13653](#): Make dfs.client.failover.random.order a per nameservice configuration. Contributed by Ekanth Sethuramalingam.
- [HADOOP-15499](#): Performance severe drops when running RawErasureCoderBenchmark with NativeRSRawErasureCoder. Contributed by Sammi Chen.
- [HDFS-13667](#): Typo: Marking all datadoes as stale. Contributed by Nanda Kumar
- [HDFS-12670](#): can't renew HDFS tokens with only the hdfs client jar. Contributed by Arpit Agarwal.
- [HDFS-13664](#): Refactor ConfiguredFailoverProxyProvider to make inheritance easier. Contributed by Chao Sun.
- [HADOOP-15516](#): Add test cases to cover FileUtil#readLink. Contributed by Giovanni Matteo Fumarola.
- [HDFS-13659](#): Add more test coverage for contentSummary for snapshottable path. Contributed by Wei-Chiu Chuang.
- [HDFS-12950](#): [oiv] ls will fail in secure cluster. Contributed by Wei-Chiu Chuang.

- [HADOOP-15217](#): FsURLConnection does not handle paths with spaces. Contributed by Joseph Fourny and Zsolt Venczel.
- [HDFS-13545](#): "guarded" is misspelled as "gaurded" in FSPermissionChecker.java. Contributed by Jianchao Jia.
- [HDFS-13650](#): Randomize baseDir for MiniDFSCluster in TestDFSStripedInputStream and TestDFSStripedInputStreamWithRandomECPolicy. Contributed by Anbang Hu.
- [HDFS-13649](#): Randomize baseDir for MiniDFSCluster in TestReconstructStripedFile and TestReconstructStripedFileWithRandomECPolicy. Contributed by Anbang Hu.
- [HDFS-13652](#): Randomize baseDir for MiniDFSCluster in TestBlockScanner. Contributed by Anbang Hu.
- [HADOOP-15137](#): ClassNotFoundException: org.apache.hadoop.yarn.server.api.DistributedSchedulingAMProtocol when using hadoop-client-minicluster. Contributed by Bharat Viswanadham
- [HDFS-13281](#) Namenode#createFile should be /.reserved/raw/ aware.. Contributed by Rushabh S Shah
- [HDFS-13155](#): BlockPlacementPolicyDefault.chooseTargetInOrder Not Checking Return Value for NULL. Contributed by Zsolt Venczel.
- [HDFS-13339](#): Volume reference can't be released and may lead to deadlock when DataXceiver does a check volume. Contributed by liaoyuxiangqin and Zsolt Venczel.
- [HDFS-13648](#): Fix TestGetConf#testGetJournalNodes on Windows due to a mismatch line separator. Contributed by Giovanni Matteo Fumarola.
- [HDFS-13651](#): TestReencryptionHandler fails on Windows. Contributed by Anbang Hu.
- MapReduce JHS should honor yarn.webapp.filter-entity-list-by-user(addendum). Contributed by Sunil Govindan.
- [HDFS-13637](#): RBF: Router fails when threadIndex (in ConnectionPool) wraps around Integer.MIN_VALUE. Contributed by CR Hota.
- [HDFS-13636](#): Cross-Site Scripting vulnerability in HttpServer2 (Contributed by Haibo Yan via Daniel Templeton)
- [HDFS-13582](#): Improve backward compatibility for HDFS-13176 (WebHdfs file path gets truncated when having semicolon (;) inside). Contributed by Zsolt Venczel.
- [HDFS-13602](#): Add checkOperation(WRITE) checks in FSNamesystem. Contributed by Chao Sun.
- [HDFS-12978](#): Fine-grained locking while consuming journal stream. Contributed by Konstantin Shvachko. (cherry picked from commit 442dd87dcd3e9913a1d97793fde879ce51d701e7)
- [HDFS-13631](#): TestDFSAdmin#testCheckNumOfBlocksInReportCommand should use a separate MiniDFSCluster path. Contributed by Anbang Hu.
- [BUG-104243](#):Dependency jackson-databind dependency related to blackduck scan.
- [HADOOP-15406](#): hadoop-nfs dependencies for mockito and junit are not test scope
- [HDFS-13611](#): Unsafe use of Text as a ConcurrentHashMap key in PBHelperClient.
- [HDFS-13629](#): Some tests in TestDiskBalancerCommand fail on Windows due to MiniDFSCluster path conflict and improper path usage. Contributed by Anbang Hu.
- [HDFS-13632](#): Randomize baseDir for MiniJournalCluster in MiniQJMHACluster for TestDFSAdminWithHA. Contributed by Anbang Hu.
- Additional check when unpacking archives. Contributed by Wilfred Spiegelenburg.
- [HADOOP-15497](#): TestTrash should use proper test path to avoid failing on Windows. Contributed by Anbang Hu.
- [HADOOP-15498](#): TestHadoopArchiveLogs (#testGenerateScript, #testPrepareWorkingDir) fails on Windows. Contributed by Anbang Hu.
- [HDFS-13591](#): TestDFSShell#testSetrepLow fails on Windows. Contributed by Anbang Hu.
- [HDFS-13627](#): TestErasureCodingExerciseAPIs fails on Windows. Contributed by Anbang Hu.
- [HADOOP-15449](#): Increase default timeout of ZK session to avoid frequent NameNode failover
- [HDFS-13628](#): Update Archival Storage doc for Provided Storage
- [HDFS-13620](#): Randomize the test directory path for TestHDFSFileSystemContract. Contributed by Anbang Hu.
- [HDFS-13619](#): TestAuditLoggerWithCommands fails on Windows. Contributed by Anbang Hu.
- [HADOOP-15473](#): Configure serialFilter in KeyProvider to avoid UnrecoverableKeyException caused by JDK-8189997. Contributed by Gabor Bota.
- [HDFS-13618](#): Fix TestDataNodeFaultInjector test failures on Windows. Contributed by Xiao Liang.
- [HADOOP-15486](#): Make NetworkTopology#netLock fair. Contributed by Nanda kumar.

- Skip the proxy user check if the ugi has not been initialized. Contributed by Daryn Sharp
- [HDFS-13540](#): DFSStripedInputStream should only allocate new buffers when reading. Contributed by Xiao Chen.
- [HADOOP-15478](#): WASB: hflush() and hsync() regression. Contributed by Thomas Marquardt.
- [HDFS-13493](#): Reduce the HttpServer2 thread count on DataNodes. Contributed by Erik Krogen.
- Additional check when unpacking archives. Contributed by Jason Lowe and Akira Ajisaka.
- [HDFS-13587](#): TestQuorumJournalManager fails on Windows. Contributed by Anbang Hu.
- [HDFS-13588](#): Fix TestFsDatasetImpl test failures on Windows. Contributed by Xiao Liang.
- [HADOOP-15450](#): Avoid fsync storm triggered by DiskChecker and handle disk full situation. Contributed by Arpit Agarwal.
- [HDFS-13601](#): Optimize ByteString conversions in PBHelper.
- [HDFS-13589](#): Add dfsAdmin command to query if upgrade is finalized. Contributed by Hanisha Koneru
- [HADOOP-15299](#): Bump Jackson 2 version to Jackson 2.9.x.
- [HDFS-13593](#): TestBlockReaderLocalLegacy#testBlockReaderLocalLegacyWithAppend fails on Windows. Contributed by Anbang Hu.
- [HDFS-13592](#): TestNameNodePrunesMissingStorages#testNameNodePrunesUnreportedStorages does not shut down cluster properly. Contributed by Anbang Hu.
- [HDFS-13586](#): Fsync fails on directories on Windows. Contributed by Lukas Majercack.
- [HDFS-13560](#): Insufficient system resources exist to complete the requested service for some tests on Windows. Contributed by Anbang Hu.
- [HDFS-13556](#): TestNestedEncryptionZones does not shut down cluster. Contributed by Anbang Hu.
- [HDFS-13554](#): TestDatanodeRegistration#testForcedRegistration does not shut down cluster. Contributed by Anbang Hu.
- [HDFS-13558](#): TestDatanodeHttpXFrame does not shut down cluster. Contributed by Anbang Hu.
- [BUG-103397](#). [HDFS-13589](#). Add dfsAdmin command to query if upgrade is finalized.
- [HADOOP-15469](#): S3A directory committer commit job fails if _temporary directory created under dest. Contributed by Steve Loughran.
- [HADOOP-15250](#): Split-DNS MultiHomed Server Network Cluster Network IPC Client Bind Addr Wrong Contributed by Ajay Kumar
- [HDFS-13570](#): TestQuotaByStorageType,TestQuota,TestDFSOutputStream fail on Windows. Contributed by Anbang Hu.
- [HDFS-13581](#): DN UI logs link is broken when https is enabled. Contributed by Shashikant Banerjee.
- [HDFS-13559](#): TestBlockScanner does not close TestContext properly. Contributed by Anbang Hu.
- [HDFS-13512](#): WebHdfs getFileStatus doesn't return ecPolicy. Contributed by Ajay Kumar.
- [HDFS-13550](#):TestDebugAdmin#testComputeMetaCommand fails on Windows. Contributed by Anbang Hu.
- [HDFS-13557](#): TestDFSAdmin#testListOpenFiles fails on Windows. Contributed by Anbang Hu.
- [BUG-103055](#) Remove aiyun, openstack from hadoop-cloud-project
- [HDFS-13567](#): TestNameNodeMetrics#testGenerateEDEKTime,TestNameNodeMetrics#testResourceCheck should use a different cluster basedir. Contributed by Anbang Hu.
- [HDFS-13548](#): TestResolveHdfsSymlink#testFcResolveAfs fails on Windows. Contributed by Anbang Hu.
- [HDFS-11700](#). TestHDFSServerPorts#testBackupNodePorts doesn't pass on Windows. Contributed by Anbang Hu.
- [HDFS-13551](#): TestMiniDFSCluster#testClusterSetStorageCapacity does not shut down cluster. Contributed by Anbang Hu.
- [HADOOP-15442](#): ITestS3AMetrics.testMetricsRegister can't know metrics source's name. Contributed by Sean Mackrory.
- [HDFS-13544](#): Improve logging for JournalNode in federated cluster.
- [HDFS-13539](#) DFSStripedInputStream NPE when reportCheckSumFailure.
- Add 2.9.1 release notes and changes documents
- [BUG-102169](#) Include Google Connector in Hadoop
- [HADOOP-15441](#): Log kms url and token service at debug level. Contributed by Gabor Bota

- [HDFS-13542](#): TestBlockManager#testNeededReplicationWhileAppending fails due to improper cluster shutdown in TestBlockManager#testBlockManagerMachinesArray on Windows. Contributed by Anbang Hu.
- [HDFS-13346](#): RBF: Fix synchronization of router quota and nameservice quota.
- [HADOOP-15454](#): TestRollingFileSystemSinkWithLocal fails on Windows. Contributed by Xiao Liang.
- [HDFS-13537](#): TestHdfsHelper does not generate jceks path properly for relative path in Windows. Contributed by Xiao Liang.
- [HDFS-12981](#): renameSnapshot a Non-Existent snapshot to itself should throw error. Contributed by Kittu Nanasi.
- [HADOOP-15446](#): WASB: PageBlobInputStream.skip breaks HBASE replication. Contributed by Thomas Marquardt
- [HADOOP-14841](#) Kms client should disconnect if unable to get output stream from connection. Contributed by Rushabh S Shah
- [HDFS-13525](#): RBF: Add unit test TestStateStoreDisabledNameservice. Contributed by Yiqun Lin.
- [HDFS-13481](#): TestRollingFileSystemSinkWithHdfs#testFlushThread: test failed intermittently (Contributed by Gabor Bota via Daniel Templeton)
- [HDFS-13512](#) . WebHdfs getFileStatus doesn't return ecPolicy. Contributed by Ajay Kumar.
- [HDFS-13488](#): RBF: Reject requests when a Router is overloaded. Contributed by Inigo Goiri.
- [HDFS-13503](#): Fix TestFscck test failures on Windows. Contributed by Xiao Liang.
- [HADOOP-15390](#): Yarn RM logs flooded by DelegationTokenRenewer trying to renew KMS tokens (xiaochen via rkanter)
- [HDFS-13283](#): Percentage based Reserved Space Calculation for DataNode. Contributed by Lukas Majercak.
- [HDFS-13509](#): Bug fix for breakHardlinks() of ReplicaInfo/LocalReplica, and fix TestFileAppend failures on Windows. Contributed by Xiao Liang.
- [HDFS-13508](#): RBF: Normalize paths (automatically) when adding, updating, removing or listing mount table entries. Contributed by Ekanth S.
- [HDFS-13484](#): RBF: Disable Nameservices from the federation. Contributed by Inigo Goiri.
- [HDFS-13434](#): RBF: Fix dead links in RBF document. Contributed by Chetna Chaudhari.
- [HDFS-13490](#): RBF: Fix setSafeMode in the Router. Contributed by Inigo Goiri.
- [HDFS-13499](#): RBF: Show disabled name services in the UI. Contributed by Inigo Goiri.
- [HDFS-13326](#): RBF: Improve the interfaces to modify and view mount tables. Contributed by Gang Li.
- [HADOOP-15411](#): AuthenticationFilter should use Configuration.getPropsWithPrefix instead of iterator. (Suma Shivaprasad via wangda)
- [HDFS-13336](#): Test cases of TestWriteToReplica failed in windows. Contributed by Xiao Liang.
- [HDFS-13356](#): Balancer:Set default value of minBlockSize to 10mb. Contributed by Bharat Viswanadham
- [BUG-101422](#). Disable HDFS and Hadoop unit tests that are flaky in Mawo.
- [HDFS-13408](#): MiniDFSCluster to support being built on randomized base directory. Contributed by Xiao Liang
- [HDFS-13433](#): webhdfs requests can be routed incorrectly in federated cluster. Contributed by Arpit Agarwal.
- [HDFS-13478](#): RBF: Disabled Nameservice store API. Contributed by Inigo Goiri.
- [HDFS-10183](#): Prevent race condition during class initialization. Contributed by Pavel Avgustinov.
- [HDFS-13453](#): RBF: getMountPointDates should fetch latest subdir time/date when parent dir is not present but / parent/child dirs are present in mount table. Contributed by Dibyendu Karmakar.
- [HADOOP-15396](#): Some java source files are executable
- [HDFS-13464](#): Fix javadoc in FsVolumeList#handleVolumeFailures. Contributed by Shashikant Banerjee
- [HADOOP-15366](#): Add a helper shutdown routine in HadoopExecutor to ensure clean shutdown. Contributed by Shashikant Banerjee.
- [HDFS-13426](#): Fix javadoc in FsDatasetAsyncDiskService#removeVolume. Contributed by Shashikant Banerjee.
- [HDFS-13463](#): Fix javadoc in FsDatasetImpl#checkAndUpdate. Contributed by Shashikant Banerjee.
- [HADOOP-15391](#): Add missing css file in hadoop-aws, hadoop-aliyun, hadoop-azure and hadoop-azure-datalake modules.
- [HDFS-12828](#): OIV ReverseXML Processor fails with escaped characters
- [BUG-100895](#). Ignore UTs failing due to environmental issues.
- [HDFS-13439](#): Add test case for read block operation when it is moved. Contributed by Ajay Kumar.

- [HDFS-13435](#): RBF: Improve the error loggings for printing the stack trace.
- [HDFS-13330](#): ShortCircuitCache#fetchOrCreate never retries. Contributed by Gabor Bota.
- [HDFS-13311](#): RBF: TestRouterAdminCLI#testCreateInvalidEntry fails on Windows. Contributed by Inigo Goiri
- [HDFS-13056](#): Expose file-level composite CRCs in HDFS which are comparable across different instances/layouts. Contributed by Dennis Huo.
- [HDFS-13466](#): RBF: Add more router-related information to the UI.
- [HDFS-13462](#): Add BIND_HOST configuration for JournalNode's HTTP and RPC Servers. Contributed by Lukas Majercak.
- [HDFS-13129](#): Add a test for DfsAdmin refreshSuperUserGroupsConfiguration. Contributed by Mukul Kumar Singh
- [HDFS-13433](#): webhdfs requests can be routed incorrectly in federated cluster. Contributed by Arpit Agarwal.
- [HDFS-11900](#): Hedged reads thread pool creation not synchronized. Contributed by John Zhuge.
- [HDFS-11043](#): TestWebHdfsTimeouts fails. Contributed by Xiaoyu Yao and Chao Sun.
- [HDFS-13315](#): Add a test for the issue reported in HDFS-11481 which is fixed by HDFS-10997. Contributed by Yongjun Zhang.
- [HADOOP-15332](#): Fix typos in hadoop-aws markdown docs. Contributed by Gabor Bota.
- [HADOOP-15331](#): Fix a race condition causing parsing error of java.io.BufferedInputStream in class org.apache.hadoop.conf.Configuration. Contributed by Miklos Szegedi.
- [HADOOP-15330](#) Remove jdk1.7 profile from hadoop-annotations module
- [HADOOP-15062](#): TestCryptoStreamsWithOpensslAesCtrCryptoCodec fails on Debian 9. Contributed by Miklos Szegedi.
- [HADOOP-14667](#): Flexible Visual Studio support. Contributed by Allen Wittenauer
- [HDFS-13438](#): Fix javadoc in FsVolumeList#removeVolume. Contributed by Shashikant Banerjee.
- [HDFS-13418](#): NetworkTopology should be configurable when enable DFSNetworkTopology. Contributed by Tao Jie.
- [HADOOP-14970](#): MiniHadoopClusterManager doesn't respect lack of format option. Contributed by Erik Krogen
- [HADOOP-15379](#): Make IrqHandler.bind() public. Contributed by Ajay Kumar
- [HDFS-13436](#): Fix javadoc of package-info.java
- [HDFS-7101](#): Potential null dereference in DFSck#doWork(). Contributed by skrho and Ted Yu.
- [HDFS-13427](#): Fix the section titles of transparent encryption document.
- [HDFS-13386](#): RBF: Wrong date information in list file(-ls) result. Contributed by Dibyendu Karmakar.
- [HDFS-13428](#): RBF: Remove LinkedList From StateStoreFileImpl.java. Contributed by BELUGA BEHR.
- [HDFS-13045](#): RBF: Improve error message returned from subcluster. Contributed by Inigo Goiri.
- [HADOOP-15346](#) S3ARetryPolicy for 400/BadArgument to be "fail". Contributed by Steve Loughran.
- [HADOOP-15313](#): TestKMS should close providers.
- [HADOOP-15357](#): Configuration.getPropsWithPrefix no longer does variable substitution. Contributed by Jim Brennan
- [HDFS-13328](#): Abstract ReencryptionHandler recursive logic in separate class. Contributed by Surendra Singh Lilhore.
- [HADOOP-15376](#): Remove double semi colons on imports that make Clover fall over.
- [HDFS-13420](#): License header is displayed in ArchivalStorage/MemoryStorage html pages. Contributed by Akira Ajisaka.
- [HADOOP-15374](#): Add links of the new features of 3.1.0 to the top page
- [HDFS-13384](#): RBF: Improve timeout RPC call mechanism. Contributed by Inigo Goiri.
- [HDFS-13410](#): RBF: Support federation with no subclusters. Contributed by Inigo Goiri.
- [HDFS-13380](#): RBF: mv/rm fail after the directory exceeded the quota limit. Contributed by Yiqun Lin.
- [HDFS-13388](#): RequestHedgingProxyProvider calls multiple configured NNs all the time. Contributed by Jinglun.
- [HDFS-13402](#): RBF: Fix java doc for StateStoreFileSystemImpl. Contributed by Yiran Wu.
- Added CHANGES/RELEASES/Jdiff for 3.1.0 release
- [HDFS-13353](#): RBF: TestRouterWebHDFSContractCreate failed. Contributed by Takanobu Asanuma.

- [HDFS-13350](#): Negative legacy block ID will confuse Erasure Coding to be considered as striped block. (Contributed by Lei (Eddy) Xu).
- [HDFS-13364](#): RBF: Support NamenodeProtocol in the Router. Contributed by Inigo Goiri.
- [HADOOP-15355](#): TestCommonConfigurationFields is broken by HADOOP-15312. Contributed by LiXin Ge.
- [HADOOP-15317](#): Improve NetworkTopology chooseRandom's loop.
- [HDFS-13289](#): RBF: TestConnectionManager#testCleanup() test case need correction. Contributed by Dibyendu Karmakar.
- [HADOOP-15253](#): Should update maxQueueSize when refresh call queue. Contributed by Tao Jie.
- [HADOOP-14999](#): AliyunOSS: provide one asynchronous multi-part based uploading mechanism. Contributed by Genmao Yu.
- [HADOOP-12862](#): LDAP Group Mapping over SSL can not specify trust store. Contributed by Wei-Chiu Chuang and Konstantin Shvachko.
- [HDFS-13087](#): Snapshotted encryption zone information should be immutable. Contributed by LiXin Ge.
- [HADOOP-15342](#): Updating ADLS connector to use the current SDK version (2.2.7). Contributed by Atul Sikaria.

HBase

This release provides HBase 2.0.2 and the listed Apache patches.

