Cloudera Manager 7.6.2

Release Notes

Date published: 2020-11-30
Date modified: 2022-05-12

https://docs.cloudera.com/

Legal Notice

Contents

Cloudera Manager 7.6.2 Release Notes

Known issues, fixed issues and new features for Cloudera Manager and CDP Private Cloud Base.

What's New in Cloudera Manager 7.6.2

New features and changed behavior for Cloudera Manager 7.6.2.

Merged keytab with load balancer principal for Kafka service.

Configuration options have been added to use load balancer in front of kafka brokers.

Validate HBase replication setup after first time setup is done

- When an HBase replication policy is created then there is a new field available in the API body validateReplicationSetup:
- By default, its value is false, which doesn't make any difference: HBase replication policy creation will be the same as before.
- When it's set to true, then in the end of the HBase replication policy creation, additional steps will check if replication works properly between the source and destination. HBase policy creation will work the same way as before, no additional action is required by users.
- If you want to validate whether HBase replication works properly, use that new field in the API body: "validateReplicationSetup": true

Enable JMX Authentication by default

In Kafka, JMX authentication is enabled by default. Passwords for monitor and control users are generated by default, but if values are set, they will be the values used.

Cloudera Manager client support for hadoop.security.group.mapping.ldap.bind.password in jceks

With this change, the 'ldap.bind.password' parameter will be appended to core-site.xml. This will enable HDFS's clients (like YARN) to access LDAP functionalities.

Automatically refresh metric filter and collection settings

A new parameter 'metric_config_auto_refresh' has been added. By setting this parameter to true, Metric Collection and Metric Filter parameters will be set automatically if they are changed without a role restart or configuration refresh.

Option for logout landing page in the Cloudera Manager Admin Console

With SAML, the logout request re-initiates a new SAML authentication request, therefore the session wasn't terminated unless SAML SLO (Single-LogOut) is enabled. We added a logout success page to break the SAML logout loop which has a Return to Login Page button for new session creation.

If Single-logout (SLO) is supported by the identity provider (IdP), you can enable this feature by turning on SLO in the SAML configuration during SAML set up, .

Configurations for Streams Replication Manager explicit topic creation

Streams Replication Manager now supports user-configured values for partition count and/or replication factor of it's internal topics. The following configuration parameters were introduced: metrics.topic.partition.count, control.topic.replication.factor, streams.replication.manager.service.remote.advertisement.topic.replication.factor, streams.replication.manager.service.streams.replication.factor.

HbaseReplicationFirstTimeSetupReset removes the HDFS files from both sides, and clears the aux entry.

Introduced a new API endpoint: hbaseReplicationFirstTimeSetupCleanAndReset, which cleans up and resets the HBase replication first time setup between the given source and destination:

Removes the jceks file

Clears the HBASE_REPLICATION_AUXILIARY_INFO configuration

Add support for JSON schema type in the registry config template s

Schema Registry now supports Avro and JSON schemas.\

Add an optional "Cluster Alias" field to the existing Kafka external account type

The Kafka External Account name no longer has to match the alias of the Kafka cluster this account is describing. A new "Cluster Alias" field has been introduced for Kafka External Accounts, providing users with freedom to name the accounts as they wish. This field is optional - if it is specified, its value will be used when connecting to the Kafka cluster, otherwise the account name will be used.

Need to update default value for keystore alias in Ranger

Default keystore alias is now configurable for Ranger Admin and is not set to the Hostname by default.

Chive: add an option to ignore certain partitionParameters when comparing and improve location comparison

When customer replicated Hive data, then the replication had poor performance because all partitions were recreated.

This fix introduces a new parameter called HIVE_IGNORED_PARTITION_PARAMETERS for hive_replication_env_safety_valve. The value is a comma separated list of Hive partition parameters that will not be compared during the import stage. This means that even if these partition parameters don't match between the exported and existing partitions, the partition will not be dropped and recreated. Since these parameters can easily differ for metadata only replications, it's safe to ignore them in those cases. By setting these parameters, Hive replication performance can be improved.

OPSAPS-62697 Need to add property javax.security.auth.useSubjectCredsOnly to JVM args

Users having issues with Atlas - Solr communication should add the required argument through the Atlas Service Environment Advanced Configuration Snippet (Safety Valve) configuration parameter. Use the following key and value:

Key: ATLAS_CUSTOM_OPTS

Value: -Djavax.security.auth.useSubjectCredsOnly=false

Kafka topic entity import into Atlas

The Import Kafka Topics Into Atlas command has been introduced. It is available from the actions list of the Kafka service or any of the brokers in Cloudera Manager. The command can be used by users who have permissions to handle service configurations in the system. The Atlas service is required for the command, otherwise the process will fail. When "not Ranger" is preferred as default authorizer, then the "kafka" service user has to be defined in the selected authorizer service.

Add "emit.consumer.metrics" config to SMM CSD, and remove (now) unused SMON host/port configs.