- [HBASE-22690](#): Deprecate / Remove OfflineMetaRepair in hbase-2+
- [HBASE-19893](#): restore_snapshot is broken in master branch when region splits
- [HBASE-21136](#): NPE in MetricsTableSourceImpl.updateFlushTime
- [HBASE-21284](#): Default limits for PressureAwareCompactionThroughputController are too low
- [HBASE-17525](#): Add space quota support to the deprecated SecureBulkLoadEndpoint
- [HBASE-15728](#): Add Per-Table metrics back
- [HBASE-21890](#): Use execute instead of submit to submit a task in RemoteProcedureDispatcher
- [HBASE-20662](#): Increasing space quota on a violated table does not remove SpaceViolationPolicy.DISABLE enforcement
- [HBASE-21634](#): Print error message when user uses unacceptable values for LIMIT while setting quotas.
- [HBASE-21225](#): Having RPC quota on a table/namespace doesn't allow space quota to be removed using 'NONE'
- [HBASE-21631](#): list_quotas should print human readable values for LIMIT
- [HBASE-21988](#): Remove hadoop2.6.1-hadoop-2.6.4 as supported on branch-2.0
- [HBASE-21915](#): Make FileLinkInputStream implement CanUnbuffer
- [HBASE-22200](#): WALSplitter.hasRecoveredEdits should use same FS instance from WAL region dir
- [HBASE-21688](#): Address WAL filesystem issues
- [HBASE-22144](#): Correct MultiRowRangeFilter to work with reverse scans
- [HBASE-21734](#): Some optimization in FilterListWithOR
- [HBASE-21620](#): Problem in scan query when using more than one column prefix filter in some cases
- [HBASE-21129](#): Clean up duplicate codes in #equals and #hashCode methods of Filter
- [HBASE-19008](#): Add missing equals or hashCode method(s) to stock Filter implementations
- [HBASE-22185](#): RAMQueueEntry#writeToCache should freeBlock if any exception encountered instead of the IOException catch block
- [HBASE-21932](#): Use Runtime.getRuntime().halt to terminate regionserver when abort timeout
- [HBASE-21775](#): The BufferedMutator doesn't ever refresh region location cache
- [HBASE-21644](#): Modify table procedure runs infinitely for a table having region replication > 1
- [HBASE-21621](#): Reversed scan does not return expected number of rows
- [HBASE-21618](#): Scan with the same startRow(inclusive=true) and stopRow(inclusive=false) returns one result
- [HBASE-21402](#): Backport parent "HBASE-21325 Force to terminate regionserver when abort hang in somewhere"
- [HBASE-22123](#): REST gateway reports Insufficient permissions exceptions as 404 Not Found
- [HBASE-21732](#): Should call toUpperCase before using Enum.valueOf in some methods for ColumnFamilyDescriptor

- [HBASE-21683](#): Reset readsEnabled flag after successfully flushing the primary region
- [HBASE-21645](#): Perform sanity check and disallow table creation/modification with region replication < 1
- [HBASE-21413](#): Empty meta log doesn't get split when restart whole cluster
- [HBASE-20220](#): [RSGroup] Check if table exists in the cluster before moving it to the specified regionserver group
- [HBASE-18484](#): VerifyRep by snapshot does not work when Yarn/SourceHBase/PeerHBase located in three different HDFS clusters
- [HBASE-21738](#): Remove all the CSLM#size operation in our memstore because it's an quite time consuming.
- [HBASE-22128](#): Move namespace region then master crashed make deadlock
- [HBASE-21551](#): Memory leak when use scan with STREAM at server side
- [HBASE-20604](#): ProtobufLogReader#readNext can incorrectly loop to the same position in the stream until the the WAL is rolled
- [HBASE-21960](#): Ensure RESTServletContainer used by RESTServer
- [HBASE-22077](#): Expose sleep time as a command line argument of IntergationTestBackupRestore
- [HBASE-20734](#): Colocate recovered edits directory with hbase.wal.dir
- [HBASE-21568](#): Disable use of BlockCache for LoadIncrementalHFiles
- [HBASE-21523](#): Avoid extra logging when the backup system table already exists
- [HBASE-21524](#): Fix logging in ConnectionImplementation.isTableAvailable()
- [HBASE-21466](#): WALProcedureStore uses wrong FileSystem if wal.dir is not under rootdir
- [HBASE-21457](#): BackupUtils#getWALFilesOlderThan refers to wrong FileSystem

HDP 3.1.0 provided HBase 2.0.2 and the following Apache patches:

- [HBASE-20723](#): Custom hbase.wal.dir results in data loss because we write recovered edits into a different place than where the recovering region server looks for them
- [HBASE-20722](#): Make RegionServerTracker only depend on children changed event
- [HBASE-20668](#): Avoid permission change if ExportSnapshot's copy fails
- [HBASE-20657](#): Retrying RPC call for ModifyTableProcedure may get stuck
- [HBASE-20635](#): Add missing shaded-pb conversion method for UserPermission
- [HBASE-20627](#): Relocate RS Group pre/post hooks from RSGroupAdminServer to RSGroupAdminEndpoint
- [HBASE-20579](#): Include original exception in wrapped exception
- [HBASE-20579](#): Improve snapshot manifest copy in ExportSnapshot
- [HBASE-20547](#): Restore from backup will fail if done from a different file system
- [HBASE-20531](#): RS may throw NPE when close meta regions in shutdown procedure.
- [HBASE-20495](#): REST unit test fails with NoClassDefFoundError against hadoop3
- [HBASE-20196](#): Maintain all regions with same size in memstore flusher
- [HBASE-20136](#): TestKeyValue misses ClassRule and Category annotations
- [HBASE-19735](#): Create a client-tarball assembly
- [HBASE-19478](#): Utilize multi-get to speed up WAL file checking in BackupLogCleaner (Toshihiro Suzuki)
- [HBASE-18135](#): Implement mechanism for RegionServers to report file archival for space quotas
- [HBASE-18133](#): Decrease quota reaction latency by HBase
- [HBASE-17549](#): HBase-Spark Module: Corrected - Incorrect log at println and unwanted comment code
- [HBASE-21440](#): Assign procedure on the crashed server is not properly interrupted
- [HBASE-21344](#): Fix TestMetaShutdownHandler (addendum)
- [HBASE-21384](#): Procedure with holdlock=false should not be restored lock when restarts
- [HBASE-21344](#): hbase:meta location in ZooKeeper set to OPENING by the procedure which eventually failed but precludes Master from assigning it forever
- [HBASE-21215](#): Figure how to invoke hbck2; make it easy to find
- [HBASE-21372](#): Set hbase.assignment.maximum.attempts to Long.MAX
- [HBASE-21363](#): Rewrite the buildingHoldCleanupTracker method in WALProcedureStore
- [HBASE-21364](#): Procedure holds the lock should put to front of the queue after restart
- [HBASE-21342](#): FileSystem in use may get closed by other bulk load call in secure bulkLoad

- [HBASE-21349](#): Do not run CatalogJanitor or Nomalizer when cluster is shutting down
- [HBASE-21338](#): Warn if balancer is an ill-fit for cluster size
- [HBASE-21073](#): Redo concept of maintenance mode
- [HBASE-20937](#): ArrayIndexOutOfBoundsException when rolling back procedure
- [HBASE-21321](#): Backport HBASE-21278 to branch-2.1 and branch-2.0 ("Do not rollback successful sub procedures when rolling back a procedure")
- [HBASE-21354](#): Procedure may be deleted improperly during master restarts resulting in 'Corrupt'(addendum)
- [HBASE-21356](#): bulkLoadHFile API should ensure that rs has the source hfile's write permissions
- [HBASE-21336](#): Addendum remove unused code in HBTU
- [HBASE-21355](#): (addendum) replace the expensive reload storefiles with reading the merge result of compacted storefiles and current storefiles
- [HBASE-21336](#): Simplify the implementation of WALProcedureMap
- [HBASE-21355](#): HStore's storeSize is calculated repeatedly which causing the confusing region split
- [HBASE-21348](#): Fix failing TestRegionBypass, broke by HBASE-21291
- [HBASE-21345](#): [hbck2] Allow version check to proceed even though master is 'initializing'.
- [HBASE-21200](#): Memstore flush doesn't finish because of seekToPreviousRow() in memstore scanner.
- [HBASE-21323](#): Should not skip force updating for a sub procedure even if it has been finished
- [HBASE-21075](#): Confirm that we can (rolling) upgrade from 2.0.x and 2.1.x to 2.2.x after HBASE-20881
- [HBASE-21292](#): IdLock.getLockEntry() may hang if interrupted
- [HBASE-21291](#): Add a test for bypassing stuck state-machine procedures
- [HBASE-21323](#): Should not skip force updating for a sub procedure even if it has been finished; ADDENDUM
- [HBASE-21288](#): HostingServer in UnassignProcedure is not accurate
- [HBASE-21323](#): Should not skip force updating for a sub procedure even if it has been finished
- [HBASE-21055](#): NullPointerException when balanceOverall() but server balance info is null
- [HBASE-21198](#): Exclude dependency on net.minidev:json-smart
- [HBASE-21310](#) and [HBASE-21311](#): Addendum fix failed UTs, some UTs are not present on branch-2.1 and some are a bit different in the implementation
- [HBASE-21327](#): Fix minor logging issue where we don't report servername if no associated SCP
- [HBASE-12320](#): [canary] Cleanup of usage and add commentary
- [HBASE-21311](#): Split TestRestoreSnapshotFromClient
- [HBASE-21263](#): Mention compression algorithm along with other storefile details
- [HBASE-21242](#): [amv2] Miscellaneous minor log and assign procedure create improvements; ADDENDUM Fix TestHRegionInfo AND TestRegionInfoDisplay
- [HBASE-21315](#): The getActiveMinProcId and getActiveMaxProcId of BitSetNode are incorrect if there are no active procedure
- [HBASE-21310](#): Split TestCloneSnapshotFromClient
- [HBASE-21266](#): Not running balancer because processing dead regionservers, but empty dead rs list
- [HBASE-21290](#): No need to instantiate BlockCache for master which not carry table
- [HBASE-21260](#): The whole balancer plans might be aborted if there are more than one plans to move a same region
- [HBASE-21309](#): Increase the waiting timeout for TestProcedurePriority
- [HBASE-21271](#): [amv2] Don't throw UnsupportedOperationException when rollback called on Assign/Unassign; spiral of death
- [HBASE-21251](#): Refactor RegionMover
- [HBASE-21259](#): [amv2] Revived deadservers; recreated serverstatenode
- [HBASE-21303](#): [shell] clear_deadservers with no args fails
- [HBASE-21289](#): Remove the log "'hbase.regionserver.maxlogs' was deprecated." in AbstractFSWAL
- [HBASE-21242](#): [amv2] Miscellaneous minor log and assign procedure create improvements
- [HBASE-21254](#): Need to find a way to limit the number of proc wal files
- [HBASE-21103](#): nightly job should make sure cached yetus will run.
- [HBASE-21287](#): Allow configuring test master initialization wait time.

- [HBASE-20764](#): build broken when latest commit is gpg signed
- [HBASE-21271](#): [amv2] Don't throw UnsupportedOperationException when rollback called on Assign/Unassign; spiral of death
- [HBASE-21250](#): Addendum remove unused modification in hbase-server module
- [HBASE-21250](#): Refactor WALProcedureStore and add more comments for better understanding the implementation
- [HBASE-21223](#): [amv2] Remove abort_procedure from shell
- [HBASE-20506](#): Add doc and test for unused RetryCounter, useful-looking utility
- [HBASE-21185](#): - WALPrettyPrinter: Additional useful info to be printed by wal printer tool, for debugability purposes
- [HBASE-21281](#): Upgrade bouncycastle to latest - addendum adds test dependency
- [HBASE-21178](#): [BC break] : Get and Scan operation with a custom converter_class not working
- [HBASE-21282](#): Upgrade to latest jetty 9.2 and 9.3 versions
- [HBASE-21281](#): Upgrade bouncycastle to latest
- [HBASE-21261](#): Add log4j.properties for hbase-rsgroup tests
- [HBASE-19275](#): TestSnapshotFileCache never worked properly
- [HBASE-21196](#): HTableMultiplexer clears the meta cache after every put operation
- [HBASE-21249](#): Add jitter for ProcedureUtil.getBackoffTimeMs
- [HBASE-21233](#): Allow the procedure implementation to skip persistence of the state after a execution
- [HBASE-21237](#): Use CompatRemoteProcedureResolver to dispatch open/close region requests to RS
- [HBASE-21228](#): Memory leak since AbstractFSWAL caches Thread object and never clean later
- Revert "[HBASE-20657](#) Retrying RPC call for ModifyTableProcedure may get stuck"
- [HBASE-21212](#): Wrong flush time when update flush metric
- [HBASE-21208](#): Bytes#toShort doesn't work without unsafe
- Amend [HBASE-20704](#): Sometimes some compacted storefiles are not archived on region close
- [HBASE-21203](#): TestZKMainServer#testCommandLineWorks won't pass with default 4lw whitelist
- [HBASE-21206](#): Scan with batch size may return incomplete cells
- [HBASE-20704](#): Sometimes some compacted storefiles are not archived on region close
- [HBASE-21097](#): Flush pressure assertion may fail in testFlushThroughputTuning
- [HBASE-21189](#): flaky job should gather machine stats
- [HBASE-21190](#): Log files and count of entries in each as we load from the MasterProcWAL store
- [HBASE-21171](#): Reimplement the retry backoff logic for ReopenTableRegionsProcedure
- [HBASE-21179](#): Fix the number of actions in responseTooSlow log
- [HBASE-21174](#): [REST] Failed to parse empty qualifier in TableResource#getScanResource
- [HBASE-21181](#): Use the same filesystem for wal archive directory and wal directory
- [HBASE-21158](#): Empty qualifier cell is always returned when using QualifierFilter
- [HBASE-21144](#): AssignmentManager.waitForAssignment is not stable
- [HBASE-21143](#): Update findbugs-maven-plugin to 3.0.4
- [HBASE-21001](#): ReplicationObserver fails to load in HBase 2.0.0
- [HBASE-20307](#): LoadTestTool prints too much zookeeper logging (Colin Garcia)
- [HBASE-21157](#): Split TableInputFormatScan to individual tests
- [HBASE-21127](#): TableRecordReader need to handle cursor result too
- [HBASE-20892](#): [UI] Start / End keys are empty on table.jsp
- [HBASE-21132](#): return wrong result in rest multiget

HDP 3.0.1 provided HBase HBase 2.0.0 and the following Apache patches:

- [HBASE-20723](#): Custom hbase.wal.dir results in data loss because we write recovered edits into a different place than where the recovering region server looks for them
- [HBASE-20722](#): Make RegionServerTracker only depend on children changed event
- [HBASE-20668](#): Avoid permission change if ExportSnapshot's copy fails
- [HBASE-20657](#): Retrying RPC call for ModifyTableProcedure may get stuck

- [HBASE-20635](#): Add missing shaded-pb conversion method for UserPermission
- [HBASE-20627](#): Relocate RS Group pre/post hooks from RSGroupAdminServer to RSGroupAdminEndpoint
- [HBASE-20579](#): Include original exception in wrapped exception
- [HBASE-20579](#): Improve snapshot manifest copy in ExportSnapshot
- [HBASE-20547](#): Restore from backup will fail if done from a different file system
- [HBASE-20531](#): RS may throw NPE when close meta regions in shutdown procedure.
- [HBASE-20495](#): REST unit test fails with NoClassDefFoundError against hadoop3
- [HBASE-20196](#): Maintain all regions with same size in memstore flusher
- [HBASE-20136](#): TestKeyValue misses ClassRule and Category annotations
- [HBASE-19735](#): Create a client-tarball assembly
- [HBASE-19478](#): Utilize multi-get to speed up WAL file checking in BackupLogCleaner (Toshihiro Suzuki)
- [HBASE-18135](#): Implement mechanism for RegionServers to report file archival for space quotas
- [HBASE-18133](#): Decrease quota reaction latency by HBase
- [HBASE-17549](#): HBase-Spark Module: Corrected - Incorrect log at println and unwanted comment cod

Hive

This release provides Hive 3.1.0 and the listed Apache patches provided in the earlier versions of HDP 3.x.

HDP 3.1.0 provided Hive 3.1.0 and the following Apache patches:

- [HIVE-20827](#): Inconsistent results for empty arrays
- [HIVE-20998](#): HiveStrictManagedMigration utility should update DB/Table location as last migration steps
- [HIVE-20542](#): Incremental REPL DUMP progress information log message is incorrect
- [HIVE-20953](#): Remove a function from function registry when it can not be added to the metastore when creating it.
- [HIVE-20884](#): Support bootstrap of tables from HDP 2.6.5 to HDP 3.0. Added managed to external testcase.
- [HIVE-20926](#): Semi join reduction hint fails when bloom filter entries are high or when there are no stats
- [HIVE-20979](#): Fix memory leak in hive streaming
- [HIVE-20838](#): Timestamps with timezone are set to null when using the streaming API
- [HIVE-20988](#): Wrong results for group by queries with primary key on multiple columns
- [HIVE-20985](#): If select operator inputs are temporary columns vectorization may reuse some of them as output
- [HIVE-20842](#): Fix logic introduced in HIVE-20660 to estimate statistics for group by
- [HIVE-20951](#): LLAP: Set Xms to 50% always
- [HIVE-20822](#): Hive JDBC Storage Handler (Postgres): "java.lang.RuntimeException: While invoking method 'public org.apache.calcite.rel.rel2sql.SqlImplementor\$Result" while executing TPCDS query GRP_2_85
- [HIVE-20884](#): Support bootstrap of tables from HDP 2.6.5 to HDP 3.0.
- [HIVE-20926](#): Semi join reduction hint fails when bloom filter entries are high or when there are no stats
- [HIVE-20976](#): Hive JDBC Storage Handler: TPCDS_GRP_3_88 doesn't output proper results
- [HIVE-20626](#): Backport Log more details when druid metastore transaction fails
- [HIVE-20954](#): Vector RS operator is not using uniform hash function for TPC-DS query 95
- [HIVE-20691](#): Fix org.apache.hadoop.hive.cli.TestMiniLlapCliDriver.testCliDriver[cttl]
- [HIVE-10296](#): Cast exception observed when hive runs a multi join query on metastore , since postgres pushes the filter into the join, and ignores the condition before applying cast
- [HIVE-20940](#): Bridge cases in which Calcite's type resolution is more stricter than Hive.
- [HIVE-20944](#): Not validate stats during query compilation
- [HIVE-20949](#): Improve PKFK cardinality estimation in physical planning
- [HIVE-20940](#): Bridge cases in which Calcite's type resolution is more stricter than Hive.
- [HIVE-20937](#): Postgres jdbc query fail with "LIMIT must not be negative"
- [HIVE-20775](#): Factor cost of each SJ reduction when costing a follow-up reduction
- [HIVE-19882](#): Fix QTestUtil session lifecycle addendum

- [HIVE-20920](#): Use SQL constraints to improve join reordering algorithm
- [HIVE-20918](#): Flag to enable/disable pushdown of computation from Calcite into JDBC connection
- [HIVE-20775](#): Factor cost of each SJ reduction when costing a follow-up reduction
- [HIVE-20910](#): Insert in bucketed table fails due to dynamic partition sort optimization
- [HIVE-20682](#): Async query execution can potentially fail if shared sessionHive is closed by master thread
- [HIVE-19552](#): Backport fix Druid kafka mini tests
- [HIVE-19882](#): Fix QTestUtil session lifecycle
- [HIVE-20822](#): Hive JDBC Storage Handler: TPCDS_GRP_3_88 doesn't output proper results
- [HIVE-20101](#): BloomKFilter: Avoid using the local byte[] arrays entirely
- [HIVE-20899](#): Keytab URI for LLAP YARN Service is restrictive to support HDFS only
- [HIVE-20744](#): Use SQL constraints to improve join reordering algorithm
- [HIVE-20898](#): For time related functions arguments may not be casted to a non nullable type
- [HIVE-20746](#): HiveProtoHookLogger does not close file at end of day.
- [HIVE-20873](#): Use Murmur hash for VectorHashKeyWrapperTwoLong to reduce hash collision
- [HIVE-20881](#): Constant propagation oversimplifies projections
- [HIVE-20853](#): Expose ShuffleHandler.registerDag in the llap daemon API
- [HIVE-20804](#): Further improvements to group by optimization with constraints
- [HIVE-20768](#): Adding Tumbling Window UDF
- [HIVE-20813](#): udf to_epoch_milli need to support timestamp without time zone as well.
- [HIVE-20868](#): SMB Join fails intermittently when TezDummyOperator has child op in getFinalOp in MapRecordProcessor
- [HIVE-20617](#): Fix type of constants in IN expressions to have correct type
- [HIVE-20486](#): Vectorization support for Kafka Storage Handler
- [HIVE-20767](#): Multiple project between join operators may affect join reordering using constraints
- [HIVE-20822](#): Improvements to push computation to JDBC from Calcite
- [HIVE-20835](#): Interaction between constraints and MV rewriting may create loop in Calcite planner
- [HIVE-20821](#): Rewrite SUM0 into SUM + COALESCE combination
- [HIVE-20820](#): MV partition on clause position
- [HIVE-20788](#): Extended SJ reduction may backtrack columns incorrectly when creating filters
- [HIVE-20477](#): OptimizedSql is not shown if the expression contains INs
- [HIVE-20834](#): Hive QueryResultCache entries keeping reference to SemanticAnalyzer from cached query
- [HIVE-20829](#): JdbcStorageHandler range split throws NPE
- [HIVE-20830](#): JdbcStorageHandler range query assertion failure in some cases
- [HIVE-20815](#): JdbcRecordReader.next shall not eat exception
- [HIVE-20792](#): Inserting timestamp with zones truncates the data
- [HIVE-20638](#): Upgrade version of Jetty to 9.3.25.v20180904
- [HIVE-20765](#): HIVE- : fetch partitions for txn stats validation in get_aggr_stats with one call
- [HIVE-20718](#): Add perf cli driver with constraints
- [HIVE-20763](#): Add google cloud storage to the exim uri schema whitelist
- [HIVE-20703](#): Put dynamic sort partition optimization under cost based decision
- [HIVE-20762](#): NOTIFICATION_LOG cleanup interval is hardcoded as 60s and is too small
- [HIVE-6980](#): Drop table by using direct sql
- [HIVE-20716](#): Set default value for hive.cbo.stats.correlated.multi.key.joins to true
- [HIVE-18986](#): Table rename will run java.lang.StackOverflowError in dataNucleus if the table contains large number of columns
- [HIVE-20720](#): Add partition column option to JDBC handler
- [HIVE-20735](#): Adding Support for Kerberos Auth, Removed start/end offset columns, remove the best effort mode and made 2pc default for EOS
- [HIVE-20754](#): Hive: java.lang.NoClassDefFoundError: org/apache/hadoop/security/SaslRpcServer exception occurs while using Hive JDBC standalone jar

- [HIVE-20761](#): [Intermittent]Hive DDL's failing with Error- Couldn't acquire the DB log notification lock because we reached the maximum # of retries: 5 retries.
- [HIVE-20644](#):HIVE- : Avoid exposing sensitive information through a Hive Runtime exception
- [HIVE-20731](#): keystore file in JdbcStorageHandler should be authorized
- [HIVE-20649](#): LLAP aware memory manager for Orc writers
- [HIVE-20648](#): LLAP: Vector group by operator should use memory per executor
- [HIVE-20696](#): msck_*.q tests are broken
- [HIVE-20537](#): Multi-column joins estimates with uncorrelated columns different in CBO and Hive
- [HIVE-20639](#): Add ability to Write Data from Hive Table/Query to Kafka Topic
- [HIVE-20561](#): Use the position of the Kafka Consumer to track progress instead of Consumer Records offsets
- [HIVE-20481](#): Add the Kafka Key record as part of the row
- [HIVE-20377](#): Hive Kafka Storage Handler
- [HIVE-20710](#): Constant folding may not create null constants without types
- [HIVE-20719](#): UPDATE with hive.optimize.sort.dynamic.partition corrupts table
- [HIVE-20660](#): Group by statistics estimation could be improved by bounding the total number of rows to source table
- [HIVE-17043](#): Remove non unique columns from group by keys if not referenced later
- [HIVE-20692](#): Enable folding of NOT x IS (NOT) [TRUE|FALSE] expressions
- [HIVE-14431](#): Recognize COALESCE as CASE
- [HIVE-20618](#): During join selection BucketMapJoin might be chosen for non bucketed tables
- [HIVE-20712](#): HivePointLookupOptimizer should extract deep cases
- [HIVE-20498](#):HIVE- Support date type for column stats autogather
- [HIVE-20657](#): pre-allocate LLAP cache at init time
- [HIVE-18871](#): hive on tez execution error due to set hive.aux.jars.path to hdfs://
- [HIVE-20583](#): Use canonical hostname only for kerberos auth in HiveConnection
- [HIVE-19993](#): Using a table alias which also appears as a column name is not possible
- [HIVE-20183](#): Inserting from bucketed table can cause data loss, if the source table contains empty bucket
- [HIVE-19850](#): Dynamic partition pruning in Tez is leading to 'No work found for tablescan' error
- [HIVE-20349](#): Implement Retry Logic in HiveDruidSplit for Scan Queries
- [HIVE-20700](#): Add config to disable rollup for druid tables
- [HIVE-20698](#): Add better message for NPE when inserting rows with null timestamp to druid
- [HIVE-20686](#): Sync query IDs between druid and Hive
- [HIVE-20684](#): Make compute stats work for Druid tables
- [HIVE-20582](#): Avoid hflush in hive proto logger
- [HIVE-20702](#): Account for overhead from datastructure aware estimations during mapjoin selection
- [HIVE-20704](#): Extend HivePreFilteringRule to support other functions
- [HIVE-20651](#): JdbcStorageHandler password should be encrypted
- [HIVE-20677](#): JDBC handler ordering problem - single split flag [HIVE-20652](#): JdbcStorageHandler push join of two different datasource to jdbc driver [HIVE-20669](#): JdbcStorageHandler push union of two different datasource to jdbc driver
- [HIVE-20455](#): Log spew from security.authorization.PrivilegeSynchronizer.run
- [HIVE-20676](#): HiveServer2: PrivilegeSynchronizer is not set to daemon status
- [HIVE-17040](#): Join elimination in the presence of FK relationship
- [HIVE-20636](#): Improve number of null values estimation after outer join
- [HIVE-20326](#): Create constraints with RELY as default instead of NO RELY
- [HIVE-20646](#): Partition condition not being pushed into metastore query if it has IS NOT NULL.
- [HIVE-18778](#): Needs to capture input/output entities in explain
- [HIVE-20623](#): Shared work: Extend sharing of map-join cache entries in LLAP
- [HIVE-20095](#): Fix feature to push computation to jdbc external tables

- [HIVE-20540](#):HIVE- : Vectorization : Support loading bucketed tables using sorted dynamic partition optimizer - II
- [HIVE-20603](#): "Wrong FS" error when inserting to partition after changing table location filesystem
- [HIVE-20620](#): manifest collisions when inserting into bucketed sorted MM tables with dynamic partitioning
- [HIVE-20612](#): Create new join multi-key correlation flag for CBO
- [HIVE-20625](#): Regex patterns not working in SHOW MATERIALIZED VIEWS '<pattern>'
- [HIVE-20593](#): Load Data for partitioned ACID tables fails with bucketId out of range: -1
- [HIVE-20522](#): HiveFilterSetOpTransposeRule may throw assertion error due to nullability of fields
- [HIVE-20517](#): Creation of staging directory and Move operation is taking time in S3
- [HIVE-20499](#): GetTablesOperation pull all the tables meta irrespective of auth.
- [HIVE-19432](#): GetTablesOperation is too slow if the hive has too many databases and tables
- [HIVE-18873](#): Skipping predicate pushdown for MR silently at HiveInputFormat can cause storage handlers to produce erroneous result
- [HIVE-20267](#): Concurrent async queries intermittently fails with LockException and cause memory leak.
- [HIVE-20632](#): Query with get_splits UDF fails if materialized view is created on queried table.
- [HIVE-20521](#): HS2 doAs=true has permission issue with hadoop.tmp.dir, with MR and S3A filesystem .
- [HIVE-1118](#): background updates fix
- [HIVE-20399](#): CTAS w/a custom table location that is not fully qualified fails for MM tables
- [HIVE-20311](#): add txn stats checks to some more paths
- [HIVE-20507](#): Beeline: Add a utility command to retrieve all uris from beeline-site.xml
- [HIVE-20503](#): Use datastructure aware estimations during mapjoin selection
- [HIVE-20513](#): Vectorization: Improve Fast Vector MapJoin Bytes Hash Tables
- [HIVE-20296](#): Improve HivePointLookupOptimizerRule to be able to extract from more sophisticated contexts
- [HIVE-20541](#): NoSuchElementException in REPL dump.
- [HIVE-19985](#): Benchmark - LLAP performance issue for fully compacted ACID tables
- [HIVE-20409](#): Hive ACID: Update/delete/merge leave behind the staging directory
- [HIVE-20491](#): Fix mapjoin size estimations for Fast implementation
- [HIVE-19097](#): related equals and in operators may cause inaccurate stats estimations
- [HIVE-20439](#): Use the inflated memory limit during join selection for llap
- [HIVE-20260](#): NDV of a column shouldn't be scaled when row count is changed by filter on another column
- [HIVE-20418](#): LLAP IO may not handle ORC files that have row index disabled correctly for queries with no columns selected
- [HIVE-20508](#):HIVE- : Hive does not support user names of type "user@realm"
- [HIVE-20193](#): cboInfo is not present in the explain plan json
- [HIVE-20510](#): Vectorization : Support loading bucketed tables using sorted dynamic partition optimizer
- [HIVE-20515](#): Empty query results when using results cache and query temp dir, results cache dir in different filesystems
- [HIVE-20476](#): CopyUtils used by REPL LOAD and EXPORT/IMPORT operations ignore distcp error
- [HIVE-20450](#): Add replication test for LOAD command on ACID table.
- [HIVE-20433](#): Implicit String to Timestamp conversion is slow
- [HIVE-20446](#): CachedStore: bug UT fixes when CachedStore is enabled
- [HIVE-19748](#): CREATE, INSERT commands fails with MetaException on HDI azure cluster due to NPE in MetaStoreUtils.decimalToDouble
- [HIVE-18624](#): Parsing time is extremely high (~10 min) for queries with complex select expressions
- [HIVE-20187](#): Incorrect query results in hive when hive.convert.join.bucket.mapjoin.tez is set to true
- [HIVE-20443](#): txn stats cleanup in compaction txn handler is unneeded
- [HIVE-20444](#): Parameter is not properly quoted in DbNotificationListener.addWriteNotificationLog
- [HIVE-20431](#): txn stats write ID check triggers on set location
- [HIVE-20329](#): Long running repl load (incr/bootstrap) causing OOM error
- [HIVE-19924](#): Tag distcp jobs run by Repl Load

- [HIVE-20411](#): Hive.loadPartition doesn't support catalogs
- [HIVE-20412](#): Hive upgrade: HBasestoragehandler tables fail to upgrade at HiveMetaStoreClient.alter_table with Null Pointer Exception - committing HIVE-20412.1.patch directly to HDP
- [HIVE-20410](#): aborted Insert Overwrite on transactional table causes "Not enough history available for..." error
- [HIVE-20407](#): Hive upgrade: Alter table(orc-bucket) owned by hive fails at seedWriteIdOnAcidConversion with Duplicate entry for key 'NEXT_WRITE_ID_IDX'
- [HIVE-20413](#): "cannot insert NULL" for TXN_WRITE_NOTIFICATION_LOG in Oracle
- [HIVE-20393](#): Semijoin Reduction : markSemiJoinForDPP behaves inconsistently
- [HIVE-20407](#): add txn stats support to upgrade tool
- [HIVE-14898](#): HS2 shouldn't log callstack for an empty auth header error
- [HIVE-20378](#): don't update stats during alter for txn table conversion
- [HIVE-20357](#): Introduce initOrUpgradeSchema option to schema tool
- [HIVE-19250](#): Schema column definitions inconsistencies in MySQL
- [HIVE-20391](#): HiveAggregateReduceFunctionsRule may infer wrong return type when decomposing aggregate function
- [HIVE-20383](#): Invalid queue name and synchronisation issues in hive proto events hook.
- [HIVE-20389](#): NPE in SessionStateUserAuthenticator when authenticator=SessionStateUserAuthenticator
- [HIVE-20345](#): Drop database may hang if the tables get deleted from a different call
- [HIVE-20033](#): PowerBI unable to connect to LLAP cluster with large Databases/Tables
- [HIVE-20098](#): Statistics: NPE when getting Date column partition statistics
- [HIVE-19860](#): HiveServer2 ObjectInspectorFactory memory leak with cachedUnionStructObjectInspector
- [HIVE-20382](#): Materialized views: Introduce heuristic to favour incremental rebuild
- [HIVE-20379](#): Rewriting with partitioned materialized views may reference wrong column
- [HIVE-18620](#): Improve error message while dropping a table that is part of a materialized view
- [HIVE-20067](#): mm_all test failing.
- [HIVE-19316](#): StatsTask fails due to ClassCastException
- [HIVE-20149](#): TestHiveCli failing/timing out
- [HIVE-20340](#): Druid Needs Explicit CASTs from Timestamp to STRING when the output of timestamp function is used as String
- [HIVE-20344](#): PrivilegeSynchronizer for SBA might hit AccessControlException
- [HIVE-20320](#): Turn on hive.optimize.remove.sql_count_check flag
- [HIVE-20354](#): HIVE- : Semijoin hints dont work with merge statements
- [HIVE-20264](#): Bootstrap repl dump with concurrent write and drop of ACID table makes target inconsistent
- [HIVE-20361](#): ReplDumpTaskTest is failing
- [HIVE-20316](#): Skip external table file listing for create table event
- [HIVE-19800](#): TestUpgradeTool.testPostUpgrade failing
- [HIVE-20337](#): Queries returning incorrect results using CachedStore. ObjectStore works fine
- [HIVE-20347](#): hive.optimize.sort.dynamic.partition should work with partitioned CTAS and MV
- [HIVE-20336](#): Masking and filtering policies for materialized views
- [HIVE-20301](#): Enable vectorization for materialized view rewriting tests
- [HIVE-20314](#): Include partition pruning in materialized view rewriting
- [HIVE-20302](#): LLAP: non-vectorized execution in IO ignores virtual columns, including ROW__ID
- [HIVE-19408](#): Improve show materialized views statement to show more information about invalidation
- [HIVE-20335](#): Add tests for materialized view rewriting with composite aggregation functions
- [HIVE-20165](#): Enable ZLIB for streaming ingest
- [HIVE-20147](#): Hive streaming ingest is contented on synchronized logging
- [HIVE-19800](#): HIVE- PreUpgrade Tool: Support upgrade from >= 2.6.0
- [HIVE-20315](#): Vectorization: Fix more NULL Wrong Results issues and avoid unnecessary casts/conversions : Vectorization: Backport HIVE-20315 to Atlantic-M05: Vectorization: Fix more NULL Wrong Results issues and avoid unnecessary casts/conversions
- [HIVE-19812](#): Disable external table replication by default via a configuration property

- [HIVE-20277](#): Vectorization: Case expressions that return BOOLEAN are not supported for FILTER
- [HIVE-20130](#): Better logging for information schema synchronizer
- [HIVE-20118](#): SessionStateUserAuthenticator.getGroupNames is always empty
- [HIVE-20292](#): Adding constraints on TPCDS schema, changes query plan for query93 and degrades its perf by 4x times
- [HIVE-19970](#): Replication dump has a NPE when table is empty
- [HIVE-20120](#): Incremental repl load DAG generation is causing OOM error
- [HIVE-20322](#): FlakyTest: TestMiniDruidCliDriver
- [HIVE-19829](#): Incremental replication load should create tasks in execution phase rather than semantic phase
- [HIVE-20322](#): FlakyTest: TestMiniDruidCliDriver
- [HIVE-19927](#): Last Repl ID set by bootstrap dump is not proper and may cause loss of data if have ACID tables.
- [HIVE-17683](#): Add explain locks <sql> command
- [HIVE-20209](#): Metastore connection fails for first attempt in repl dump
- [HIVE-20294](#): Vectorization: Fix NULL Wrong Results issues in COALESCE ELT
- [HIVE-20245](#): Vectorization: Fix NULL Wrong Results issues in BETWEEN IN
- [HIVE-20207](#): Vectorization: Fix NULL Wrong Results issues in Filter Compare
- [HIVE-20281](#): SharedWorkOptimizer fails with 'operator cache contents and actual plan differ'
- [HIVE-20252](#): Semijoin Reduction : Cycles due to semi join branch may remain undetected if small table side has a map join upstream.
- [HIVE-14493](#): Partitioning support for materialized views
- [HIVE-20241](#): Support partitioning spec in CTAS statements
- [HIVE-20212](#): Hiveserver2 in http mode emitting metric default.General.open_connections incorrectly
- [HIVE-20105](#): Druid-Hive: tpchs query on timestamp throws java.lang.IllegalArgumentException: Cannot create timestamp, parsing error
- [HIVE-20263](#): Typo in HiveReduceExpressionsWithStatsRule variable
- [HIVE-20090](#): Extend creation of semijoin reduction filters to be able to discover new opportunities
- [HIVE-19568](#): Active/Passive HS2 HA: Disallow direct connection to passive HS2 instance
- [HIVE-20082](#): HiveDecimal to string conversion doesn't format the decimal correctly
- [HIVE-20204](#): Type conversion during IN () comparisons is using different rules from other comparison operations
- [HIVE-20240](#): Semijoin Reduction : Use local variable to check for external table condition
- [HIVE-19313](#): TestJdbcWithDBTokenStoreNoDoAs tests are failing
- [HIVE-20228](#): configure repl configuration directories based on user running hiveserver2
- [HIVE-19935](#): Hive WM session killed: Failed to update LLAP tasks count
- [HIVE-19800](#): TestUpgradeTool.testPostUpgrade failing
- [HIVE-19891](#): inserting into external tables with custom partition directories may cause data loss
- [HIVE-4367](#): enhance TRUNCATE syntax to drop data of external table
- [HIVE-19992](#): Vectorization: Follow-on to HIVE-19951 --> add call to SchemaEvolution.isOnlyImplicitConversion to disable encoded LLAP I/O for ORC only when data type conversion is not implicit : Vectorization: Backport HIVE-19992 to Atlantic-M05: Follow-on to HIVE-19951 --> add call to SchemaEvolution.isOnlyImplicitConversion to disable encoded LLAP I/O for ORC only when data type conversion is not implicit
- [HIVE-20069](#): Fix reoptimization in case of DPP and Semijoin optimization
- [HIVE-20152](#): reset db state, when repl dump fails, so rename table can be done
- [HIVE-20197](#): Vectorization: Add DECIMAL_64 testing, add Date/Interval/Timestamp arithmetic, and add more GROUP BY Aggregation tests : Vectorization: Backport HIVE-20197 to Atlantic-M05: Add DECIMAL_64 testing, add Date/Interval/Timestamp arithmetic, and add more GROUP BY Aggregation tests
- [HIVE-20174](#): Vectorization: Fix NULL Wrong Results issues in GROUP BY Aggregation Functions
- [HIVE-20192](#): HiveServer2 with embedded metastore is leaking JDOPersistenceManager objects
- [HIVE-19104](#): When test MetaStore is started with retry the instances should be independent
- [HIVE-20164](#): Murmur Hash : Make sure CTAS and IAS use correct bucketing version
- [HIVE-20116](#):HIVE- - TestOperationLoggingAPIWithTez failing

- [HIVE-20088](#): Beeline config location path is assembled incorrectly
- [HIVE-20011](#): Move away from append mode in proto logging hook
- [HIVE-19360](#): CBO: Add an "optimizedSQL" to QueryPlan object
- [HIVE-20127](#):HIVE- : fix some issues with LLAP Parquet cache
- [HIVE-19387](#):HIVE- Upgrade: Post the second truncate call, select query fails to fetch new results if query.results.cache is enabled
- [HIVE-20129](#): Revert to position based schema evolution for orc tables
- [HIVE-20126](#): OrcInputFormat does not pass conf to orc reader options
- [HIVE-20123](#): Fix masking tests after HIVE-19617
- [HIVE-19628](#):HIVE- : possible NPE in LLAP testSigning
- [HIVE-19759](#): Flaky test: TestRpc#testServerPort
- [HIVE-20111](#): Making storage handler tables external by default breaks HBase
- [HIVE-20112](#): Making storage handler tables external by default breaks Accumulo
- [HIVE-20102](#): Ranger row-level filter policies don't work with CTAS
- [HIVE-20001](#): With doas set to true, running select query as hrt_qa user on external table fails due to permission denied to read warehouse/tablespace/managed directory
- [HIVE-20100](#): OpTraits : Select Optraits should stop when a mismatch is detected
- [HIVE-20085](#): review comments and fix drop table to honour external table purge property, enable ctas by default on external tables
- [HIVE-20039](#): Bucket pruning: Left Outer Join on bucketed table gives wrong result
- [HIVE-20085](#): Make all Druid Tables as EXTERNAL tables
- [HIVE-20067](#): fix InsertEvent on mm tables to not cause failing capability checks
- [HIVE-20025](#): Clean-up of event files created by HiveProtoLoggingHook
- [HIVE-17840](#): HiveMetaStore eats exception if transactionalListeners.notifyEvent fail
- [HIVE-20051](#): Skip authorization for temp tables
- [HIVE-19938](#): Upgrade scripts for information schema
- [HIVE-19326](#): stats auto gather: incorrect aggregation during UNION queries
- [HIVE-20059](#): Hive streaming should try shade prefix unconditionally on exception
- [HIVE-20028](#): Metastore client cache config is used incorrectly
- [HIVE-20038](#): Updates/deletes on non-bucketed acid tables throw NPE
- [HIVE-20002](#): Shipping jdbd-storage-handler dependency jars in LLAP
- [HIVE-19967](#): SMB Join : Need Optraits for PTFOperator ala GBY Op
- [HIVE-20004](#):HIVE-" Wrong scale used by ConvertDecimal64ToDecimal results in incorrect results
- [HIVE-19995](#): Aggregate row traffic for acid tables
- [HIVE-20009](#): Fix runtime stats for merge statement
- [HIVE-19989](#): Metastore uses wrong application name for HADOOP2 metrics
- [HIVE-20035](#): Serialize boolean as long for druid
- [HIVE-20013](#): Add an Implicit cast to date type for to_date function
- [HIVE-18140](#): Partitioned tables statistics can go wrong in basic stats mixed case
- [HIVE-18729](#): Druid Time column type
- [HIVE-19938](#): Upgrade scripts for information schema
- [HIVE-20010](#): Fix create view over literals
- [HIVE-19753](#): Strict managed tables mode in Hive
- [HIVE-19972](#): Followup to HIVE-19928 : Fix the check for managed table
- [HIVE-19481](#): Tablesample uses incorrect logic to pick files corresponding to buckets.
- [HIVE-19981](#): Managed tables converted to external tables by the HiveStrictManagedMigration utility should be set to delete data when the table is dropped
- [HIVE-19898](#): Disable TransactionalValidationListener when the table is not in the Hive catalog
- [HIVE-19904](#): Load data rewrite into Tez job fails for ACID
- [HIVE-19775](#): Schematool should use HS2 embedded mode in privileged auth mode

- [HIVE-19800](#): HIVE-19598 + adjustments. needs HIVE-19750 for post upgrade file rename to work properly
- [HIVE-19440](#): fix merge conflicts that got dropped
- [HIVE-19258](#): add originals support to MM tables
- [HIVE-19476](#): Fix failures in TestReplicationScenariosAcidTables, TestReplicationOnHDFSEncryptedZones and TestCopyUtils
- [HIVE-17657](#): export/import for MM tables is broken
- [HIVE-19312](#): MM tables don't work with BucketizedHIF
- [HIVE-19825](#): HiveServer2 leader selection shall use different zookeeper znode
- [HIVE-19728](#): beeline with USE_BEELINE_FOR_HIVE_CLI fails when trying to set hive.aux.jars.path
- [HIVE-19485](#): dump directory for non native tables should not be created
- [HIVE-19440](#): Make StorageBasedAuthorizer work with information schema
- [HIVE-19465](#): Upgrade ORC to 1.5.0
- [HIVE-19631](#): reduce epic locking in AbstractService
- [HIVE-19258](#): add originals support to MM tables
- [HIVE-19660](#): update branch-3 to be version 3.1 and fix storage-api mismatch
- [HIVE-14247](#): Disable parallel query execution within a session
- [HIVE-18193](#): Migrate existing ACID tables to use table Write ID instead of global txn ID
- [HIVE-19476](#): Fix failures in TestReplicationScenariosAcidTables, TestReplicationOnHDFSEncryptedZones and TestCopyUtils
- [HIVE-19435](#): Incremental replication cause data loss if a table is dropped followed by create and insert-into with different partition type
- [HIVE-19322](#): Broken test: TestNegativeMinimrCliDriver#testCliDriver[minimr_broken_pipe]
- [HIVE-17657](#): Export/import for MM tables is broken
- [HIVE-19304](#): Update templates.py based on config changes in YARN-7142 and YARN-8122
- [HIVE-19194](#): addendum for branch-3: fix imports
- [HIVE-19193](#): TestActivePassiveHA fails
- [HIVE-19014](#): Utilize YARN-8028 (queue ACL check) in Hive Tez session pool
- [HIVE-19072](#): incorrect token handling for LLAP plugin endpoint
- [HIVE-18971](#): add HS2 WM metrics for use in Grafana and such
- [HIVE-18952](#): Tez session disconnect and reconnect on HS2 HA failover INTERNAL COMMIT
- [HIVE-18717](#): Avoid transitive dependency on jetty 6.x
- [HIVE-18414](#): upgrade to tez-0.9.1
- [HIVE-18037](#): align kerberos keytab and principal with YARN Service
- [HIVE-14498](#): Freshness period for query rewriting using materialized views
- [HIVE-18096](#): add a user-friendly show plan command
- [HIVE-18037](#): plug-in llapstatus with YARN Service
- [HIVE-18037](#): fix issue with replacement of % %(queue.string)s with correct queue value
- [HIVE-15393](#): Update Guava version
- [HIVE-18150](#): Upgrade Spark Version to 2.2.0

HDP 3.0.0 and HDP 3.0.1 provided Hive 3.1.0 and no additional Apache patches.

Kafka

This release provides Kafka 2.0.0 and no additional Apache patches.

HDP 3.0.1 provided Kafka 1.1.1 and no additional Apache patches.

HDP 3.0.0 provided Kafka 1.0.1 and no additional Apache patches. A new feature was added to capture producer and topic partition level metrics. This feature is available in HDP 3.0.0. It is not available in Apache Kafka 1.0.1.

Knox

This release provides Knox 1.0.0 and no Apache patches.

HDP 3.0.0 and HDP 3.0.1 provided Knox 1.0.0 and no additional Apache patches.

Livy

This release provides Livy 0.5.0 and no additional patches.

HDP 3.0.0 and HDP 3.0.1 provided Livy 0.5.0 and no additional Apache patches.

Oozie

This release provides Oozie 4.3.1 and the additional Apache patches.

- [OOZIE-3194](#): [tools] Oozie should set proper permissions to sharelib after upload
- [OOZIE-3439](#): Hive2 action is not parsing application ID for TEZ from log file properly
- [OOZIE-3440](#) :[action] Oozie Spark action replaces path symlink # to %23
- [OOZIE-3370](#): Property filtering is not consistent across job submission

HDP 3.0.1, and HDP 3.1.0 provided Oozie 4.3.1 and the following Apache patches:

- [OOZIE-3304](#): Parsing sharelib timestamps is not threadsafe (dionusos, matijhs via andras.piros)
- [OOZIE-3135](#): Configure log4j2 in SqoopMain
- [OOZIE-3072](#): oozie.service.HadoopAccessorService.action.configurations should overwrite default values set in Hadoop's configuration files (gezapeti)
- [OOZIE-2916](#): Set a job name for the MR Action's child job (asasvari)
- [OOZIE-2494](#): Cron syntax not handling DST properly (kmarion via andras.piros)
- [OOZIE-2726](#): Flaky test due to daylight saving changes (sasishsaley, andras.piros via gezapeti)
- [OOZIE-3173](#) : Coordinator job with frequency using cron syntax creates only one action in catchup mode (andras.piros)
- [OOZIE-2691](#): Show workflow action retry information in UI
- [OOZIE-3109](#): [log-streaming] Escape HTML-specific characters (dionusos via andras.piros)
- [OOZIE-2815](#): Oozie not always display job log
- [OOZIE-3075](#): Follow-up on OOOZIE-3054: create the lib directory if it doesn't exist (pbacsko)
- [OOZIE-3054](#): Disable erasure coding for sharelib if Oozie runs on Hadoop 3 (pbacsko)
- [OOZIE-3048](#): Check El Functions for the coordinator action (satishsaley)
- [OOZIE-3004](#) : Forked action retry info is not working
- [OOZIE-2691](#): Show workflow action retry information in UI
- [OOZIE-3009](#): Number of Oozie tests executed dropped after OOOZIE-2854 (andras.piros via asasvari)
- [OOZIE-2854](#): Oozie should handle transient database problems (andras.piros via gezapeti)
- [OOZIE-3026](#): fix openjpa enhancer stage during build for logging (dbdist13, andras.piros via pbacsko)
- [OOZIE-2816](#): Strip out the first command word from Sqoop action if its "sqoop" (harsh)
- [OOZIE-3113](#): Retry for ZK lock release (satishsaley)
- [OOZIE-3208](#): "It should never happen" error messages should be more specific to root cause (kmarion via andras.piros)
- [OOZIE-2940](#): Possible NullPointerException in WorkflowActionBean (dionusos via gezapeti)
- [OOZIE-3297](#): Retry logic does not handle the exception from BulkJPAExecutor properly (andras.piros)
- [OOZIE-2733](#): change org.apache.hadoop.fs.permission.AccessControlException to org.apache.hadoop.security.AccessControlException (gezapeti)
- [OOZIE-3021](#) : Error on job or SLA event listening: WARN logs instead of DEBUG (andras.piros via asasvari)
- [OOZIE-2387](#): Oozie is Unable to handle Spaces in file/archive tag. (asasvari via gezapeti)

- [OOZIE-2818](#): Can't overwrite `oozie.action.max.output.data` on a per-workflow basis (asasvari via pbacsko)
- [OOZIE-2041](#): Add an admin command to run the `PurgeXCommand` (abhishekbafna)
- [OOZIE-2875](#): Typo in ssh action twiki docs (Dongying Jiao via gezapeti)
- [OOZIE-1401](#): `PurgeCommand` should purge the workflow jobs w/o `end_time` (asasvari)
- [OOZIE-3156](#): Retry SSH action check when cannot connect to remote host (txsing via andras.piros)
- [OOZIE-3183](#): Better logging for `SshActionExecutor` and extended HA capability when calling to remote host (andras.piros)
- [OOZIE-2769](#): Extend FS action to allow setrep on a file (Artem Ervits via gezapeti)
- [OOZIE-3085](#): Improve logging in `ActionExecutors`: add log entry for `start()`, `end()` and `kill()` (kmarion via andras.piros)

HDP 3.0.0 provided Oozie 4.3.1 and the following Apache patches:

- [OOZIE-2802](#):: Spark action failure on Spark 2.1.0 due to duplicate sharelibs.
- [OOZIE-2658](#):: `--driver-class-path` can overwrite the classpath in `SparkMain`.
- [OOZIE-2787](#):: Oozie distributes application jar twice making the spark job fail.
- [OOZIE-2606](#):: Set `spark.yarn.jars` to fix Spark 2.0 with Oozie.
- [OOZIE-2436](#):: Fork/join workflow fails with "`oozie.action.yarn.tag must not be null`".
- [OOZIE-2578](#):: Oozie example distcp job fails to run within an encrypted zone with checksum match error.
- [OOZIE-2362](#):: SQL injection in `BulkJPAExecutor`.
- [OOZIE-2577](#):: Flaky tests `TestCoordActionInputCheckXCommand.testTimeout` and `testTimeoutWithException`.
- [OOZIE-2570](#):: remove `-PtestPatchCompile` from patch testing as there is no such profile.
- [OOZIE-2504](#):: Create a `log4j.properties` under `HADOOP_CONF_DIR` in Shell Action.
- [OOZIE-2567](#):: HCat connection is not closed while getting heat cred.
- [OOZIE-2547](#):: Add `mapreduce.job.cache.files` to spark action.
- [OOZIE-2550](#):: Flaky tests in `TestZKUUIDService.java`.
- [OOZIE-2445](#):: Doc for - Specifying coordinator input datasets in more logical ways.
- [OOZIE-2541](#):: Possible resource leak in `Hive2Credentials`.
- [OOZIE-2563](#):: Pass `spark-defaults.conf` to spark action.
- [OOZIE-2556](#):: `TestAbandonedCoordChecker.testCatchupJob` is flaky.
- [OOZIE-2522](#):: There can be multiple coord submit from bundle in case of ZK glitch.
- [OOZIE-2553](#):: Cred tag is required for all actions in the workflow even if an action does not require it.
- [OOZIE-2503](#):: show `ChildJobURLs` to spark action.
- [OOZIE-2551](#):: Feature request: epoch timestamp generation.
- [OOZIE-2542](#):: Option to disable OpenJPA `BrokerImpl` finalization.
- [OOZIE-2447](#):: Illegal character `0x0` oozie client.
- [OOZIE-2548](#):: Flaky test `TestZKLocksService.testLockReaper`.
- [OOZIE-2546](#):: Improperly closed resources in `OozieDBCLI`.
- [OOZIE-2476](#):: When one of the action from fork fails with transient error, WF never joins .
- [OOZIE-2475](#):: Oozie does not cleanup action dir of killed actions.
- [OOZIE-2535](#):: User can't disable uber mode.
- [OOZIE-2482](#):: Pyspark job fails with Oozie.
- [OOZIE-2467](#):: Oozie can shutdown itself on long GC pause.
- [OOZIE-2537](#):: `SqoopMain` does not set up `log4j` properly.
- [OOZIE-2532](#):: patch apply does not handle binary files.
- [OOZIE-2330](#):: Spark action should take the global `jobTracker` and `nameNode` configs by default and allow file and archive elements.
- [OOZIE-2510](#):: `TestLogStreamer.testStreamLog` fails in other timezones.
- [OOZIE-2531](#):: Prevent Spark trying for token which is already available.
- [OOZIE-2509](#):: SLA job status can stuck in running state.
- [OOZIE-2529](#):: Support adding secret keys to `Credentials` of `Launcher`.
- [OOZIE-1402](#):: Increase retry interval for non-progressing coordinator action with fix value .

- [OOZIE-2512](#):: ShareLibservice returns incorrect path for jar.
- [OOZIE-2471](#):: Show child job url tab for distcp.
- [OOZIE-2511](#):: SubWorkflow missing variable set from option if config-default is present in parent workflow.
- [OOZIE-2391](#):: spark-opts value in workflow.xml is not parsed properly.
- [OOZIE-2489](#):: XML parsing is vulnerable.
- [OOZIE-2485](#):: Oozie client keeps trying to use expired auth token.
- [OOZIE-2490](#):: Oozie can't set hadoop.security.token.service.use_ip.
- [OOZIE-2474](#):: <job-xml> is not being applied to the launcher job.
- [OOZIE-2486](#):: TestSLAEventsGetForFilterJPAExecutor is flakey.
- [OOZIE-2481](#):: Add YARN_CONF_DIR in the Shell action.
- [OOZIE-2492](#):: JSON security issue in js code.
- [OOZIE-2429](#):: TestEventGeneration test is flakey.
- [OOZIE-2466](#):: Repeated failure of TestMetricsInstrumentation.testSamplers.
- [OOZIE-2470](#):: Remove infinite socket timeouts in the Oozie email action.
- [OOZIE-2246](#):: CoordinatorInputCheckCommand does not behave properly when har file is one of data dependency and doesn't exist.
- [OOZIE-2446](#):: Job does not fail during submission if non existent credential is specified .
- [OOZIE-2283](#):: Documentation should not say that System.exit is not allowed in Java Action .
- [OOZIE-2400](#):: Workflow xml configuration parser cannot deal with namespace prefix.
- [OOZIE-2452](#):: Coordinator Functional Specification - EL Constants Typo.
- [OOZIE-2173](#):: DISCLAIMER.txt still refers to Apache Incubator.
- [OOZIE-2312](#):: oozie doesn't purge audit and error log.
- [OOZIE-2431](#):: Support hive.metastore.uris as alias for hcat.metastore.uri.
- [OOZIE-1976](#):: Specifying coordinator input datasets in more logical ways.
- [OOZIE-2444](#):: Need conditional logic in bundles.
- [OOZIE-2394](#):: Oozie can execute command without holding lock.
- [OOZIE-1922](#):: MemoryLocksService fails if lock is acquired multiple times in same thread and released.
- [OOZIE-2432](#):: TestPurgeXCommand fails.
- [OOZIE-2434](#):: inconsistent coord action status and workflow job status.
- [OOZIE-2438](#):: Oozie client "jobs -filter" diagnostic message clarification.
- [OOZIE-2165](#):: Job log fetching can fail in Oozie HA mode when using doAs impersonation.
- [OOZIE-2430](#):: Add root logger for hive,sqoop action.
- [OOZIE-2441](#):: SubWorkflow action with propagate-configuration but no global section throws NPE on submit.
- [OOZIE-2370](#):: Docs for Coordinator Action Status Notification has wrong property name.
- [OOZIE-2419](#):: HBase credentials are not correctly proxied.
- [OOZIE-2439](#):: FS Action no longer uses name-node from global section or default NN.
- [OOZIE-2435](#):: TestCoordChangeXCommand is flakey.
- [OOZIE-2428](#):: TestSLAService, TestSLAEventGeneration flakey tests.
- [OOZIE-2437](#):: Exclude additional Hadoop dependencies from Spark Sharelib.
- [OOZIE-2030](#):: Configuration properties from global section is not getting set in Hadoop job conf when using sub-workflow action in Oozie workflow.xml.
- [OOZIE-2380](#):: Oozie Hive action failed with wrong tmp path.
- [OOZIE-2222](#):: Oozie UI parent job should be clickable.
- [OOZIE-2407](#):: AbandonedService should not send mail if there is no abandoned coord.
- [OOZIE-2402](#):: oozie-setup.sh sharelib create takes a long time on large clusters.
- [OOZIE-2185](#):: Make oozie cli source conf/oozie-client-env.sh.
- [OOZIE-2413](#):: Kerberos credentials can expire if the KDC is slow to respond.
- [OOZIE-2411](#):: Add BCC to oozie email action.
- [OOZIE-2397](#):: LAST_ONLY and NONE don't properly handle READY actions.
- [OOZIE-2401](#):: Typo in twiki docs with FilesAchives instead of FilesArchives.

- [OOZIE-2168](#):: Oozie flow and action names have 50 char limit.
- [OOZIE-2346](#):: Add sub-workflow information like the super-parent id and workflow depth into the 'oozie.job.info' property.
- [OOZIE-2303](#):: Typo in documentation.
- [OOZIE-2328](#):: Coordinator endtime change should check if the last action is in database .
- [OOZIE-2367](#):: fs delete should support skipTrash option.
- [OOZIE-2368](#):: coord:dateOffset and coord:dateTzOffset can't be used for coord initial-instance.
- [OOZIE-2369](#):: coord:offset doesn't resolve correctly.
- [OOZIE-2251](#):: Expose instrumental matrices in Realtime Graphing tool.
- [OOZIE-2314](#):: Unable to kill old instance child job by workflow or coord rerun by Launcher .
- [OOZIE-2386](#):: org.apache.oozie.util.TestXConfiguration.testSubstituteVar is flakey.
- [OOZIE-2385](#):: org.apache.oozie.TestCoordinatorEngineStreamLog.testCoordLogStreaming is flakey.
- [OOZIE-2382](#):: org.apache.oozie.action.hadoop.TestPigMain.testPig_withNullExternalID is flakey.
- [OOZIE-2379](#):: org.apache.oozie.command.coord.TestCoordPushDependencyCheckXCommand.testRequeueOnException is flakey.
- [OOZIE-2378](#):: org.apache.oozie.client.TestOozieCLI.testAdminInstrumentation is flakey.
- [OOZIE-2377](#):: Hive2 Action should not propagate oozie.hive2.* properties to Beeline.
- [OOZIE-2376](#):: Default action configs not honored if no <configuration> section in workflow.
- [OOZIE-2365](#):: oozie fail to start when smtp password not set.
- [OOZIE-2360](#):: Spark Action fails due to missing mesos jar.
- [OOZIE-2345](#):: Parallel job submission for forked actions.
- [OOZIE-2358](#):: Coord rerun cleanup should reuse heat connections.
- [OOZIE-2356](#):: Add a way to enable/disable credentials in a workflow.
- [OOZIE-2355](#):: Hive2 Action doesn't pass along oozie configs to jobconf.
- [OOZIE-2318](#):: Provide better solution for specifying SSL truststore to Oozie Client.
- [OOZIE-2344](#):: Enabling 'oozie.action.jobinfo.enable' doesn't inject the job information into the map/reduce job's configuration..
- [OOZIE-2350](#):: Package changes for release.
- [OOZIE-2347](#):: Remove unnecessary new Configuration./new jobConf. calls from oozie.
- [OOZIE-2348](#):: Recovery service keeps on recovering coord action of suspended jobs.
- [OOZIE-2277](#):: Honor oozie.action.sharelib.for.spark in Spark jobs.
- [OOZIE-2322](#):: Oozie Web UI doesn't work with Kerberos in Internet Explorer 10 or 11 and curl .
- [OOZIE-2343](#):: Shell Action should take Oozie Action config and setup HADOOP_CONF_DIR.
- [OOZIE-2245](#):: Service to periodically check database schema.
- [OOZIE-2332](#):: Add ability to provide Hive and Hive 2 Action queries inline in workflows .
- [OOZIE-2329](#):: Make handling yarn restarts configurable.
- [OOZIE-2228](#):: Statustransit service doesn't pick bundle with suspend status.
- [OOZIE-2325](#):: Shell action fails if user overrides oozie.launcher.mapreduce.map.env.
- [OOZIE-2324](#):: A syntax error in the kill node causes the workflow to get stuck and other problems.
- [OOZIE-2309](#):: Enable the coord:dateOffset. function in /coordinator-app/datasets/dataset/@initial-instance.
- [OOZIE-2305](#):: Compile Oozie with Hive-1.2.0.
- [OOZIE-2320](#):: TestZKXLogStreamingService.testStreamingWithMultipleOozieServers_coordActionList is failing.
- [OOZIE-2293](#):: Oozie 4.0.1 build failed while building Catalog.
- [OOZIE-2308](#):: Add support for bundle:conf. function.
- [OOZIE-2315](#):: TestOozieCLI.testshareLibUpdate_withSecurity fails with Hadoop 2.
- [OOZIE-2291](#):: Hive2 workflow.xml.security should have "cred" in action tag instead of "hive2" tag.
- [OOZIE-2289](#):: hive-jdbc dependency in core/pom.xml should be compile.
- [OOZIE-2300](#):: TestAuthFilterAuthOozieClient.testClientAuthTokenCache fails with Hadoop 2.7.0 and later.

- [OOZIE-1837](#):: LauncherMainHadoopUtils sensitive to clock skew.
- [OOZIE-2187](#):: Add a way to specify a default JT/RM and NN.
- [OOZIE-2272](#):: Use Hadoop's CredentialProvider for passwords in oozie-site.
- [OOZIE-2287](#):: Add support for deleting hcat partitions in fs action delete.
- [OOZIE-2285](#):: Change in concurrency should trigger coord action ready command.
- [OOZIE-2284](#):: HBaseCredentials should only add hbase-default.xml and hbase-site.xml to actionConf.
- [OOZIE-2286](#):: Update Log4j and Log4j-extras to latest 1.2.x release.
- [OOZIE-2250](#):: Show log for WAITING and TIMEDOUT coord actions.
- [OOZIE-2262](#):: Fix log streaming from other server with start/end filter.
- [OOZIE-2159](#):: 'oozie validate' command should be moved server-side.
- [OOZIE-2271](#):: Upgrade Tomcat to 6.0.44.
- [OOZIE-2266](#):: Fix 'total' actions returned in coordinator job info.
- [OOZIE-2264](#):: Fix coord:offset. to resolve correct data set.
- [OOZIE-2178](#):: fix javadoc to compile on JDK8.
- [OOZIE-2268](#):: Update ActiveMQ version for security and other fixes.
- [OOZIE-2215](#):: Support glob in FS EL function.
- [OOZIE-2270](#):: Audit log doesn't get loaded when "audit log" tab is clicked for coord job.
- [OOZIE-2257](#):: BundleStatusTransitXCommand doesn't update last modified time.
- [OOZIE-2260](#):: Only set "java.io.tmpdir" to "./tmp" for the AM.

Phoenix

This release provides Phoenix 5.0.0 and the following additional Apache patches.

- [PHOENIX-5388](#): Incorrect current_date()/now() when query involves subquery
- [PHOENIX-5060](#): DeleteFamily cells are getting skipped while building Index data after HBASE-21158
- [PHOENIX-4822](#): Ensure the provided timezone is used client-side
- [PHOENIX-5070](#): NPE when upgrading Phoenix 4.13.0 to Phoenix 4.14.1 with hbase-1.x branch in secure setup
- [PHOENIX-5178](#): SYSTEM schema is not getting cached at MetaData server
- [PHOENIX-4872](#): BulkLoad has bug when loading on single-cell-array-with-offsets table.
- [PHOENIX-3067](#): Phoenix metrics system should not be started in mini-cluster mode
- [PHOENIX-3499](#): Enable null value for quote character for CSVBulkLoad tool
- [PHOENIX-3633](#): Null pointer exception when subquery for not exists returns empty result set
- [PHOENIX-3662](#): PhoenixStorageHandler throws ClassCastException
- [PHOENIX-3710](#): Cannot use lowername data table name with indextool
- [PHOENIX-4983](#): Allow using a connection with a SCN set to write data to tables EXCEPT transactional tables or mutable tables with indexes or tables with ROW_TIMESTAMP column.
- [PHOENIX-5309](#): Skip adding log4j and slf4j to phoenix-hive jar to avoid logging in hive-server2.err

HDP 3.0.0, HDP 3.0.1, and HDP 3.1.0 provided Phoenix 5.0.0 and no additional Apache patches.

Pig

This release provides Pig 0.16.0 and the listed Apache patches.

- [PIG-5175](#): Upgrade jruby to 1.7.26
- [PIG-5159](#): Fix Pig not saving grunt history
- [PIG-5287](#): bump jython to 2.7.1

HDP 3.0.0, HDP 3.0.1, HDP 3.1.0 provided Pig 0.16.0 and the following Apache patches:

- [PIG-4947](#): LOAD with HBaseStorage using a mix of pure wildcards and prefixed wildcards results in empty maps for the pure wildcarded column families.

- [PIG-4948](#) : Pig on Tez AM use too much memory on a small cluster.
- [PIG-4921](#): Kill running jobs on InterruptedException (partial).
- [PIG-4916](#): Pig on Tez fail to remove temporary HDFS files in some cases.
- [PIG-5287](#): Bump jython to 2.7.1.

Ranger

This release provides Ranger 1.2.0 and the listed Apache patches.

- [RANGER-2521](#): Masking policies not picked from the zone of the accessed resource
- [RANGER-2513](#): Unable to delete user if he has references in new ref tables
- [RANGER-2511](#): default tag based service is getting created for the tag based service
- [RANGER-2509](#): Add validation message for Importing non JSON file on import action
- [RANGER-2481](#): Create a tag service when a resource service is created and link it to resource service, Parts 1 and 2.
- [RANGER-2496](#): Upgrade Spring Security version in Ranger
- [RANGER-2466](#): Improvement in setting cluster Name in RangerAccessRequest
- [RANGER-2479](#): Change test connection preferred SQL statement for Oracle DB Flavor
- [RANGER-2475](#): Replacing bootstrap accordion with jquery SlideToggle.
- [RANGER-2474](#):Policy version and details in access audits wrong when deny condition added to policy
- [RANGER-2473](#):Upgrade of Ranger Admin to the current version fails in PatchForAtlasResourceAndAccessTypeUpdate_J10016
- [RANGER-2468](#): Upgrade jQuery version in Ranger.
- [RANGER-2458](#): Cluster property name changes in Ranger Plugin code
- [RANGER-2465](#):Create a PolicyCondition to apply if all given tags are present for the accessed resource-addon-sampleTagDef
- [RANGER-2469](#): Adding absolute-ordering to web.xml to fix More than one fragment error during Ranger start
- [RANGER-2465](#):Create a PolicyCondition to apply if all given tags are present for the accessed resource
- [RANGER-2464](#): Upgrade spring, zookeeper, c3p0, jackson-databind, tomcat libraries
- [RANGER-2463](#):Ranger admin authorization audits fails intermittently to fetch from Solr
- [RANGER-2394](#): Add ability to search for multiple users, and exclude multiple users from audit logs
- [RANGER-2455](#): When service created inside a zone landing page that service gets created in unzoned landing page.
- [RANGER-2454](#): Remove the trailing slash in Ranger URL in RangerAdminJersey2RESTClient
- [RANGER-2451](#): fix for ClassCastException in RangerPolicyEngineImpl.convertFromSingleResource()
- [RANGER-2453](#): Tag data-masking policy should allow only one tag as resource
- [RANGER-2456R](#):Upgrade of Ranger Admin to the current version fails in PatchForKafkaServiceDefUpdate_J10025
- [RANGER-2437](#):Update grant/revoke error message to provide more information about the principal type
- [RANGER-2449](#): If service part of zone is not present then null pointer exception is thrown
- [RANGER-2420](#): Ranger spends 36% of CPU in ObjectMapper
- [RANGER-2445](#): Import of TagService for zone
- [RANGER-2446](#): Include security zone details as part of admin audit for policy update
- [RANGER-2435](#): The breadcrumb for Security zone name on policy creation/policy edit page should be stagnant.
- [RANGER-2421](#): add missing httpclient libraries to the atlas plugin
- [RANGER-2436](#): Custom condition: Access from cluster
- [RANGER-2434](#): Remove dependency from com.google.common.base.Objects
- [RANGER-2430](#) : Zoneadmin User is able to create policy for those services which is not associated to zone
- [RANGER-2439](#): Unable to view policy details from access audits when policy has policy condition at policy level
- [RANGER-2435](#): Add support for sticky breadcrumbs.
- [RANGER-2438](#): Legacy PublicAPI REST API to get all policies fails

- [RANGER-2417](#): Set Atlas Entity owner to RangerAccessResource ownerUser attribute for Atlas Ranger Plugin
- [RANGER-2427](#): Tag policies are not evaluated if no security zones are configured
- [RANGER-2423](#): Ranger KnoxSSO authentication in Ranger HA environment

HDP 3.1.0 provided Ranger 1.2.0 and the listed Apache patches:

- [RANGER-2292](#): Test case fix for RANGER-2276
- [RANGER-2244](#): Tomcat Security Vulnerability Alert. The version of the tomcat for ranger should upgrade to 7.0.91 or later.
- [RANGER-2288](#): Sqoop repository config missing 'Common Name for Certificate'
- [RANGER-2284](#): Unable to build image using docker
- [RANGER-2276](#): Email Address should be verified when Add New User in Ranger Admin
- [RANGER-2267](#): Add a icon to differentiate the status of the service
- [RANGER-2277](#): Kylin repository config missing 'Common Name for Certificate'
- [RANGER-2280](#): The emptyText of User Sync and Plugin Status should be reasonable
- [RANGER-2286](#): Ranger install may be prevented by leftover DB entry
- [RANGER-2283](#): User is getting total count of groups even if he is assigned to one group due to which pagination is breaking
- [RANGER-2279](#): Reduce the time spent changing passwords during Ranger Admin install
- [RANGER-2278](#): Unable to delete user if he has references in new ref tables
- [RANGER-2248](#): Sorting does not work in AbstractPredicateUtil.java
- [RANGER-2257](#): Add policyID to error message when click the Access log of Audit
- [RANGER-2258](#): Improve the policy list page to prompt users when the service is disabled
- [RANGER-2264](#): Kafka default policies for new resources are not showing up in UI when upgrade is done from older version
- [RANGER-2251](#): Need to provide options for making java heap size memory configurable in Ranger services
- [RANGER-2259](#): Need to provide appropriate permissions for unix-auth files.
- [RANGER-2252](#): Permission Kafka Admin should not be part of Topic resource in Ranger Kafka resource definition
- [RANGER-2209](#): Added Ranger service definition for ABFS
- [RANGER-2220](#): minify JavaScript files during packaging
- [RANGER-2238](#): String comparison should not use '==' in ServiceUtil.java
- [RANGER-2249](#): Ranger Audit not flushed immediately to hdfs
- [RANGER-2235](#): Modify the login session detail page as a modal
- [RANGER-2241](#): Removed mds and sha1 generation
- [RANGER-2222](#): Apache RangerKafkaPlugin support to handle Kafka Cluster as a new resource
- [RANGER-2210](#): Ranger support for Apache Kafka 2.0.0

HDP 3.0.0 and HDP 3.0.1 provided Ranger 1.1.0 and the following Apache patches.

- [RANGER-1715](#): Enhance Ranger Hive Plugin to support authorization on Hive replication Tasks
- [RANGER-1715](#): Enhance Ranger Hive Plugin to support authorization on Hive replication Tasks - repladmin permission
- [RANGER-1841](#): Audit log record for 'use dbName' hive command contains large number of tags
- [RANGER-1828](#): Good coding practice-add additional headers in ranger
- [RANGER-1833](#): Update Ranger to use 0.8.1 Atlas version
- [RANGER-1825](#): BUG-89494 : Ranger Tagsync start is failing
- [RANGER-91119](#): Exclude jackson jaxrs library from ranger-admin packaging
- [RANGER-1886](#): TagSync updates for changes in Atlas API
- [RANGER-1886](#): TagSync updates for changes in Atlas API - Part 2
- [RANGER-1883](#): Part 2 - All tag sources and sinks should use the same kerberos identity and reuse it
- [RANGER-1897](#): TagSync should replace use of V1 Atlas APIs with V2 APIs for efficient tag-download from Atlas

- [RANGER-1895](#): Simplify Storm dependencies
- [RANGER-1797](#): Tomcat Security Vulnerability Alert. The version of the tomcat for ranger should upgrade to 7.0.82.
- [RANGER-1906](#): Simplify Atlas plugin dependency management
- [RANGER-1898](#): Simplify Knox plugin dependency management
- [RANGER-1893](#): Update HBase dependency to 1.2.6
- [RANGER-1922](#): Simplify HBase dependencies
- [RANGER-1926](#): Upgrade Storm dependency to 1.1.1
- [RANGER-1905](#): Fix assembly creation - it's a black magic. As ranger-plugins-audit doesn't exclude guava from hadoop common, an older guava is leaked into embeddedwebserver classpath, which somehow filtered out by the assembly plugin Fix the atlas plugin assembly, so the guava is actually coming from the 'atlas-plugin's dependency, and not from the ranger-plugins-audit's dependency. And revert "RANGER-1858:ranger-admin start fail when policymgr_http_enabled equal to false."
- [RANGER-1984](#): Hbase audit log records may not show all tags associated with accessed column
- [RANGER-1999](#): Ranger policy engine updates to support list-of-values in access resource
- [RANGER-2000](#): Policy effective dates to support time-bound and temporary authorization
- [RANGER-1987](#): Upgrade to Storm 1.2.0 in Ranger
- [RANGER-2019](#): Handle upgrade scenario to rename the old ATLAS service def and use the new service def
- [RANGER-2019](#): Handle upgrade scenario to rename the old ATLAS service def and use the new service def
- [RANGER-2027](#): Evaluate grantor's group membership in the plugin for grant/revoke request
- [RANGER-2026](#): Update Hbase plugin to handle default namespace
- [RANGER-2018](#): Upgrade to Spring 4
- [RANGER-1994](#): Tomcat Security Vulnerability Alert. The version of the tomcat for ranger should upgrade to 7.0.85.
- [RANGER-2035](#): fix handling of null implClass in servicedefs (Oracle returns null for stored empty strings)
- [RANGER-2043](#): Ranger KMS KeyProvider and HSM KeyProvider should have more debug logs
- [RANGER-2115](#): Atlas service def creation failing in Ranger manual installation

Spark

This release provides Spark 2.3.2 and the listed Apache patches.

- [SPARK-21783](#): Turn on `native` ORC impl and PPD by default
- [SPARK-23456](#): Turn on `native` ORC impl and PPD by default
- [SPARK-23228](#): Add Python Created jsparkSession to JVM's defaultSession
- [SPARK-23510](#): Support Hive 2.2 and Hive 2.3 metastore
- [SPARK-23518](#): Avoid metastore access when the users only want to read and write data frames
- [SPARK-23635](#): Spark executor env variable is overwritten by same name AM env variable
- [SPARK-23787](#): Fix file download test in SparkSubmitSuite for Hadoop 2.9.
- [SPARK-23355](#): convertMetastore should not ignore table properties
- [SPARK-24110](#): Avoid UGI.loginUserFromKeytab in STS
- [SPARK-24149](#): Retrieve all federated namespaces tokens
- [SPARK-24209](#): Automatic retrieve proxyBase from Knox headers
- [SPARK-24312](#): Upgrade to 2.3.3 for Hive Metastore Client 2.3
- [SPARK-24377](#): make --py-files work in non pyspark application
- [SPARK-24479](#): Added config for registering streamingQueryListeners
- [SPARK-24518](#): Using Hadoop credential provider API to store password
- [SPARK-24660](#): Show correct error pages when downloading logs in SHS
- [SPARK-23654](#): remove jets3t as a dependency of spark
- [SPARK-25126](#): Avoid creating Reader for all orc files
- [SPARK-23679](#): Setting RM_HA_URLS for AmIpFilter to avoid redirect failure in YARN mode

- [SPARK-25306](#): Avoid skewed filter trees to speed up `createFilter` in ORC

HDP 3.0.0 and HDP 3.0.1 provided Spark 2.3.1 and the following Apache patches.

- [SPARK-24495](#): SortMergeJoin with duplicate keys wrong results.
- [SPARK-207](#): Remove hardcoded FS scheme from Spark archive.

(Backport from 2.3.2)

- [SPARK-24455](#): fix typo in TaskSchedulerImpl comment.
- [SPARK-24369](#): Correct handling for multiple distinct aggregations having the same argument set.
- [SPARK-24468](#): Handle negative scale when adjusting precision for decimal operations.
- [SPARK-23732](#): Fix source links in generated scaladoc.
- [SPARK-24502](#): flaky test: UnsafeRowSerializerSuite.
- [SPARK-24531](#): Remove version 2.2.0 from testing versions in HiveExternalCatalogVersionsSuite.
- [SPARK-24506](#): Add UI filters to tabs added after binding.
- [SPARK-23754](#): Move UDF stop iteration wrapping from driver to executor.
- Remove unnecessary sort in UnsafeInMemorySorterSuite.
- Fix typo in serializer exception.
- Revert [SPARK-21743](#) top-most limit should not cause memory leak.
- [SPARK-24531](#): Replace 2.3.0 version with 2.3.1.

(Backport from 2.4)

- [SPARK-21783](#): Turn on `native` ORC impl and PPD by default.
- [SPARK-23456](#): Turn on `native` ORC impl and PPD by default.
- [SPARK-23228](#): Add Python Created jsparkSession to JVM's defaultSession.
- [SPARK-23510](#): Support Hive 2.2 and Hive 2.3 metastore.
- [SPARK-23518](#): Avoid metastore access when the users only want to read and write data frames.
- [SPARK-23635](#): Spark executor env variable is overwritten by same name AM env variable.
- [SPARK-23787](#): Fix file download test in SparkSubmitSuite for Hadoop 2.9..
- [SPARK-23355](#) convertMetastore should not ignore table properties.
- [SPARK-24110](#) Avoid UGI.loginUserFromKeytab in STS.
- [SPARK-24149](#): Retrieve all federated namespaces tokens.
- [SPARK-24209](#): Automatic retrieve proxyBase from Knox headers.
- [SPARK-24312](#): Upgrade to 2.3.3 for Hive Metastore Client 2.3.
- [SPARK-24377](#): make --py-files work in non pyspark application.
- [SPARK-24479](#): Added config for registering streamingQueryListeners.
- [SPARK-24518](#): Using Hadoop credential provider API to store password.
- [SPARK-24660](#): Show correct error pages when downloading logs in SHS.

Sqoop

This release provides Sqoop 1.4.7 and no additional Apache patches.

HDP 3.0.0 and HDP 3.0.1 provided Sqoop 1.4.7 and no additional Apache patches.

Storm

This release provides Storm 1.2.1 and the following Apache patches.

- [STORM-3123](#): add support for Kafka security config in storm-kafka-monitor
- [STORM-3046](#): Ensure KafkaTridentSpoutEmitter handles empty batches correctly
- [STORM-3026](#): Upgrade ZK instance for security
- [STORM-3052](#): Allow for blobs to be unzipped/untarred

- [STORM-3110](#): Skip the user while checking isProcessAlive
- [STORM-2911](#): add serialVersionUID to storm-kafka SpoutConfig
- [STORM-2896](#): Add tool to help users migrate offsets from storm-kafka to storm-kafka-client
- [STORM-2952](#): Deprecate storm-kafka
- [STORM-2946](#): Upgrade to HBase 2.0
- [STORM-2960](#): Stress importance of setting up proper OS account for Storm processes

HDP 3.0.0, HDP 3.0.1, and HDP 3.1.0 provided Storm 1.2.1 and the following Apache patches:

- [STORM-1839](#): Kinesis Spout
- [STORM-1694](#): Kafka Spout Trident Implementation Using New Kafka Consumer API
- [STORM-1960](#): Add CORS support to STORM UI Rest api
- [STORM-1959](#): Add missing license header to KafkaPartitionOffsetLag
- [STORM-1950](#): Change response json of "Topology Lag" REST API to keyed by spoutId, topic, partition.
- [STORM-1833](#): Simple equi-join in storm-sql standalone mode
- [STORM-1866](#): Update Resource Aware Scheduler Documentation
- [STORM-1930](#): Kafka New Client API - Support for Topic Wildcards
- [STORM-1924](#): Adding conf options for Persistent Word Count Topology
- [STORM-1956](#): Disabling Backpressure by default
- [STORM-1934](#): Fix race condition between sync-supervisor and sync-processes
- [STORM-1919](#): Introduce FilterBolt on storm-redis
- [STORM-1742](#): More accurate 'complete latency'
- [STORM-1945](#): Fix NPE bugs on topology spout lag for storm-kafka-monitor
- [STORM-1888](#): add description for shell command
- [STORM-1902](#): add a simple & flexible FileNameFormat for storm-hdfs
- [STORM-1914](#): Storm Kafka Field Topic Selector
- [STORM-1907](#): PartitionedTridentSpoutExecutor has incompatible types that cause ClassCastException
- [STORM-1925](#): Remove Nimbus thrift call from Nimbus itself
- [STORM-1909](#): Update HDFS spout documentation
- [STORM-1136](#): Command line module to return kafka spout offsets lag and display in storm ui
- [STORM-1911](#): IClusterMetricsConsumer should use seconds to timestamp unit
- [STORM-1849](#): HDFSFileTopology should use the 3rd argument as topologyName
- [STORM-1906](#): Window count/length of zero should be disallowed
- [STORM-1893](#): Support OpenTSDB for storing timeseries data.
- [STORM-1723](#): Introduce ClusterMetricsConsumer
- [STORM-1700](#): Introduce 'whitelist' / 'blacklist' option to MetricsConsumer
- [STORM-1698](#): Asynchronous MetricsConsumerBolt
- [STORM-1705](#): Cap number of retries for a failed message
- [STORM-1884](#): Prioritize pendingPrepare over pendingCommit
- [STORM-1575](#): fix TwitterSampleSpout NPE on close
- [STORM-1874](#): Update logger private permissions
- [STORM-1865](#): update command line client document
- [STORM-1771](#): HiveState should flushAndClose before closing old or idle Hive connections
- [STORM-1882](#): Expose TextFileReader public
- [STORM-1873](#): Implement alternative behaviour for late tuples
- [STORM-1719](#): Introduce REST API: Topology metric stats for stream
- [STORM-1887](#): Fixed help message for set_log_level command
- [STORM-1878](#): Flux can now handle IStatefulBolts
- [STORM-1864](#): StormSubmitter should throw respective exceptions and log respective errors for registered submitter hook invocation
- [STORM-1868](#): Modify TridentKafkaWordCount to run in distributed mode

- [STORM-1859](#): Ack late tuples in windowed mode
- [STORM-1851](#): Fix default nimbus impersonation authorizer config
- [STORM-1848](#): Make KafkaMessageId and Partition serializable to support
- [STORM-1862](#): Flux ShellSpout and ShellBolt can't emit to named streams
- [STORM-1728](#): TransactionalTridentKafkaSpout error
- [STORM-1850](#): State Checkpointing Documentation update
- [STORM-1674](#): Idle KafkaSpout consumes more bandwidth than needed
- [STORM-1842](#): Forward references in storm.thrift cause tooling issues
- [STORM-1730](#): LocalCluster#shutdown() does not terminate all storm threads/thread pools.
- [STORM-1841](#): Address a few minor issues in windowing and doc
- [STORM-1709](#): Added group by support in storm sql standalone mode
- [STORM-1720](#): Support GEO in storm-redis

Tez

This release provides Tez 0.9.1 and the listed Apache patches.

- [TEZ-3976](#): Batch ShuffleManager error report events
- [TEZ-3984](#): Shuffle: Out of Band DME event sending causes errors
- [TEZ-3980](#): ShuffleRunner: the wake loop needs to check for shutdown
- [TEZ-3958](#): Add internal vertex priority information into the tez dag.dot debug information
- [TEZ-3978](#): DAGClientServer Socket exception when localhost name lookup failures
- [TEZ-3974](#): Correctness regression of TEZ-955 in TEZ-2937
- [TEZ-3934](#): LegacySpeculator sometime issues wrong number of speculative attempts
- [TEZ-3942](#): RPC getTask writable optimization invalid in hadoop 2.8+
- [TEZ-3965](#): TestMROOutput: Fix the hard-coded /tmp/output paths
- [TEZ-3916](#): Add hadoop-azure-datalake jar to azure profile
- [TEZ-3970](#): NullPointerException in Tez ShuffleHandler Ranged Fetch
- [TEZ-3964](#): Inflater not closed in some places
- [TEZ-3912](#): Fetchers should be more robust to corrupted inputs
- [TEZ-3963](#): Possible InflaterInputStream leaked in TezCommonUtils and related classes
- [TEZ-3954](#): Reduce Tez Shuffle Handler Memory needs for holding TezIndexRecords
- [TEZ-3960](#): Better error handling in proto history logger and add doAs support.
- [TEZ-3962](#): Configuration decode leaks an Inflater object
- [TEZ-3953](#): Restore API-compat for DAGClient for TEZ-3951
- [TEZ-3951](#): TezClient wait too long for the DAGClient for prewarm; tries to shut down the wrong DAG
- [TEZ-3944](#): TestTaskScheduler times-out on Hadoop3
- [TEZ-3938](#): Task attempts failing due to not making progress
- [TEZ-3949](#): TestATSHistoryV15 is failing with hadoop3+
- [TEZ-3946](#): NoClassDefFoundError, org.apache.hadoop.mapred.ShuffleHandler.
- [TEZ-3929](#): Upgrade Jersey to 1.19
- [TEZ-3947](#): TestATSHistoryWithACLs fails with Hadoop3 and Jersey 1.19
- [TEZ-3943](#): TezClient leaks DAGClient for prewarm
- [TEZ-3939](#): Remove performance hit of precondition check in AM for register running task attempt
- [TEZ-3940](#): Reduce time to convert TaskFinishedEvent to string
- [TEZ-3902](#): Upgrade to netty-3.10.5.Final.jar
- [TEZ-3935](#): DAG aware scheduler should release unassigned new containers rather than hold them
- [TEZ-3937](#): Empty partition BitSet to byte[] conversion creates one extra byte in rounding error
- [TEZ-3824](#): MRCombiner creates new JobConf copy per spill
- [TEZ-3933](#): Remove sleep from test TestExceptionPropagation

- [TEZ-3911](#): Optional min/max/avg aggr. task counters reported to HistoryLoggingService at final counter aggr
- [TEZ-3932](#): TaskSchedulerManager can throw NullPointerException during DAGAppMaster container cleanup race
- [TEZ-3931](#): TestExternalTezServices fails on Hadoop3
- [TEZ-3930](#): TestDagAwareYarnTaskScheduler fails on Hadoop 3
- [TEZ-3927](#): TestReduceProcessor fails on Hadoop 3.x
- [TEZ-3924](#): TestDefaultSorter fails intermittently due random keys and interaction with RLE and partition collisions
- [TEZ-3914](#): Recovering a large DAG fails to size limit exceeded
- [TEZ-3887](#): Tez Shuffle Handler should support Index Cache configuration
- [TEZ-3914](#): Recovering a large DAG fails to size limit exceeded
- [TEZ-3817](#): DAGs can hang after more than one uncaught Exception during doTransition.
- [TEZ-3915](#): Create protobuf based history event logger.
- [TEZ-3909](#): DAG can hang if vertex with no tasks is killed
- [TEZ-3907](#): Improve log message to include the location the writers decide to spill output
- [TEZ-3874](#): NPE in TezClientUtils when "yarn.resourcemanager.zk-address" is present in Configuration.
- [TTEZ-3892](#): getClient API for TezClient
- [TEZ-3897](#): Tez Local Mode hang for vertices with broadcast input.
- [TEZ-3888](#): Update Jetty to org.eclipse.jetty 9.x
- [TEZ-3898](#): TestTezCommonUtils fails when compiled against hadoop version >= 2.8
- [TEZ-3896](#): TestATSV15HistoryLoggingService#testNonSessionDomains is failing
- [TEZ-3893](#): Tez Local Mode can hang for cases.
- [TEZ-3894](#): Tez intermediate outputs implicitly rely on permissive umask for shuffle
- [TEZ-3895](#): Missing name for local mode task scheduler service async request handler thread
- [TEZ-3770](#): DAG-aware YARN task scheduler
- [TEZ-3880](#): Do not count rejected tasks as killed in vertex progress
- [TEZ-3877](#): Delete unordered spill files once merge is done

HDP 3.0.0 and HDP 3.0.1 provided Tez 0.9.1 and no additional Apache patches.

Zeppelin

This release provides Zeppelin 0.8.0 and no additional Apache patches.

HDP 3.0.0 and HDP 3.0.1 provided Zeppelin 0.8.0 and no additional Apache patches.

ZooKeeper

This release provides ZooKeeper 3.4.6 and the following patches.

- [ZOOKEEPER-2498](#): Potential resource leak in C client when processing unexpected / out of order response
- [ZOOKEEPER-2693](#): DOS attack on wchp/wchc four letter words (4lw)
- [ZOOKEEPER-1045](#): Support Quorum Peer mutual authentication via SASL
- [ZOOKEEPER-2726](#): Patch for ZOOKEEPER-2693 introduces potential race condition
- [ZOOKEEPER-1392](#): Should not allow to read ACL when not authorized to read node

HDP 3.1.0 provides ZooKeeper 3.4.6 and the patches listed for HDP 3.0.0 and HDP 3.0.1

HDP 3.0.1 provided ZooKeeper 3.4.6 with the following Apache patches:

- [ZOOKEEPER-2662](#): Export a metric for txn log sync times
- [ZOOKEEPER-3019](#): Add a metric to track number of slow fsyncs

HDP 3.0.0 provided ZooKeeper 3.4.6 and the following Apache patches:

- [ZOOKEEPER-2141](#): ACL cache in DataTree never removes entries
- [ZOOKEEPER-1897](#): ZK Shell/Cli not processing commands (stack via michim)
- [ZOOKEEPER-1926](#): Unit tests should only use build/test/data for data
- [ZOOKEEPER-2064](#): Prevent resource leak in various classes
- [ZOOKEEPER-2053](#): Make scripts use ZOOKEEPER_HOME if set.
- [ZOOKEEPER-1876](#): Add support for installing windows services in .cmd scripts
- [ZOOKEEPER-1848](#): [WINDOWS] Java NIO socket channels does not work with Windows ipv6 on JDK6
- [ZOOKEEPER-1574](#): Set svn:eol-style property to native.
- [ZOOKEEPER-1575](#): adding .gitattributes to prevent CRLF and LF mismatches for source and text files (Raja Aluri via michim)
- [ZOOKEEPER-1888](#): ZkCli.cmd commands fail with "'java' is not recognized as an internal or external command" (Ivan Mitic via michim)
- [ZOOKEEPER-1791](#): ZooKeeper package includes unnecessary jars that are part of the package. (mahadev)
- [ZOOKEEPER-1506](#): Re-try DNS hostname -> IP resolution if node connection fails
- [ZOOKEEPER-1952](#): Default log directory and file name can be changed
- [ZOOKEEPER-706](#): large numbers of watches can cause session re-establishment to fail (Chris Thunes)
- [ZOOKEEPER-1576](#): Zookeeper cluster - failed to connect to cluster if one of the provided IPs causes java.net.UnknownHostException (Edward Ribeiro via camille)

Fixed Common Vulnerabilities and Exposures

Common Vulnerabilities and Exposures (CVE) that are addressed in this release.

CVE-2018-11770

Component: Spark

Summary: Update Rest Server docs defaults.

Severity: Important

Vendor: Hortonworks

Versions Affected: HDP versions before 3.1.4.0

Users Affected: Users who use Apache Spark.

Impact: See [BUG-120486](#) and [SPARK-25088](#). From version 1.3.0, Apache Spark's standalone master exposes a REST API for job submission. However, the REST API does not use any authentication mechanism, and this is not adequately documented.

Recommended Action: Upgrade to HDP 3.1.4+

CVE-2016-5017

Component: ZooKeeper

Summary: Buffer overflow in the C CLI shell in Apache Zookeeper.

Severity: Low

Vendor: Hortonworks

Versions Affected: HDP versions before 3.1.4.0

Users Affected: Users who use Apache ZooKeeper with C client.

Impact: See [BUG-120252](#) and [ZOOKEEPER-2498](#). Buffer overflow in the C cli shell in Apache Zookeeper before 3.4.9 and 3.5.x before 3.5.3, when using the "cmd:" batch mode syntax, allows attackers to have unspecified impact via a long command string.

Recommended Action: Upgrade to HDP 3.1.4+

CVE-2017-5637

Component: ZooKeeper

Summary: Two four letter word commands "wchp/wchc" are CPU intensive and could cause spike of CPU utilization on Apache ZooKeeper server if abused.

Severity: Low

Vendor: Hortonworks

Versions Affected: HDP versions before 3.1.4.0

Users Affected: Users who use Apache ZooKeeper in public infrastructure.

Impact: See BUG-120252 and ZOOKEEPER-2693. Two four letter word commands "wchp/wchc" are CPU intensive and could cause spike of CPU utilization on Apache ZooKeeper server if abused.

Recommended Action: Upgrade to HDP 3.1.4+

CVE-2018-8012

Component: ZooKeeper

Summary: No authentication/authorization is enforced when a server attempts to join a quorum in Apache ZooKeeper

Severity: Important

Vendor: Hortonworks

Versions Affected: HDP versions before 3.1.4.0

Users Affected: Users who use Apache ZooKeeper.

Impact: See BUG-120252 and ZOOKEEPER-1045, ZOOKEEPER-2726. No authentication/authorization is enforced when a server attempts to join a quorum in Apache ZooKeeper.

Recommended Action: Upgrade to HDP 3.1.4+

CVE-2019-0201

Component: ZooKeeper

Summary: ZooKeeper's getACL() command doesn't check any permission when retrieves the ACLs

Severity: Important

Vendor: Hortonworks

Versions Affected: HDP versions before 3.1.4.0

Users Affected: Users who use Apache ZooKeeper.

Impact: See BUG-120252 and ZOOKEEPER-1392. ZooKeeper's getACL() command doesn't check any permission when retrieves the ACLs of the requested node and returns all information contained in the ACL Id field as plaintext string.

Recommended Action: Upgrade to HDP 3.1.4+

Related Concepts

[Component Versions](#)

Fixed Issues

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

Incorrect-results

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|---|------------------|--|
| BUG-121291 | HIVE-21746 | Hive | ArrayIndexOutOfBoundsException during dynamically partitioned hash join, with CBO disabled |
| BUG-121285 | HIVE-21862 | Hive | ORC ppd produces wrong result with timestamp |
| BUG-120240 | HIVE-21509 | Hive | LLAP may cache corrupted column vectors and return wrong query result |
| BUG-120048 | HIVE-21291 | Hive | Restore historical way of handling timestamps in Avro while keeping the new semantics at the same time |
| BUG-119841 | HIVE-20274 | Hive | HiveServer2 ObjectInspectorFactory leaks for Struct and List object inspectors |
| BUG-119788 | SPARK-23986 SPARK-24495 SPARK-24589 SPARK-25368 | Spark | Spark integrity and correctness problems with Spark 2.3.0 |
| BUG-119025 | HIVE-21460 | Hive | ACID: Load data followed by a select * query results in incorrect results |
| BUG-118705 | TEZ-3982 | Tez | DAGAppMaster and tasks should not report negative or invalid progress |

Other

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|--|------------------|--|
| BUG-121579 | BUG-121509 | Atlas | Atlas upgrade will likely hit out-of-memory in environments having large amount of data |
| BUG-121516 | BUG-121509 | Atlas | Atlas split2: test_nested_entity_filter_search_OR_of_ORs_and_ANDs |
| BUG-121509 | ATLAS-3033 | Atlas | ATLAS Hive hook fails for CTAS query involving UDFs |
| BUG-121414 | N/A | HBase | Incremental and restore backup fails on AWS cluster |
| BUG-121375 | ATLAS-3045 | Atlas | Notification processing to provide an option to ignore Hive dummy databases and tables |
| BUG-121193 | ATLAS-3045 , ATLAS-3045 , ATLAS-3045 , ATLAS-3045 | Atlas | Incorrect column lineage: each output column has input from *all columns* of the input table |
| BUG-121149 | ATLAS-3333 | Atlas | Notification processing to provide an option to ignore Hive dummy databases and tables |
| BUG-121095 | HBASE-21156 | HBase | Fix HBCK2 to work with Kerberized deployment of HBase |
| BUG-121018 | N/A | HWC | Disable pruning and pushdowns for spark v2.3.x by default |
| BUG-120946 | HIVE-21888 | Hive | Set hive.parquet.timestamp.skip.conversion default to true |
| BUG-120941 | PHOENIX-5309 | Phoenix | No job progress after adding phoenix-hive JAR under AUX JAR PATH |
| BUG-120843 | HBASE-22617 | HBase | Recovered WAL directories not getting cleaned up |
| BUG-120840 | RANGER-2479 | Ranger | Change test connection preferred SQL statement for Oracle database Flavor |

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|--|------------------|--|
| BUG-120817 | ATLAS-3006 | Atlas | Option to ignore/prune metadata for temporary/staging Hive tables |
| BUG-120782 | ATLAS-3218 | Atlas | Atlas user interface - entity detail page columns sorting empty value check |
| BUG-120734 | N/A | HWC | Turn off smart execution by default |
| BUG-120717 | RANGER-2475 | Ranger | Replacing bootstrap accordion with jquery slideToggle. |
| BUG-120706 | HIVE-21827 | Hive | Multiple calls in SemanticAnalyzer do not go through getTableObjectByName method |
| BUG-120705 | HBASE-21284 | HBase | Adopt less stringent compaction throughput defaults |
| BUG-120652 | ATLAS-3034 ATLAS-3043 ATLAS-3051 ATLAS-3053 ATLAS-3055 ATLAS-3056 ATLAS-3063 ATLAS-3065 | Atlas | Kafka message size reduced |
| BUG-120588 | N/A | HWC | Incorrect data when writes happen in parallel |
| BUG-120561 | n/a | Phoenix | Upgrade Phoenix jetty 9.3 |
| BUG-120411 | ATLAS-3218 | Atlas | Atlas user interface - entity detail page columns sorting empty value check |
| BUG-120371 | ATLAS-3129 | Atlas | Fix SSL Truststore reloader leak from AtlasBaseClient |
| BUG-120292 | ATLAS-3129 | Atlas | Fix SSL Truststore reloader leak from AtlasBaseClient |
| BUG-120220 | ATLAS-2634 ATLAS-2891 ATLAS-2975 ATLAS-3006 | Atlas | Improve Atlas NotificationhookConsumer performance |
| BUG-120082 | ATLAS-3193 | Atlas | Hive bootstrap fails with atlas exception |
| BUG-120081 | HIVE-21668 | Hive | Backport HIVE-21668 to 3.1 - Addressing vulnerabilities |
| BUG-120070 | HIVE-21722 | Hive | REPL::END event not printed in the background thread pool |
| BUG-120055 | RANGER-2423 | Ranger | Ranger KnoxSSO authentication in Ranger HA environment |
| BUG-120046 | HIVE-21619 | Hive | Print timestamp type without precision in SQL explain extended |
| BUG-120033 | HIVE-21717 | Hive | Hive bootstrap instance fails when table in source database is updated when bootstrap is running |
| BUG-120032 | ATLAS-3045 | Atlas | Hive replication fails with Atlas Exception |
| BUG-120022 | RANGER-2419 | Ranger | Improve SQL script to skip statements when atlas service def is not supported |
| BUG-119969 | HBASE-20604 | HBase | ProtobufLogReader#readNext can incorrectly loop to the same position in the stream until the the WAL is rolled |
| BUG-119898 | HBASE-21688 , HBASE-22200 | HBase | Backport fixes from separate filesystems for WAL and storefiles |
| BUG-119897 | HBASE-20662 , HBASE-21225 , HBASE-21631 , HBASE-21634 | HBase | Backport quota-related fixes |

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|---|------------------|---|
| BUG-119896 | HBASE-19008 , HBASE-21129 , HBASE-21620 , HBASE-21734 , HBASE-22144 | HBase | Backport of upstream Apache HBase issues around filters |
| BUG-119895 | HBASE-21402 , HBASE-21618 , HBASE-21621 , HBASE-21644 , HBASE-21775 , HBASE-21890 , HBASE-21932 , HBASE-22185 | HBase | Upstream Apache HBase backports |
| BUG-119894 | HBASE-18484 , HBASE-20220 , HBASE-21413 , HBASE-21645 , HBASE-21683 , HBASE-21732 , HBASE-21915 , HBASE-22123 , HBASE-22128 | HBase | Backport of upstream Apache HBase operator tooling and administrative cleanliness |
| BUG-119893 | HBASE-21738 | HBase | Backport of upstream Performance related fixes |
| BUG-119892 | ATLAS-3176 ATLAS-3193 | Atlas | Atlas import fails if the target database is already created |
| BUG-119870 | HIVE-21677 | Hive | Using strict managed tables for ACID table fixes |
| BUG-119848 | HIVE-20291 HIVE-21674 | Hive | Replicate streaming ingestion with transactional batch size as 1. |
| BUG-119819 | HIVE-21651 | Hive | Move protobuf into hive-exec. |
| BUG-119809 | HIVE-21232 , HIVE-21294 , HIVE-21564 , ORC-477 , ORC-491 | Hive2 | Backport of upstream Hive & ORC fixes - HIVE-21232, HIVE-21294, HIVE-21564, ORC-491,ORC-477 |
| BUG-119802 | ATLAS-3177 | Atlas | Regression, Export : changeMarker is not set right - set to older value compared to entity's lastModifiedTime |
| BUG-119795 | ATLAS-3176 | Atlas | Regression , Export : Export zip doesn't contain all entities which are listed in atlas-export-order.json |
| BUG-119793 | RANGER-2445 | Ranger | Import of Tagservice for zone |
| BUG-119792 | N/A | Hive | External table location is not preserved at target when base dir is set as /. |
| BUG-119746 | RANGER-2407 | Ranger | Update default header values sent from Ranger |
| BUG-119733 | ATLAS-3168 | | Atlas fails with NPE(createTableIfNotExists exception) |
| BUG-119693 | N/A | Druid | Druid Historic Nodes going down |
| BUG-119632 | RANGER-2400 , RANGER-2402 | Ranger | Policy name needs to be unique within security zone and service |
| BUG-119546 | STORM-3068 | Storm | STORM_JAR_JVM_OPTS are not passed to storm-kafka-monitor properly |
| BUG-119544 | STORM-3123 | Storm | Storm Kafka Monitor does not work with Kafka over two-way SSL |
| BUG-119535 | RANGER-2404 | Ranger | Issue of delegated admin for zone and unzoned policy associated to same resource |

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|--|------------------|---|
| BUG-119514 | CALCITE-2929 , HIVE-21342 , HIVE-21544 | Calcite | Backport the Calcite issue where analyze compute stats for column leave behind staging dir on HDFS |
| BUG-119484 | ATLAS-3116 | Atlas | Fixes the issue when Atlas is configured to use both Knox SSO and LDAP authentication. LDAP login fails, once a sign-in is made through Knox SSO. |
| BUG-119481 | ATLAS-2842 | Atlas | Entities with different guids does not get referenced in array with SET cardinality |
| BUG-119463 | ATLAS-3182 ATLAS-3188 | Atlas | Hive Replication fails during Atlas export (3.1 - 3.1) |
| BUG-119374 | HBASE-22128 | HBase | Backport Apache HBASE-22128 "Move namespace region then master crashed make deadlock" |
| BUG-119354 | CALCITE-2976 | Hive | Improve materialized view rewriting coverage with disjunctive predicates |
| BUG-119330 | RANGER-2408 | Ranger | Restrict Ranger User's capabilities according to their role |
| BUG-119255 | HBASE-21915 | HBase | FileLink\$FileLinkInputStream doesn't implement CanUnbuffer |
| BUG-119253 | RANGER-2377 | Ranger | Make solr bootstrapping configurable |
| BUG-119251 | RANGER-2392 | Ranger | Description of zone is not available on the UI |
| BUG-119241 | KNOX-1848 , KNOX-1851 | Knox | Remote Alias Discovery Service Not Working |
| BUG-119229 | N/A | HWC | Hive Warehouse Connector is changing External table as Managed table in Overwrite mode |
| BUG-119189 | RANGER-2387 | Ranger | Add public api v2 for security zones |
| BUG-119174 | RANGER-2392 | Ranger | Create / Update zone to have provision to associate Tag based service with zone |
| BUG-119159 | ATLAS-3054 | Atlas | hive_column state is not getting changed to "DELETED" after a column is replaced |
| BUG-119137 | ATLAS-3136 | Atlas | Assigning a term to an invalid entity GUID throws 500, expected is 400. |
| BUG-119120 | N/A | Hive | Backport HIVE-21404 MSSQL upgrade script alters the wrong column |
| BUG-119109 | HIVE-21362 | Hive | HIVE-21362 - Add an input format and serde to read from protobuf files. |
| BUG-119078 | RANGER-2384 | Ranger | Get All Zones API is returning response in raw format,proper response object is required. |
| BUG-119018 | HIVE-21500 | Hive | Disable conversion of managed table to external and vice versa at source. |
| BUG-118999 | CALCITE-2946 | Hive | RelBuilder wrongly skips creation of Aggregate that prunes columns if input produces one row at most |
| BUG-118890 | PHOENIX-5178 | Phoenix | [PHOENIX-5178] - AccessDeniedException while connecting to Phoenix with read access to SYSTEM name space |
| BUG-118858 | KNOX-1830 | Knox | Fix sporadic Illegal character error in TestHashicorpVaultAliasService |
| BUG-118833 | RANGER-2399 | Ranger | User's listing page hits users API call twice from UI |
| BUG-118820 | RANGER-2372 | Ranger | Remove non-existing URL entries from spring config file |
| BUG-118781 | RANGER-2383 | Ranger | Incorrect response when trying to delete user attached to a security zone |

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|--|------------------|--|
| BUG-118779 | HADOOP-16104 | Hadoop Common | WASB tests to downgrade to skip when test a/c is namespace enabled |
| BUG-118772 | HADOOP-16169 | Hadoop Common | ABFS: Bug fix for getPathProperties |
| BUG-118770 | HADOOP-16068 | Hadoop Common | ABFS Authentication and Delegation Token plugins to optionally be bound to specific URI of the store |
| BUG-118764 | HADOOP-16163 | Hadoop Common | NPE in setup/teardown of ITestAbfsDelegationTokens |
| BUG-118763 | HADOOP-15813 | Hadoop Common | Enable more reliable SSL connection reuse |
| BUG-118729 | KNOX-1826 | Knox | Replace zip4j with ShrinkWrap |
| BUG-118672 | HADOOP-16182 | Hadoop Common | Update ABFS storage back-end with "close" flag when application is done writing to a file |
| BUG-118667 | N/A | Knox | Let end-users customize Knox without editing start scripts |
| BUG-118663 | RANGER-2358 | Ranger | Update Ranger libraries version |
| BUG-118644 | HIVE-21435 | Hive | LlapBaseInputFormat should get task number from TASK_ATTEMPT_ID conf if present, while building SubmitWorkRequestProto |
| BUG-118633 | CALCITE-2789 | Calcite | Upgrade jackson-databind to 2.9.8 |
| BUG-118615 | HIVE-21214 | Hive | Backport HIVE-21214 - HiveServer2 split 2 empty bucket tests fail due to missing data from output |
| BUG-118599 | RANGER-2446 | Ranger | Include security zone details as part of admin audit for policy update |
| BUG-118583 | RANGER-2354 | Ranger | Enhance Ranger UI Policy Page to have "Add Policy Condition" |
| BUG-118573 | N/A | Knox | The Knox Gateway truststore should be configurable |
| BUG-118571 | RANGER-2340 | Ranger | Add Policy Version to the Ranger Audit log |
| BUG-118530 | PHOENIX-5070 | Phoenix | Create local index command failed with NPE |
| BUG-118447 | N/A | Hive | Propagate exception from all the APIs using executeUpdate() |
| BUG-118444 | HIVE-21403 | Hive | Error code for failure in case of bootstrap twice should qualify for failed admin case. |
| BUG-118440 | RANGER-2347 | Ranger | Audits Tab : Zone Auditor should see audits only related to his own zones |
| BUG-118439 | RANGER-2355 | Ranger | Reports page: policy listing to have column of Zone name |
| BUG-118428 | KNOX-1787 | Knox | Create PasswordUtils class in gateway-util-common |
| BUG-118389 | KNOX-1801 | Knox | Master secret is incorrectly assumed when a custom truststore is not specified when clientauth is enabled |
| BUG-118292 | KNOX-1793 | Knox | DefaultKeystoreService should not validate the signing key on initialization |
| BUG-118282 | KNOX-1791 | Knox | MasterService should be a field on GatewayServices |
| BUG-118252 | N/A | Ranger | Fix Ranger vulnerabilities |
| BUG-118205 | HIVE-21329 | Hive | Custom Tez runtime unordered output buffer size depending on operator pipeline |
| BUG-118194 | KNOX-1756 | Knox | Knox Gateway TLS Keystore and Alias Should be Configurable |
| BUG-118168 | RANGER-2382 , RANGER-2385 | Ranger | Improvement to Audit page |
| BUG-118090 | KNOX-1687 | Knox | Hashicorp Vault alias credential provider |

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|--|------------------|--|
| BUG-118089 | KNOX-1789 | Knox | Refactor RemoteAliasService to use service loading |
| BUG-118065 | HBASE-21960 | HBase | Backport to-be-disclosed HBase security issue |
| BUG-118042 | ATLAS-3045 , ATLAS-3046 , ATLAS-3048 | Atlas | Fixes an issue where you see duplicate tables and databases in Atlas |
| BUG-118035 | KNOX-1802 | Knox | Plugins API support for ranger is missing in Knox. |
| BUG-117987 | N/A | Storm | Upgrade Apache Thrift Java client library to 0.12.0 due to CVE-2018-1320 |
| BUG-117906 | HIVE-21278 | Hive | Fixes an ambiguity in grammar warnings at compilation time |
| BUG-117880 | PHOENIX-4872 | Phoenix | Inconsistent/missing results during Phoenix CsvBulkLoadTool |
| BUG-117870 | N/A | Ranger | Export policies based on zone. |
| BUG-117860 | HIVE-21529 | Hive | Hive support bootstrap of ACID/MM tables on an existing policy. |
| BUG-117855 | STORM-3046 | Storm | Getting a NPE leading worker to die when starting a topology. |
| BUG-117854 | HIVE-21286 | Hive | Hive should support clean-up of previously bootstrapped tables when retry from different dump. |
| BUG-117807 | HIVE-21281 | Hive | Repl checkpointing doesn't work when retry bootstrap load with partitions of external tables. |
| BUG-117758 | N/A | Knox | Add back Livy Session API after RMP-12915 |
| BUG-117703 | HIVE-21167 | Hive | Bucketing: Bucketing version 1 is incorrectly partitioning data |
| BUG-117639 | HBASE-21200 | HBase | "Memstore flush doesn't finish because of seekToPreviousRow() in memstore scanner." to branch-1 |
| BUG-117634 | HIVE-21188 | Hive | Backport HIVE-21188: Hive query on view fails with SemanticException when there is a Ranger masked policy in place |
| BUG-117588 | RANGER-2351 | Ranger | Implement Import / Export of Policies by Zone |
| BUG-117561 | ATLAS-3050 | Atlas | HDFS replication instance failed intermittently due to Atlas replication failure. |
| BUG-117559 | HADOOP-16041 | Hadoop Common | UserAgent string for ABFS |
| BUG-117424 | RANGER-2333 , RANGER-2339 | Ranger | All Users should be able to view "Security Zone" tab and should get access to security zones. |
| BUG-117420 | RANGER-2339 | Ranger | Hidden users and groups are getting listed in Zone administration section during zone creation. |
| BUG-117377 | HADOOP-16036 | Hadoop Common | WASB: Disable jetty logging configuration announcement |
| BUG-117370 | HADOOP-15876 | Hadoop Common | Use keySet().removeAll() to remove multiple keys from Map in AzureBlobFileSystemStore |
| BUG-117369 | HADOOP-15872 | Hadoop Common | HADOOP-15872: REST version for ADLS Gen 2 |
| BUG-117261 | RANGER-1937 | Ranger | Ranger Tagsync not syncing tags in Atlas |
| BUG-117192 | N/A | Knox | Knox HDP-3.0.100.0 build failure |
| BUG-117159 | RANGER-2234 | Ranger | Cannot add or update a child row,a foreign key constraint fails when installing ranger-admin |
| BUG-117146 | KNOX-1750 | Knox | Knox Admin UI - Unable to view descriptor service params |

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|--|------------------|---|
| BUG-117134 | N/A | HWC | HWC unittests failure - TestHive[Streaming Server2]CredentialProvider |
| BUG-117116 | CALCITE-2388 , CALCITE-2733 | Calcite | Use catalog and schema from JDBC connection string to retrieve tables if specified |
| BUG-117115 | CALCITE-2732 | Calcite | Upgrade Postgresql driver version |
| BUG-117114 | CALCITE-2673 | Calcite | Default implementation of SqlDialect should not support pushing all functions (addendum) |
| BUG-117106 | SPARK-26118 | Spark | Make Jetty's requestHeaderSize configurable in Spark |
| BUG-117104 | KNOX-1705 , KNOX-1707 , KNOX-1708 , KNOX-1731 , KNOX-1735 | Knox | Knox Admin UI improvements |
| BUG-116782 | RANGER-2411 | Ranger | Admin role user is able to create Zone for KMS service with REST API |
| BUG-116737 | HIVE-21236 | Hive | Hive JDBC Storage Handler: TPCDS_GRP_3_88 doesn't output proper results. |
| BUG-116676 | ATLAS-3006 | Atlas | Option to ignore/prune metadata for temporary/staging Hive tables |
| BUG-116638 | KNOX-1339 , KNOX-1728 | Knox | Allow custom parameters to be passed to dispatches |
| BUG-116480 | RANGER-2313 | Ranger | Customer is getting duplicate tables and databases in Atlas |
| BUG-116359 | HIVE-18661 | Hive | CachedStore: Use metastore notification log events to update cache |
| BUG-116358 | N/A | HWC | SparkRHWC: test a set of user scenarios for Zeppelin, Livy, Spark Shell/Submit |
| BUG-116357 | N/A | HWC | SparkR support for HWC |
| BUG-116346 | HIVE-20772 | Hive | Record per-task CPU counters in LLAP |
| BUG-116344 | HIVE-20210 | Hive | Simple Fetch optimizer should lead to MapReduce when filter on non-partition column and conversion is minimal |
| BUG-116324 | KNOX-1679 | Knox | Admin API to add alias for a given topology |
| BUG-116295 | HIVE-20568 | Hive | HWC: Append SaveMode fails with Ranger |
| BUG-116287 | HIVE-21295 | Hive | Hive JDBC Storage Handler (Oracle): Few TPCDS queries producing improper results |
| BUG-116193 | KNOX-1683 | Knox | Ensure truststorePass is set in KnoxSession |
| BUG-116145 | N/A | Knox | Add Knox Cert to KnoxToken Service Response and KnoxSession Clients |
| BUG-116082 | N/A | Hive | Hive JDBC Storage Handler : TPCDS_GRP_4_22 OOM (mysql, postgres) |
| BUG-116018 | RANGER-2232 | Ranger | Implement search for Zone available in Ranger Admin on Security Zone page |
| BUG-116011 | HIVE-21043 | Hive | Enable move optimization for cloud replication with strict managed tables. |
| BUG-115992 | HIVE-19081 , HIVE-20859 , HIVE-20941 , HIVE-20960 , HIVE-21159 | Hive | List of Hive backports |
| BUG-115986 | KNOX-1623 | Knox | Fix SecureKnoxShellTest |

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|---|------------------|--|
| BUG-115974 | HBASE-21551 | HBase | Memory leak when use scan with STREAM at server side |
| BUG-115967 | CALCITE-2713 | Hive | Hive JDBC Storage Handler (Postgres): STDJOIN_66, UnionDistinct_3 and GRP_2_66 fails with "ERROR: type "pg_catalog.varchar_UTF-16LE" does not exist" |
| BUG-115911 | RANGER-2232 | Ranger | Add zone name in Policy detail popup in Access Tab |
| BUG-115908 | RANGER-2309 | Ranger | Partial search for group lookup on policy edit page not working properly. |
| BUG-115899 | YARN-9114 | YARN | [YARN-9114] [UI2] log service redirect url should also support user name |
| BUG-115829 | HIVE-21029 | Hive | Bootstrap of external tables for existing DLM clusters |
| BUG-115825 | RANGER-2232 | Ranger | Add select zone option during policy creation from Ranger Admin UI |
| BUG-115744 | N/A | Pig | Cherry-pick missing pig patches from HDP 2.6-maint to 3.1-maint |
| BUG-115724 | KNOX-1364 | Knox | Cookie scoping path should contain the topology name |
| BUG-115701 | STORM-2993 | Storm | Storm HDFS bolt throws ClosedChannelException when Time rotation policy is used |
| BUG-115700 | HIVE-21055 | Hive | HDP 2.6 to 3.1 Hive onprem-onprem replication running in serial mode. |
| BUG-115620 | HBASE-21568 | HBase | Permission denied exception on /mnt/hbase/cache.data on HDI environments |
| BUG-115596 | ATLAS-2969 | Atlas | Hive replication failed with Atlas import failure. |
| BUG-115585 | N/A | HWC | HWC: MergeBuilder |
| BUG-115580 | CALCITE-2719 | Hive | Backport CALCITE-2719: In JDBC adapter for MySQL, fix cast to INTEGER and BIGINT |
| BUG-115443 | HIVE-21013 | Hive | Hive JDBC Storage Handler (Oracle): select min(salary) failing with "java.io.IOException: java.io.IOException: Cannot find partitionColumn:salary in [EMP_NO, SALARY, FROM_DATE, TO_DATE]" |
| BUG-115436 | CALCITE-2733 HIVE-1383 HIVE-21468 | Hive | Hive JDBC Storage Handler (Oracle): Seeing "java.lang.IllegalArgumentException: Multiple entries with same key: APEX_ACTIVITY_LOG=JdbcTable {APEX_ACTIVITY_LOG} and APEX_ACTIVITY_LOG=JdbcTable {APEX_ACTIVITY_LOG}" |
| BUG-115404 | KNOX-1623 | Knox | Kerberos support for KnoxShell |
| BUG-115070 | KNOX-1474 , KNOX-1530 , KNOX-1531 , KNOX-1532 , KNOX-1533 , KNOX-1534 , KNOX-1535 | Knox | Backport KNOX-1530 - Improve Gzip Compression Handling Performance |
| BUG-115029 | RANGER-2291 | Ranger | Make optimized db schema script idempotent for all DB Flavors |
| BUG-114938 | ATLAS-2995 | | Term Entity disassociation doesn't happen via Knox proxy |
| BUG-114621 | RANGER-2295 | Ranger | Set specific Ranger version in optimized schema |
| BUG-114459 | ATLAS-2969 | Atlas | DLM Policy instances are failing if no atlas incremental data present |
| BUG-114234 | RANGER-2291 | Ranger | Core DB may be created partially in HA setup |

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|--|------------------|---|
| BUG-113835 | HADOOP-15855 | Hadoop Common | Backport HADOOP-15855: Review hadoop credential doc, including object store details |
| BUG-113745 | HIVE-21085 | EDWS | Materialized views registry starts non-external tez session |
| BUG-113612 | YARN-8961 | YARN | Flow Run End Time shows 'Invalid date' |
| BUG-113546 | N/A | EDWS | change service API to conform to REST approach |
| BUG-113270 | HADOOP-15466 | Hadoop Common | HADOOP-15466: Correct units in adl.http.timeout |
| BUG-112708 | N/A | Hive | REPL DUMP to output the included tables list in a metadata file _tables. |
| BUG-112707 | HIVE-21764 | Hive | REPL DUMP should detect and bootstrap any rename table events where old table was excluded but renamed table is included. |
| BUG-112702 | ATLAS-2921 | Atlas | Atlas hive-bridge import-hive.sh returns 0 in case of java error |
| BUG-112593 | N/A | Hive | EDWS: store cluster status in cluster list when creating |
| BUG-112536 | RANGER-2373 | Ranger | User creation POST and PUT response not showing groupIdList and groupNameList with expected data |
| BUG-112056 | HIVE-20911 | Hive | Bootstrap and incremental replication for External Tables in Hive. |
| BUG-112038 | HIVE-21763 | Hive | Incremental replication to allow changing include/exclude tables list in replication policy. |
| BUG-112034 | HIVE-21762 | Hive | REPL DUMP to support new format for replication policy input to take included tables list. |
| BUG-112028 | HIVE-21761 | Hive | Support table level replication in Hive |
| BUG-112015 | ATLAS-2885 ATLAS-2891 | Atlas | Atlas replication taking longer time |
| BUG-111962 | N/A | Hive | Support external table replication in Hive |
| BUG-111655 | N/A | YARN | RegistryDNS: when ambari user does not have permission to execute "su root" then starting/stopping registry dns fails |
| BUG-111386 | YARN-8803 | YARN | \Show flow runs in the order of recently created time in graph widgets |
| BUG-110789 | N/A | Hive | EDWS: implement db storage mechanism |
| BUG-110358 | ATLAS-3019 | Atlas | Export fails on setting replicatedTo value to null in export JSON. |
| BUG-110170 | ATLAS-3019 | Atlas | Using transforms (replace) with both supertype and subtype fails with NPE during import |
| BUG-109937 | ATLAS-2842 | Atlas | 2.6-maint - "Slow UI load and REST improvement for entities with ownedRef" and "Creating entity containing attribute of type set with duplicate values" |
| BUG-107745 | HIVE-16812 , HIVE-17231 , HIVE-17917 , HIVE-20635 , HIVE-20664 , HIVE-20694 , HIVE-20730 , HIVE-20738 | Hive | DeleteReaderValue optimization to load only the right bucket and only relevant events |
| BUG-95263 | YARN-7824 | YARN | [UI2] [YARN-7824] Yarn Component Instance page should include link to container logs |
| BUG-92696 | N/A | Hive | Replicate Dynamic Partition and CTAS operations on ACID tables |

Performance

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|--|------------------|---|
| BUG-121232 | HIVE-21980 | Hive | Parsing time can be high in case of deeply nested subqueries |
| BUG-121157 | HIVE-21934 | Hive | Materialized view on top of Druid not pushing everything |
| BUG-121142 | HIVE-21912 | Hive | Implement BlacklistingLlapMetricsListener |
| BUG-121080 | HADOOP-16341 | Hadoop Common | ShutDownHookManager: Regressed performance on Hook removals after HADOOP-15679 |
| BUG-120873 | HIVE-21232 , HIVE-21332 , HIVE-21509 , HIVE-21686 , HIVE-21805 , HIVE-21805 , HIVE-21332 , HIVE-21686 , HIVE-21509 . Revert HIVE-21232 | Hive | Fixes for HIVE-21232, HIVE-21332, HIVE-21686, HIVE-21509. Revert HIVE-21232 |
| BUG-120708 | HIVE-21834 | Hive | Avoid unnecessary calls to simplify filter conditions |
| BUG-120707 | HIVE-21818 | Hive | CBO: Copying TableRelOptHiveTable has metastore traffic |
| BUG-120375 | HIVE-21793 | Hive | CBO retrieves column stats even if hive.stats.fetch.column.stats is set to false |
| BUG-120295 | HIVE-21965 | Hive | HiveStrictManagedMigration process can run for days!! |
| BUG-120275 | ORC-498 | Hive | ReaderImpl and RecordReaderImpl open separate file handles |
| BUG-120047 | HIVE-21685 | Hive | Wrong simplification in query with multiple IN clauses |
| BUG-119970 | HIVE-21631 | Hive | Backport HIVE-21631 to 3.1 - Enhance metastore API to allow bulk-loading materialized views |
| BUG-119841 | HIVE-20274 | Hive | HiveServer2 ObjectInspectorFactory leaks for Struct and List object inspectors |
| BUG-119460 | HIVE-21568 | Hive | HiveRelOptUtil.isRowFilteringPlan should skip Project |
| BUG-118459 | HADOOP-16126 | Hadoop Common | Backport to 3.1 line - ipc.Client.stop() may sleep too long to wait for all connections |
| BUG-116888 | HADOOP-15940 | Hadoop Common | ABFS: For HNS account, avoid unnecessary get call when doing Rename |
| BUG-116750 | HADOOP-15972 | Hadoop Common | HADOOP-15972: ABFS: update LIST_MAX_RESULTS |
| BUG-115948 | ORC-397 , ORC-435 | Hive | Backport ORC-397 and ORC-435 to 3.1-maint |
| BUG-115898 | HIVE-21206 | Hive | Hive onprem-onprem replication is slow on HDP 3.1 cluster |
| BUG-115311 | HIVE-20989 | Hive | JDBC - The GetOperationStatus + log can block query progress via sleep() |
| BUG-115185 | MAPREDUCE-7158 | MapReduce | Backport MAPREDUCE-7158: Inefficient Flush Logic in JobHistory EventWriter |

Potential-data-loss

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|----------------------------|------------------|---|
| BUG-117071 | HIVE-16907 | Hive | "INSERT INTO" overwrite old data when destination table encapsulated by backquote |

Query-failure

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|----------------------------|------------------|--|
| BUG-121291 | HIVE-21746 | Hive | ArrayIndexOutOfBoundsException during dynamically partitioned hash join, with CBO disabled |
| BUG-120380 | N/A | HWC | executeQuery() blocks when dataframe is cached and in some other scenarios as well |

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|------------------------------|------------------|--|
| BUG-120161 | CALCITE-3052 | Hive | Error while applying rule MaterializedViewAggregateRule(Project-Aggregate): ArrayIndexOutOfBoundsException |
| BUG-120159 | CALCITE-3066 | Hive | RelToSqlConverter may incorrectly throw an AssertionError for some decimal literals |
| BUG-120041 | HIVE-21600 | Hive | GenTezUtils.removeSemiJoinOperator may throw out of bounds exception for TS with multiple children |
| BUG-120040 | HIVE-21613 | Hive | Queries with join condition having timestamp or timestamp with local time zone literal throw SemanticException |
| BUG-119962 | HIVE-21679 | Hive | Replicating a CTAS event creating an MM table fails. |
| BUG-119171 | HIVE-21539 | Hive | GroupBy + where clause on same column results in incorrect query rewrite |
| BUG-117334 | HIVE-21107 | Hive2 | Backport HIVE-21107 to 3.1-maint - "Cannot find field" error during dynamically partitioned hash join |
| BUG-117260 | SPARK-26379 | Spark | Use dummy TimeZoneId for CurrentTimestamp to avoid UnresolvedException in CurrentBatchTimestamp |

Security

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|--|------------------|---|
| BUG-120559 | N/A | Accumulo | Update Jetty to 9.2.28.v20190418 |
| BUG-120546 | N/A | Oozie | Add Oozie security headers to HDP-3.1.x line |
| BUG-120506 | KNOX-1596 , KNOX-1857 | Knox | Fixes for KNOX-1596 and KNOX-1857 |
| BUG-120369 | N/A | Oozie | Eliminate Jetty from Oozie RPMs |
| BUG-120252 | ZOOKEEPER-1045 , ZOOKEEPER-1392 , ZOOKEEPER-2498 , ZOOKEEPER-2693 , ZOOKEEPER-2726 | Zookeeper | Patch CVEs for Zookeeper 3.4.6 |
| BUG-120210 | N/A | Spark | Upgrade to jackson-databind version 2.9.8 |
| BUG-119572 | N/A | Knox | Upgrade bouncycastle libraries to 1.61 - knox |
| BUG-119509 | ZOOKEEPER-1392 | Zookeeper | Should not allow to read ACL when not authorized to read node |
| BUG-119333 | HIVE-21573 | Hive | Binary transport shall ignore principal if auth is set to delegationToken |
| BUG-118965 | HADOOP-15994 | Hadoop Common | Upgrade Jackson2 to 2.9.8 |
| BUG-118773 | N/A | Oozie | Upgrade to jackson-databind version 2.9.8 for Oozie |
| BUG-118256 | KNOX-1678 | Knox | Upgrade to jackson-databind version 2.9.8 |
| BUG-118233 | N/A | Druid | Upgrade to jackson-databind version 2.9.8 |
| BUG-118198 | YARN-9334 | YARN | YARN Service Client does not work with SPNEGO when knox is configured |
| BUG-116998 | OOZIE-3370 | Oozie | Property filtering is not consistent across job submission |

Stability

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|----------------------------|------------------|--|
| BUG-121463 | HIVE-22001 | Hive | Backport HIVE-22001 to 3.1-maint - ACID: getAcidState() should cache a recursive dir listing locally |

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|--|------------------|--|
| BUG-121363 | YARN-6695 | YARN | Race condition in RM for publishing container events vs appFinished events causes NPE |
| BUG-121355 | HADOOP-16376 | Hadoop Common | ABFS: Override access() to no-op for now |
| BUG-121297 | HADOOP-16461 | Hadoop Common | FileSystem cache lock parses XML within the lock |
| BUG-121233 | HIVE-18786 | Hive | NPE in Hive windowing functions |
| BUG-121142 | HIVE-21912 | Hive | Implement BlacklistingLlapMetricsListener |
| BUG-121140 | HDFS-12748 | HDFS | NameNode memory leak when accessing webhdfs GETHOMEDIRECTORY |
| BUG-121076 | HIVE-21815 | Hive | Stats in ORC file are parsed twice |
| BUG-121058 | HIVE-21992 | Hive | REPL DUMP throws NPE when dumping Create Function event. |
| BUG-120969 | HIVE-21966 | Hive | Llap external client - Arrow Serializer throws ArrayIndexOutOfBoundsException in some cases |
| BUG-120952 | HIVE-21963 | Hive | TransactionalValidationListener.validateTableStructure should check the partition directories in the case of partitioned tables |
| BUG-120912 | HADOOP-16350 | Hadoop Common | Ability to tell HDFS client not to request KMS Information from NameNode |
| BUG-120873 | HIVE-21232 , HIVE-21332 , HIVE-21509 , HIVE-21686 , HIVE-21768 , HIVE-21805 | Hive | Backport of HIVE-21805, HIVE-21332, HIVE-21686, HIVE-21509. Revert HIVE-21232 |
| BUG-120841 | N/A | HWC | close() method at session level to clean up resources |
| BUG-120784 | ZOOKEEPER-2247 | Zookeeper | Zookeeper quorum goes down sporadically |
| BUG-120740 | HIVE-21864 | Hive | LlapBaseInputFormat#closeAll() throws ConcurrentModificationException |
| BUG-120719 | HDFS-14323 | HDFS | Distcp fails in Hadoop 3.x when 2.x source webhdfs url has special characters in hdfs file path |
| BUG-120689 | N/A | HWC | Add a clean up method when spark application ends |
| BUG-120576 | YARN-9205 | YARN | Backport to "HDP-3.1-maint" - When using custom resource type, application will fail to run due to the CapacityScheduler throws InvalidResourceRequestException(GREATER_THAN_MAX_ALLOCATION) |
| BUG-120207 | CALCITE-3189 , HIVE-21976 | Hive | Hive JDBC Storage Handler (Oracle): TPCDS queries failing with 'ORA-00907: missing right parenthesis' |
| BUG-120206 | CALCITE-3189 , HIVE-22007 | Hive | Hive JDBC Storage Handler (Oracle): TPCDS queries failing with 'java.sql.SQLException: ORA-00904: "TRUE": invalid identifier' |
| BUG-120204 | CALCITE-3189 , HIVE-22007 | Hive | Hive JDBC Storage Handler (Oracle): TPCDS queries failing with "java.sql.SQLException: ORA-00905: missing keyword" |
| BUG-120203 | CALCITE-3189 | Hive | Hive JDBC Storage Handler (Oracle): TPCDS queries failing with "specified length too long for its datatype" |
| BUG-120017 | HIVE-21678 | Hive | CTAS creating a Full ACID partitioned table fails because of no writeId. |
| BUG-120007 | HIVE-21706 | Hive | Hive external table incremental repl dump fails if table dropped during replication in-progress |
| BUG-119988 | HIVE-21729 | Hive | Hive Warehouse Connector is offsetting timestamp column by exactly one second |
| BUG-119841 | HIVE-20274 | Hive | HiveServer2 ObjectInspectorFactory leaks for Struct and List object inspectors |

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|---|------------------|--|
| BUG-119495 | HIVE-21430 | Hive | INSERT into a dynamically partitioned table with hive.stats.autogather = false throws a MetaException |
| BUG-119391 | HIVE-21602 | Hive | Dropping an external table created by migration case should delete the data directory. |
| BUG-119144 | HIVE-21540 | Hive | Query with join condition having date literal throws SemanticException. |
| BUG-119143 | HIVE-21564 | Hive | Load data into a bucketed table is ignoring partitions specs and loads data into default partition. |
| BUG-119124 | HIVE-21497 | Hive | Direct SQL exception thrown by PartitionManagementTask |
| BUG-118996 | CALCITE-2727 | Hive | MV rewriting bails out incorrectly when a view does not contain any table reference |
| BUG-118827 | HIVE-21471 | Hive | Replicating conversion of managed to external table leaks HDFS files at target. |
| BUG-118682 | ZOOKEEPER-1342 , ZOOKEEPER-1907 , ZOOKEEPER-2029 , ZOOKEEPER-2186 , ZOOKEEPER-602 | Zookeeper | Backport Hotfix request for bugs fix for issue where zookeeper node Fails to Re-join Zookeeper Quorum |
| BUG-118374 | YARN-9336 | YARN | JobHistoryServer leaks CLOSE_WAIT tcp connections when using LogAggregationIndexedFileController |
| BUG-118335 | HIVE-21376 | Hive | Incompatible change in Hive bucket computation |
| BUG-117863 | HIVE-21385 | Hive | Hive JDBC Storage Handler (MySQL): Couple of tpccs queries get stuck in running state. |
| BUG-117703 | HIVE-21167 | Hive | Bucketing: Bucketing version 1 is incorrectly partitioning data |
| BUG-117634 | HIVE-21188 | Hive | Backport HIVE-21188: Hive query on view fails with SemanticException when there is a Ranger masked policy in place |
| BUG-117558 | HADOOP-16074 | Hadoop Common | WASB: Update container not found error code |
| BUG-117555 | HADOOP-16048 | Hadoop Common | ABFS: Fix Date format parser |
| BUG-117363 | HADOOP-16079 | Hadoop Common | Token.toString faulting if any token listed can't load |
| BUG-117221 | HIVE-21186 | Hive | External table REPL throwing NPE |
| BUG-117141 | HIVE-21156 | Hive | SharedWorkOptimizer may preserve filter in TS incorrectly |
| BUG-117140 | HIVE-21006 | Hive | Extend SharedWorkOptimizer to remove semijoins when there is a reutilization opportunity |
| BUG-116767 | HIVE-21255 | Hive | Hive JDBC Storage Handler - ArrayIndexOutOfBoundsException |
| BUG-116760 | HADOOP-16040 | Hadoop Common | ABFS: Bug fix for tolerateOobAppends configuration |
| BUG-116686 | HIVE-21253 | Hive | Hive JDBC Storage Handler (DB2): TPCDS queries failing with CNF exception |
| BUG-116684 | HIVE-16100 | Hive2 | Backport HIVE-16100 to 3.1-maint - Dynamic Sorted Partition optimizer loses sibling operators |
| BUG-116670 | HBASE-20647 | HBase | HBase is stuck in initializing master state |
| BUG-115948 | ORC-397 , ORC-435 | Hive | Backport ORC-397 and ORC-435 to 3.1-maint |
| BUG-115867 | HADOOP-15808 | Hadoop Common | Harden Token service loader use |
| BUG-115543 | HADOOP-15975 | Hadoop Common | ABFS: remove timeout check for DELETE and RENAME |
| BUG-115480 | HADOOP-15968 | Hadoop Common | ABFS: add try catch for UGI failure when initializing ABFS |

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|------------------------------|------------------|---|
| BUG-114163 | HIVE-20839 | Hive2 | "Cannot find field" error during dynamically partitioned hash join |
| BUG-112704 | HADOOP-15434 | Hadoop Common | Backport HADOOP-15434: Upgrade to ADLS SDK that exposes current timeout |
| BUG-112703 | HADOOP-15356 | Hadoop Common | Backport HADOOP-15356: Make HTTP timeout configurable in ADLS Connector |

Supportability

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|---|------------------|--|
| BUG-120079 | HIVE-21061 , HIVE-21398 , HIVE-21499 | Hive | Backport Multiple Hive fixes |
| BUG-120035 | KNOX-1827 , KNOX-1828 | Knox | Backport Knox fixes to HDP-3.1-maint |
| BUG-119380 | CALCITE-2858 | Hive | Improvements in JSON writer and reader for plans |
| BUG-119022 | HIVE-21507 | Hive | Hive swallows NPE if no delegation token found |
| BUG-118997 | CALCITE-2943 | Hive | Materialized view rewriting logic calls getApplicableMaterializations each time the rule is triggered |
| BUG-118995 | CALCITE-2942 | Hive | Materialized view rewriting logic instantiates RelMetadataQuery each time the rule is triggered |
| BUG-118827 | HIVE-21471 | Hive | Replicating conversion of managed to external table leaks HDFS files at target. |
| BUG-118682 | ZOOKEEPER-1343 , ZOOKEEPER-1907 , ZOOKEEPER-2029 , ZOOKEEPER-2186 , ZOOKEEPER-602 | Zookeeper | Backport Hotfix request for bugs fix for issue where zookeeper node Fails to Re-join Zookeeper Quorum |
| BUG-117872 | SQOOP-3420 | Sqoop | Invalid ERROR message initiates false alarms(Could not rename /tmp/sqoop-root/compile/xx.java to /root/./xx.java. Error: Destination '/root/./xx.java' already exists) |
| BUG-117557 | HADOOP-16044 | Hadoop Common | ABFS: Better exception handling of DNS errors followup |
| BUG-117556 | HADOOP-15662 | Hadoop Common | ABFS: Better exception handling of DNS errors |
| BUG-117118 | HIVE-20822 , HIVE-21046 | Hive | Improvements to push computation to JDBC from Calcite |
| BUG-116171 | HADOOP-16004 | Hadoop Common | ABFS: Convert 404 error response in AbfsInputStream and AbfsOutPutStream to FileNotFoundException |
| BUG-115965 | HIVE-19711 | Hive | Refactor HiveSchemaTool |
| BUG-115918 | HIVE-21032 | Hive | Refactor HiveMetaTool |

Usability

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|--|------------------|---|
| BUG-121380 | HIVE-15177 | Hive | Authentication with hive fails when kerberos auth type is set to fromSubject and principal contains _HOST |
| BUG-121351 | HIVE-18735 HIVE-21157 | Hive | Backport HIVE-18735 & HIVE-21957 |
| BUG-121180 | HIVE-19113 | Hive | Bucketing: Make CLUSTERED BY do CLUSTER BY if no explicit sorting is specified |
| BUG-121092 | HIVE-21822 , HIVE-21823 , HIVE-21832 , HIVE-21846 | Hive | Get the LlapDaemon metrics from TezAM |

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|--|------------------|--|
| BUG-121077 | HIVE-21225 | Hive | ACID: getAcidState() should cache a recursive dir listing locally |
| BUG-121075 | HIVE-21863 | Hive | Vectorization: CASE result type casting (non-cbo path) |
| BUG-121074 | HIVE-21742 | Hive | Vectorization: CASE result type casting |
| BUG-121043 | HIVE-21958 | Hive | The list of table expression in the inclusion and exclusion list should be separated by ' ' instead of comma. |
| BUG-121016 | HIVE-21886 | Hive | Handle rename events during replace policy in table level replication |
| BUG-120940 | N/A | HWC | HWC: onApplicationEnd() should not cause spark application to fail |
| BUG-120832 | HIVE-21913 | Hive | Backport HIVE-21913: GenericUDTFGetSplits should handle usernames in the same way as LLAP |
| BUG-120807 | N/A | Hive | Backport HIVE-20707 and related patches to 3.1-maint |
| BUG-120777 | HIVE-21837 | Hive | Hive exception "VoidColumnVector is not supported" for certain queries |
| BUG-120604 | HIVE-21938 HIVE-21948 | Hive | Hive 3 upgrade issues with large number of tables |
| BUG-120384 | HIVE-21754 | Hive | Metastore schema different precision can cause comparison issues |
| BUG-120374 | HIVE-21794 | Hive | Add materialized view parameters to sqlStdAuthSafeVarNameRegexes |
| BUG-120195 | HIVE-21681 | Hive | Describe formatted shows incorrect information for multiple primary keys |
| BUG-120158 | KNOX-1577 | Knox | Knox automatically derived dispatch whitelist doesn't seem to actually match the Knox domain |
| BUG-120044 | HIVE-21645 | Hive | Include CBO json plan in explain formatted |
| BUG-120043 | HIVE-21631 | Hive | Enhance metastore API to allow bulk-loading materialized views |
| BUG-119955 | OOZIE-3194 | Oozie | Permission issue- r+x for others under hadoop fs -ls /user/oozie/share/lib/lib_2017XXX |
| BUG-119906 | HIVE-21538 | Hive | Beeline: password source though the console reader did not pass to connection param |
| BUG-119619 | HIVE-20462 | Hive | "CREATE VIEW IF NOT EXISTS" fails if view already exists |
| BUG-119491 | HIVE-20555 | Hive | HiveServer2: Preauthenticated subject for http transport is not retained for entire duration of http communication in some cases |
| BUG-119480 | HIVE-20550 | Hive | Switch WebHcat to use beeline to submit Hive queries |
| BUG-119461 | HIVE-21592 | Hive | OptimizedSql is not shown when the expression contains CONCAT |
| BUG-119414 | HIVE-21591 | Hive | Backport HIVE-21591: Using triggers in non-LLAP mode should not require wm queue |
| BUG-119388 | SQOOP-3134 | Sqoop | --class-name should override default Avro schema name |
| BUG-119375 | HIVE-21587 | Hive | Explain formatted CBO should write row type in JSON |
| BUG-119257 | HADOOP-15997 , HADOOP-16199 | Hadoop Common | KMS client uses wrong UGI after HADOOP-14445 (3.1-maint) |
| BUG-119191 | HIVE-21290 | Hive | Restore historical way of handling timestamps in Parquet while keeping the new semantics at the same time |

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|--|------------------|---|
| BUG-119152 | ZOOKEEPER-2775 ZOOKEEPER-3253 | Zookeeper | Zookeeper server will not rejoin quorum/ Zookeeper Lost Lock |
| BUG-119143 | HIVE-21564 | Hive | Load data into a bucketed table is ignoring partitions specs and loads data into default partition. |
| BUG-119067 | N/A | Hive | HWC: For SaveMode overwrite, introduce truncate option |
| BUG-119056 | HIVE-21254 | Hive | Backport HIVE-21254: Pre-upgrade tool should handle exceptions and skip db/tables |
| BUG-118961 | N/A | Sqoop | Modify Sqoop to use Beeline instead of Hive command |
| BUG-118564 | HADOOP-16174 | Hadoop Common | Disable wildfly logs to the console |
| BUG-118491 | OOZIE-3439 | Oozie | Hive2 action is not parsing application ID for TEZ from log file properly |
| BUG-118405 | OOZIE-3440 | Oozie | Oozie Spark action unable to parse # character |
| BUG-118374 | YARN-9336 | YARN | [STAMPY] JobHistoryServer leaks CLOSE_WAIT tcp connections when using LogAggregationIndexedFileController |
| BUG-118340 | HIVE-21385 | Hive | Allow disabling pushdown of non-splittable computation to JDBC sources |
| BUG-118321 | HIVE-21389 | Hive | Hive: Unable to run queries due to "Error getting policies for serviceName=c11_hiveresponse=null" |
| BUG-118283 | HIVE-21184 | Hive | Add explain and explain formatted CBO plan with cost information |
| BUG-118281 | HIVE-18920 | Hive | CBO: Initialize the Janino providers ahead of 1st query |
| BUG-118119 | HIVE-21301 | Hive | "show tables" command does not include views |
| BUG-117974 | HIVE-21307 | Hive | Need to set GzipJSONMessageEncoder as default config for EVENT_MESSAGE_FACTORY. |
| BUG-117860 | HIVE-21529 | Hive | Hive support bootstrap of ACID/MM tables on an existing policy. |
| BUG-117809 | N/A | Hive | Backport HIVE-20679 to prevent large RDBMS footprint |
| BUG-117737 | HIVE-21388 | Hive | Constant UDF is not pushed to JDBCStorage Handler |
| BUG-117715 | HIVE-21261 | Hive | Incremental replication adds redundant COPY and MOVE tasks for external table events. |
| BUG-117379 | HADOOP-15860 | Hadoop Common | ABFS: Throw IllegalArgumentException when Directory/ File name ends with a period(.) |
| BUG-117378 | HADOOP-15954 | Hadoop Common | ABFS: Enable owner and group conversion for MSI and login user using OAuth |
| BUG-117361 | HDFS-13868 | Hadoop Common | WebHDFS: GETSNAPSHOTDIFF API NPE when param "snapshotname" is given but "oldsnapshotname" is not. |
| BUG-117267 | SQOOP-3438 | Sqoop | Sqoop Import with create hcatalog table for ORC will not work in HDP 3.x as the table created would be a ACID table and transactional |
| BUG-116990 | HIVE-21018 | Hive | Backport HIVE-21018 to 3.1-maint - Grouping sets size cannot be greater than 64 in HDP 3 |
| BUG-116348 | TEZ-3957 | Tez | Report TASK_DURATION_MILLIS as a Counter for completed tasks |
| BUG-115988 | HIVE-16957 | Hive | Support CTAS for auto gather column stats |
| BUG-115980 | HIVE-21039 | Hive | CURRENT_TIMESTAMP returns value in UTC time zone |
| BUG-115603 | SQOOP-3423 | Sqoop | Sqoop support for HiveServer2 password based authentication |

| Bug ID | Apache JIRA | Apache Component | Summary |
|------------|-------------------------------|------------------|--|
| BUG-115481 | HADOOP-15969 | Hadoop Common | ABFS: getNamespaceEnabled can fail blocking user access thru ACLs |
| BUG-115311 | HIVE-20989 | Hive | JDBC - The GetOperationStatus + log can block query progress via sleep() |
| BUG-115298 | HADOOP-15957 | Hadoop Common | WASB: Add asterisk wildcard support for PageBlobDirSet |
| BUG-114163 | HIVE-20839 | Hive2 | "Cannot find field" error during dynamically partitioned hash join |
| BUG-113460 | TEZ-4014 | Hive | EDWS: Read AM plugin descriptors from JSON |
| BUG-113341 | HIVE-20967 | Hive | Handle alter events when replicate from HDP 2.6.5 to HDP 3.0 |
| BUG-113252 | N/A | Hive | HWC: Detect client-mode vs. cluster-mode for Kerberized cluster usage |
| BUG-109059 | N/A | Zeppelin | In yarn cluster mode, zeppelin does not allow specifying additional artifacts to be uploaded via distributed cache |
| BUG-96263 | ZEPPELIN-3340 | Zeppelin | Streaming support in built-in visualizations |

Related Concepts

Component Versions

Known Issues

Summary of known issues for this release.

| Hortonworks Bug ID | Apache JIRA | Apache component | Summary |
|--------------------|-------------|------------------|--|
| BUG-123169 | N/A | Ambari and Oozie | <p>Description of the problem or behavior</p> <p>Oozie-Server won't start after "Install Packages" when upgrading from HDP-3.1.x to HDP-3.1.4 or later due to incompatible Tomcat version. The start fails with below error:</p> <pre>org.apache.catalina.startup.Catalina start org.apache.tomcat.util.digester.Digester startElement SEVERE: Begin event threw exception java.lang.ClassNotFoundException: org.apache.catalina.mbeans.ServerLifecycleList</pre> <p>Workaround</p> <ol style="list-style-type: none"> Downgrade bigtop-tomcat. <pre>yum downgrade bigtop-tomcat</pre> Start oozie-server. During the HDP upgrade, upgrade the package manually. Otherwise, the HDP upgrade will fail. <pre>yum upgrade bigtop-tomcat</pre> |


| Hortonworks Bug ID | Apache JIRA | Apache component | Summary |
|--------------------|-------------|--|---|
| BUG-79238 | N/A | Documentation, HBase, HDFS, Hive, MapReduce, Zookeeper | <p>Description of the problem or behavior</p> <p>SSL is deprecated and its use in production is not recommended. Use TLS.</p> <p>Workaround</p> <p>In Ambari: Use <code>ssl.enabled.protocols=TLSv1 TLSv1.1 TLSv1.2</code> and <code>security.server.disabled.protocols=SSL SSLv2 SSLv3</code>. For help configuring TLS for other components, contact customer support. Documentation will be provided in a future release.</p> |
| BUG-106494 | N/A | Documentation, Hive | <p>Description of Problem</p> <p>When you partition a Hive column of type double, if the column value is 0.0, the actual partition directory is created as "0". An AIOB exception occurs.</p> <p>Associated error message</p> <pre> 2018-06-28T22:43:55,498 ERROR 441773a0-851c-4b25-9e47-729183946a26 main exec.StatsTask: Failed to run stats task org.apache.hadoop.hive.ql.metadata.HiveException: java.lang.IndexOutOfBoundsException: Index: 8, Size: 8 at org.apache.hadoop.hive.ql.metadata.Hive.setPartit ~hive-exec-4.0.0-SNAPSHOT.jar:4.0.0- SNAPSHOT at org.apache.hadoop.hive.ql.stats.ColStatsProcesso ~hive-exec-4.0.0-SNAPSHOT.jar:4.0.0- SNAPSHOT at org.apache.hadoop.hive.ql.stats.ColStatsProcesso ~hive-exec-4.0.0-SNAPSHOT.jar:4.0.0- SNAPSHOT at org.apache.hadoop.hive.ql.exec.StatsTask.execut hive-exec-4.0.0-SNAPSHOT.jar:4.0.0- SNAPSHOT at org.apache.hadoop.hive.ql.exec.Task.executeTask hive-exec-4.0.0-SNAPSHOT.jar:4.0.0- SNAPSHOT at org.apache.hadoop.hive.ql.exec.TaskRunner.runSec hive-exec-4.0.0-SNAPSHOT.jar:4.0.0- SNAPSHOT at org.apache.hadoop.hive.ql.Driver.launchTask(Driv hive-exec-4.0.0-SNAPSHOT.jar:4.0.0- SNAPSHOT at org.apache.hadoop.hive.ql.Driver.execute(Driver hive-exec-4.0.0-SNAPSHOT.jar:4.0.0- SNAPSHOT at org.apache.hadoop.hive.ql.Driver.run </pre> <p>Workaround</p> <p>Do not partition columns of type double.</p> |

| Hortonworks Bug ID | Apache JIRA | Apache component | Summary |
|--------------------|---|---------------------|--|
| BUG-106379 | N/A | Documentation, Hive | <p>Description of the Problem</p> <p>The upgrade process fails to perform necessary compaction of ACID tables and can cause permanent data loss.</p> <p>Workaround</p> <p>If you have ACID tables in your Hive metastore, enable ACID operations in Ambari or set Hive configuration properties to enable ACID. If ACID operations are disabled, the upgrade process does not convert ACID tables. This causes permanent loss of data; you cannot recover data in your ACID tables later.</p> |
| BUG-106286 | N/A | Documentation, Hive | <p>Description of the Problem</p> <p>The upgrade process might fail to make a backup of the Hive metastore, which is critically important.</p> <p>Workaround</p> <p>Manually make a manual backup of your Hive metastore database before upgrading. Making a backup is especially important if you did not use Ambari to install Hive and create the metastore database, but highly recommended in all cases. Ambari might not have the necessary permissions to perform the backup automatically. The upgrade can succeed even if the backup fails, so having a backup is critically important.</p> |
| BUG-101082 | N/A | Documentation, Hive | <p>Description of the problem or behavior</p> <p>When running Beeline in batch mode, queries killed by the Workload Management process can on rare occasions mistakenly return success on the command line.</p> <p>Workaround</p> <p>There is currently no workaround.</p> |
| BUG-103495 | HBASE-20634 , HBASE-20680 , HBASE-20700 | HBase | <p>Description of the problem or behavior</p> <p>Because the region assignment is refactored in HBase, there are unclear issues that may affect the stability of this feature. If you rely on RegionServer Groups feature, you are recommended to wait until a future HDP 3.x release, which will return the stability of this features as it was available in HBase 1.x/HDP 2.x releases.</p> <p>Workaround</p> <p>There is currently no workaround.</p> |
| BUG-98727 | N/A | HBase | <p>Description of the problem or behavior</p> <p>Because the region assignment is refactored in HBase, there are unclear issues that may affect the stability of this feature. If you rely on Region replication feature, you are recommended to wait until a future HDP 3.x release, which will return the stability of this features as it was available in HBase 1.x/HDP 2.x releases.</p> <p>Workaround</p> <p>There is currently no workaround.</p> |
| BUG-105983 | N/A | HBase | <p>Description of the problem or behavior</p> <p>An HBase service (Master or RegionServer) stops participating with the rest of the HBase cluster.</p> <p>Associated error message</p> <p>The service's log contains stack traces that contain "Kerberos principal name does NOT have the expected hostname part..."</p> <p>Workaround</p> <p>Retrying the connection solves the problem.</p> |

| Hortonworks Bug ID | Apache JIRA | Apache component | Summary |
|--------------------|-----------------------------|------------------|--|
| BUG-121749 | HBASE-20662 | HBase | <p>Description of the problem or behavior</p> <p>HBase space quota policy state sometimes changes automatically from Disable to Observance.</p> <p>Workaround</p> <p>Disable the space quota policy on the table by setting LIMIT => NONE.</p> |
| BUG-96402 | HIVE-18687 | Hive | <p>Description of the problem or behavior</p> <p>When HiveServer2 is running in HA (high-availability) mode in HDP 3.0.0, resource plans are loaded in-memory by all HiveServer2 instances. If a client makes changes to a resource plan, the changes are reflected (pushed) only in the HiveServer2 to which the client is connected.</p> <p>Workaround</p> <p>In order for the resource plan changes to be reflected on all HiveServer2 instances, all HiveServer2 instances has to be restarted so that they can reload the resource plan from metastore.</p> |
| BUG-88614 | N/A | Hive | <p>Description of the problem or behavior</p> <p>RDMBS schema for Hive metastore contains an index HL_TXNID_INDEX defined as</p> <pre>CREATE INDEX HL_TXNID_INDEX ON HIVE_LOCKS USING hash (HL_TXNID);</pre> <p>Hash indexes are not recommended by PostgreSQL. For more information, see https://www.postgresql.org/docs/9.4/static/indexes-types.html</p> <p>Workaround</p> <p>It's recommended that this index is changed to type BTREE.</p> |
| BUG-60904 | KNOX-823 | Knox | <p>Description of the problem or behavior</p> <p>When Ambari is being proxied by Apache Knox, the QuickLinks are not rewritten to go back through the gateway. If all access to Ambari is through Knox in the deployment, the new Ambari QuickLink profile may be used to hide and/or change URLs to go through Knox permanently. Future release will make these reflect the gateway appropriately.</p> <p>Workaround</p> <p>There is currently no workaround.</p> |
| BUG-107399 | N/A | Knox | <p>Description of the problem or behavior</p> <p>After upgrade from previous HDP versions, certain topology deployments may return a 503 error. This includes, but may not be limited to, knoxsso.xml for the KnoxSSO enabled services.</p> <p>Workaround</p> <p>When this is encountered, a minor change through Ambari (whitespace even) to the knoxsso topology (or any other with this issue) and restart of the Knox gateway server should eliminate the issue.</p> |

| Hortonworks Bug ID | Apache JIRA | Apache component | Summary |
|--------------------|---------------------------|------------------|---|
| BUG-110463 | KNOX-1434 | Knox | <p>Description of the problem or behavior</p> <p>Visiting Knox Admin UI in any browser (Firefox / Chrome) sets the HTTP Strict Transport Security (HSTS) header for the host where Knox is running. Any subsequent request to other service on the same host (e.g. Graphana, Ranger etc.) over HTTP would get redirected to HTTPS due to this header.</p> <p>Please note that, this HSTS header is disabled in all Knox topologies by default.</p> <p>For more information, see https://knox.apache.org/books/knox-1-1-0/user-guide.html#HTTP+Strict+Transport+Security</p> <p>Impact</p> <p>All the non-SSL requests to other services get redirected automatically to HTTPS and would result in SSL errors like: <code>SSL_ERROR_RX_RECORD_TOO_LONG</code> or some other error.</p> <p>Workaround</p> <p>Use the manager.xml topology and remove the setting from the WebAppSec provider. You can do this using the Knox Admin UI. After you have removed the setting, close your browser or clear the cookies.</p> |
| BUG-121014 | N/A | Oozie | <p>Description of the problem or behavior</p> <p>If you are using a non-rpm based Linux distribution, for example, Debian, Ubuntu, Oozie cannot start after upgrade due to incorrect Apache Tomcat server version present in your operating system.</p> <p>Workaround</p> <p>Install Apache Tomcat 7 or later manually after you finished the upgrade.</p> |
| N/A | N/A | Spark | <p>Description of the problem or behavior</p> <p>Spark Structured Streaming job is intermittently reading from the beginning instead of the last offset maintained by the checkpoint.</p> <p>Associated error message</p> <pre>WARNid [Executor task launch worker for task lllid] <DATE_TIME> WARN ConsumerConfig: The configuration 'request.required.acks' was supplied but isn't a known config.</pre> <p>Workaround</p> <p>Upgrade to Cloudera Data Platform (CDP).</p> |

| Hortonworks Bug ID | Apache JIRA | Apache component | Summary |
|--------------------|-------------|------------------|--|
| BUG-101227 | N/A | Spark | <p>Description of the problem or behavior</p> <p>When Spark Thriftserver has to run several queries concurrently, some of them can fail with a timeout exception when performing broadcast join.</p> <p>Associated error message</p> <pre>Caused by: java.util.concurrent.TimeoutException: Futures timed out after [300 seconds] at scala.concurrent.impl.Promise \$DefaultPromise.ready(Promise.scala:219) at scala.concurrent.impl.Promise \$DefaultPromise.result(Promise.scala:223) at scala.concurrent.Await\$\$anonfun \$result\$1.apply(package.scala:107) at scala.concurrent.BlockContext \$DefaultBlockContext \$.blockOn(BlockContext.scala:53) at scala.concurrent.Await \$.result(package.scala:107) at org.apache.spark.sql.execution.joins.BroadcastH</pre> <p>Workaround</p> <p>You can resolve this issue by increasing the spark.sql.broadcastTimeout value.</p> |
| BUG-109979 | N/A | Spark | <p>Description of the problem or behavior</p> <p>YARN NodeManagers fail to start after a Spark patch upgrade due to YarnShuffleService CNF.</p> <p>Workaround</p> <p>To resolve this problem you must:</p> <p>Replace "{{spark2_version}}" with "\${hdp.version}" in "yarn.nodemanager.aux-services.spark2_shuffle.classpath" property value. For example, old value "{{stack_root}}/{{spark2_version}}/spark2/aux/*" -> new value "{{stack_root}}/\${hdp.version}/spark2/aux/*"</p> |

| Hortonworks Bug ID | Apache JIRA | Apache component | Summary |
|--------------------|-----------------------------|------------------|---|
| BUG-65977 | SPARK-14922 | Spark | <p>Description of the problem or behavior</p> <p>Since Spark 2.0.0, `DROP PARTITION BY RANGE` is not supported grammatically. In other words, only `=` is supported while `<`, `>`, `<=`, `>=` aren't.</p> <p>Associated error message</p> <pre>scala> sql("alter table t drop partition (b<1) ").show org.apache.spark.sql.catalyst.parser.ParseException: mismatched input '<' expecting '{'}', ','}(line 1, pos 31)</pre> <p>== SQL ==</p> <pre>alter table t drop partition (b<1) -----^^^^</pre> <p>Workaround</p> <p>To drop partition, use the exact match with `=`.</p> <pre>scala> sql("alter table t drop partition (b=0) ").show</pre> |
| BUG-114383 | N/A | Storm | <p>Description of the problem or behavior</p> <p>Submitting a topology to Storm fails. You see an error when you submit a topology to Storm.</p> <p>Associated error message</p> <p>The following error message is displayed when submitting a topology with the stack trace org.apache.storm.hack:</p> <pre>Exception in thread "main" java.lang.IllegalArgumentException at org.apache.storm.hack.shade.org.objectweb.asm. Source) at org.apache.storm.hack.shade.org.objectweb.asm. Source) at org.apache.storm.hack.shade.org.objectweb.asm. Source)</pre> <p>Workaround</p> <p>Find if `client.jartransformer.class` is present in the Storm configuration using the Ambari user interface. If the config is present, please set to ` ` and restart Storm service to take effect.</p> <p> Note: The value that you have to provide is one whitespace, Ambari restricts the configuration to have a value, but does not allow multiple whitespaces.</p> |

| Hortonworks Bug ID | Apache JIRA | Apache component | Summary |
|--------------------|-------------|------------------|--|
| BUG-106917 | N/A | Sqoop | <p>Description of the problem or behavior</p> <p>In HDP 3, managed Hive tables must be transactional (hive.strict.managed.tables=true). Transactional tables with Parquet format are not supported by Hive. Hive imports with --as-parquetfile must use external tables by specifying --external-table-dir.</p> <p>Associated error message</p> <pre>Table db.table failed strict managed table checks due to the following reason: Table is marked as a managed table but is not transactional.</pre> <p>Workaround</p> <p>When using --hive-import with --as-parquetfile, users must also provide --external-table-dir with a fully qualified location of the table:</p> <pre>sqoop import ... --hive-import --as-parquetfile --external-table-dir hdfs:///path/to/table</pre> |
| BUG-102672 | N/A | Sqoop | <p>Description of the problem or behavior</p> <p>In HDP 3, managed Hive tables must be transactional (hive.strict.managed.tables=true). Writing transactional table with HCatalog is not supported by Hive. This leads to errors during HCatalog Sqoop imports if the specified Hive table does not exist or is not external.</p> <p>Associated error message</p> <p>Store into a transactional table db.table from Pig/Mapreduce is not supported</p> <p>Workaround</p> <p>Before running the HCatalog import with Sqoop, the user must create the external table in Hive. The --create-hcatalog-table does not support creating external tables.</p> |
| BUG-109607 | N/A | YARN | <p>Description of the problem or behavior</p> <p>With wire encryption enabled with containerized Spark on YARN with Docker, Spark submit fails in "cluster" deployment mode. Spark submit in "client" deployment mode works successfully.</p> <p>Workaround</p> <p>There is currently no workaround.</p> |

| Hortonworks Bug ID | Apache JIRA | Apache component | Summary |
|--------------------|-------------------------------|------------------|--|
| BUG-110192 | N/A | YARN | <p>Description of the problem or behavior</p> <p>When YARN is installed and configured with KNOX SSO alone, Application Timeline Server web endpoint blocks remote REST calls from YARN UI and displays a 401 Unauthorized error.</p> <p>Associated error message</p> <p>401 Unauthorized error.</p> <p>Workaround</p> <p>Administrator needs to configure Knox authentication handler for Timeline Server and existing hadoop level configuration.</p> <p>Administrator needs to tune the following cluster specific configurations. Values for the last two property is in the hadoop.authentication.* properties file.</p> <pre> <property> <name>yarn.timeline-service.http- authentication.type</name> <value>org.apache.hadoop.security.authentication value> </property> <property> <name>yarn.timeline-service.http- authentication.authentication.provider.url</ name> <value>https://ctr- e138-1518143905142-455650-01-000002.hwx.site:44 gateway/knoxssso/api/v1/webssso</ value> </property> <property> <name>yarn.timeline-service.http- authentication.public.key.pem</name> <value>public.key.pem</value> </property> </pre> |
| RMP-11408 | ZEPPELIN-2170 | Zeppelin | <p>Description of the problem or behavior</p> <p>Zeppelin does not show all WARN messages thrown by spark-shell at the Zeppelin's notebook level.</p> <p>Workaround</p> <p>There is currently no workaround for this.</p> |
| N/A | N/A | N/A | <p>Description of the problem or behavior</p> <p>Open JDK 8u242 is not supported as it causes Kerberos failure.</p> <p>Workaround</p> <p>Use a different version of Open JDK.</p> |

Table 2: Technical Service Bulletin

| Technical Service Bulletin | Apache JIRA | Apache component | Summary |
|----------------------------|---------------------------|------------------|---|
| TSB-327 | HDFS-5698 | HDFS | <p>CVE-2018-11768: HDFS FSImage Corruption (potential DoS, file/dir takeover)</p> <p>In very large clusters, the in-memory format to store the user, group, acl, and extended attributes may exceed the size of the on disk format, causing corruption of fsImage.</p> <p>For the latest update on this issue, see the corresponding Knowledge article: CVE-2018-11768: HDFS FSImage Corruption (potential DoS, file/dir takeover)</p> |
| TSB-405 | N/A | N/A | <p>Impact of LDAP Channel Binding and LDAP signing changes in Microsoft Active Directory</p> <p>Microsoft has introduced changes in LDAP Signing and LDAP Channel Binding to increase the security for communications between LDAP clients and Active Directory domain controllers. These optional changes will have an impact on how 3rd party products integrate with Active Directory using the LDAP protocol.</p> <p>Workaround</p> <p>Disable LDAP Signing and LDAP Channel Binding features in Microsoft Active Directory if they are enabled</p> <p>For the latest update on this issue see the corresponding Knowledge article: TSB-2021 405: Impact of LDAP Channel Binding and LDAP signing changes in Microsoft Active Directory</p> |
| TSB-406 | N/A | HDFS | <p>CVE-2020-9492 Hadoop filesystem bindings (ie: webhdfs) allows credential stealing</p> <p>WebHDFS clients might send SPNEGO authorization header to remote URL without proper verification. A maliciously crafted request can trigger services to send server credentials to a webhdfs path (ie: webhdfs://...) for capturing the service principal</p> <p>For the latest update on this issue see the corresponding Knowledge article: TSB-2021 406: CVE-2020-9492 Hadoop filesystem bindings (ie: webhdfs) allows credential stealing</p> |

| Technical Service Bulletin | Apache JIRA | Apache component | Summary |
|----------------------------|-----------------------------|------------------|---|
| TSB-450 | HBASE-21000 | HBase | <p>Default limits for PressureAwareCompactionThroughputController are too low</p> <p>HDP and CDH releases suffer from low compaction throughput limits, which cause storefiles to back up faster than compactions can re-write them.</p> <p>For the latest update on this issue see the corresponding Knowledge article: Cloudera Customer Advisory: Default limits for PressureAwareCompactionThroughputController are too low</p> |
| TSB-458 | N/A | HDFS | <p>Possible HDFS Erasure Coded (EC) Data Files Corruption in EC Reconstruction</p> <p>Cloudera has detected two bugs that can cause corruption of HDFS Erasure Coded (EC) files during the data reconstruction process.</p> <p>For the latest update on this issue see the corresponding Knowledge article: Cloudera Customer Advisory: Possible HDFS Erasure Coded (EC) Data Files Corruption in EC Reconstruction</p> |
| TSB-463 | N/A | HBase | <p>HBase Performance Issue</p> <p>The HDFS short-circuit setting <code>dfs.client.read.shortcircuit</code> is overwritten to disabled by <code>hbase-default.xml</code>. HDFS short-circuit reads bypass access to data in HDFS by using a domain socket (file) instead of a network socket. This alleviates the overhead of TCP to read data from HDFS which can have a meaningful improvement on HBase performance (as high as 30-40%).</p> <p>For the latest update on this issue see the corresponding Knowledge article: TSB 2021-463: HBase Performance Issue</p> |
| TSB-480/2 | HIVE-24224 | Hive | <p>Hive ignores the property to skip a header or footer in a compressed file</p> <p>Incorrect results can occur running SELECT queries if count value is greater than 0.</p> <p>For the latest update on this issue see the corresponding Knowledge article: TSB 2021-480.2: Hive ignores the property to skip a header or footer in a compressed file</p> |

| Technical Service Bulletin | Apache JIRA | Apache component | Summary |
|----------------------------|--|------------------|--|
| TSB-494 | <ul style="list-style-type: none"> HBASE-20781 HBASE-25459 PHOENIX-5250 | HBase | <p>Accumulated WAL Files Cannot be Cleaned up When Using Phoenix Secondary Global Indexes</p> <p>The Write-ahead-log (WAL) files for Phoenix tables that have secondary global indexes defined on them, cannot be automatically cleaned up by HBase, leading to excess storage usage and possible error due to filling up the storage.</p> <p>Workaround</p> <p>Perform rolling restart of HBase if the number of znodes under hbase-secure/splitWAL in ZooKeeper is greater than 8000.</p> <p>For the latest update on this issue see the corresponding Knowledge article: TSB 2021-494: Accumulated WAL Files Cannot be Cleaned up When Using Phoenix Secondary Global Indexes</p> |
| TSB-508 | HIVE-25159 | Hive | <p>Removing support for Order By statement from HWC LLAP mode</p> <p>The Order By statement is no longer supported for HWC LLAP mode.</p> <p>Workaround</p> <ol style="list-style-type: none"> Set <code>hive.llap.external.splits.order.by.force.single.split</code> and restart HiveServer Interactive (HSI). Add an explicit Order By on the dataframe as it no longer will be enforced by <code>executeQuery()</code> after the above configuration is set. <p>For the more information on this issue, see the corresponding Knowledge article: TSB 2021-508: Removing support for Order By statement from HWC LLAP mode</p> |

| Technical Service Bulletin | Apache JIRA | Apache component | Summary |
|----------------------------|-------------|------------------|--|
| TSB-497 | N/A | Solr | <p>CVE-2021-27905: Apache Solr SSRF vulnerability with the Replication handler</p> <p>The Apache Solr ReplicationHandler (normally registered at "/replication" under a Solr core) has a "masterUrl" (also "leaderUrl" alias) parameter. The "masterUrl" parameter is used to designate another ReplicationHandler on another Solr core to replicate index data into the local core. To help prevent the CVE-2021-27905 SSRF vulnerability, Solr should check these parameters against a similar configuration used for the "shards" parameter.</p> <p>For the latest update on this issue, see the corresponding Knowledge article: TSB 2021-497: CVE-2021-27905: Apache Solr SSRF vulnerability with the Replication handler</p> |
| TSB-512 | N/A | HBase | <p>HBase MOB data loss</p> <p>HBase tables with the MOB feature enabled may encounter problems which result in data loss.</p> <p>For the latest update on this issue, see the corresponding Knowledge article: TSB 2021-512: HBase MOB data loss</p> |

Legal Information

Hortonworks Data Platform 3.1.4

Copyright information for Hortonworks Data Platform components may be found within the documentation accompanying each component in a particular release.

Hortonworks Data Platform incorporates software from various open source projects released primarily under the Apache Software License 2.0 ("ASLv2"). Other software included may be released under the terms of alternative ASLv2 compatible open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Hortonworks Data Platform page for more information on Hortonworks Data Platform technology. For more information on Hortonworks support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Hortonworks reserves the right to change any products described herein at any time, and without notice. Hortonworks assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Hortonworks.

Trademark: Hortonworks Data Platform is/are trademark/s of Hortonworks, Inc. in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH HORTONWORKS, HORTONWORKS DOES NOT MAKE OR GIVE ANY REPRESENTATION, WARRANTY, OR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH HORTONWORKS TECHNOLOGY

OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. HORTONWORKS DOES NOT WARRANT THAT Hortonworks Data Platform WILL OPERATE UNINTERRUPTED OR THAT IT WILL BE FREE FROM DEFECTS OR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION OR UNAVAILABILITY, OR THAT # WILL MEET ALL OF CUSTOMER'S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, HORTONWORKS EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.