The cm.metrics.service.monitor.host" and "cm.metrics.service.monitor.port" Streams Messaging Manger configuration properties have been removed. These properties are no longer required because Streams Messaging Manager automatically detects the ServiceMonitor's location.

The following new configuration parameter for Streams Messaging Manager has been added: "emit.consumer.metrics": When this is set to false, Streams Messaging Manager does not emit historic ConsumerGroup metrics into the ServiceMonitor, meaning that historic metrics (for group Lag and CommittedOffset) are not available for Groups in Streams Messaging Manager. These metrics are used to populate the charts at the bottom of the ConsumerGroupDetail page, or accessed via the api/v2/admin/metrics/consumers/group/{groupId} REST API endpoint.

Allow HTTP Response Headers to be Configured for Kafka Connect

When a request was made to Kafka connect, the response did not contain a HSTS header. With this, the HSTS header has been added as default to the Kafka Connect REST API response.

Enable setting offset in Schema Registry database

Schema Registry offset ranges can now be configured via Cloudera Manager: minimum and maximum value can be set.

Cloudera Manager diagnostic bundle now includes Cloudera Manager database information

A Cloudera Manager diagnostic bundle will now include an additional file named cm_db_dump_stats.txt, in the cm_db_dump directory of the generated bundle. This file contains statistical data of the Cloudera Manager database dump thread responsible for collecting table data. The purpose of this is to help tally the following:

- Number of tables in the database
- Total collected tables - Number of records of each table and how many were collected
- Size of the records and size of the collection
- Status of the collection (in-progress or completed)
- Time taken to collect each table

Introduce allowed nexus urls config for Kafka Connect

Users are now able to configure which nexus URLs are allowed for Stateless Nifi connector configurations using the kafka.connect.allowed.nexus.urls property. By default, this is set to empty, which means to "allow everything". During connector creation / modification / validation, the nexus.url property will be validated against this list.

Add custom Kerberos path to agent and Cloudera Manager

Customers are facing the following issues when modifying the default path of krb5.conf in Cloudera Manager :

- The credential generation for roles are failing because KDC authentication with Cloudera Manager server fails.
- Services are failing to authenticate with Cloudera Manager Agent once manually getting services up by applying hacks (i.e adding relevant JVM arguments or environment variables)
- A few services like HDFS, Livy, HiveServer and Knox are reportedly failing as they are unable to locate the new Kerberos path.

To set a custom path, follow the steps in this Knowledge Base article: How to use a custom Kerberos configuration path for a cluster running with Kerberos (MIT)

Configuration property for enabling HTTP Strict Transport Security

Fixed an issue where customers were unable to configure Cruise Control to include Strict Transport Security headers in the responses of the API. A new configuration property, webserver.ssl.sts.enabled has been added for Cruise Control in Cloudera Manager. Setting this value to true configures Cruise Control to include the Strict Transport Security header in the web server responses when SSL is enabled.

Add flags to force Connectors to override the JAAS, and restrict the usage of the Worker principal

Kafka Connect now allows users to force Connectors to override the JAAS configuration of the Kafka connection, and also forbids using the same Kerberos credentials as the Connect Worker is using.

Add OpDB Agent to Knox

Configuration for autodiscovery of the OpDB Agent has been added to Knox. The OpDB Agent is a new service for CDP Operational Database that needs to be discovered by Knox.

Ranger server work directory is now configurable

Ranger Admin / KMS / KMS-KTS server work directory can now be configured through the parameter ranger.tomcat.work.dir.

Ranger RMS server work directory can now be configured through the parameter ranger-rms.tomcat.work.dir.

Ranger Raz server work directory can now be configured through the parameter ranger.raz.tomcat.work.dir.

authzmigrator : Skipping policy item creation for {OWNER}

After the Sentry migration to Ranger there are lots of {OWNER} policies being created, which are very difficult to administer. This occurs during CDH to CDP migration. If you want to skip {OWNER} policies, add the following properties in authorization-migration-site.xml

```
<property>
    <name>authorization.migration.skip.owner.policy</name>
    <value>true</value>
</property>
```

Fixed Issues in Cloudera Manager 7.6.2

Fixed issues in Cloudera Manager 7.6.2

Cloudera Bug: OPSAPS-48098: Cannot cancel Impala queries from the Cloudera Manager Admin Console or impalaQueries API on a non-kerberized cluster if Cloudera Manager/impalad LDAP authentication is enabled

Cloudera Manager might not be able to cancel an Impala query if LDAP is enabled in the Impala service and is different from Cloudera Manager's LDAP authentication. Two new Impala configurations have been added to allow an administrator to add a LDAP username and password to the Cloudera Manager's Impala configuration, when such a LDAP username and password is provided and when Impala is configured with LDAP.

Cloudera Bug: OPSAPS-61278: The Streams Replication Manager Client’s secure storage fails to generate correctly in FIPS-enabled clusters

The Streams Replication Manager client configuration secure storage did not work on FIPS-enabled clusters, which caused problems when trying to use the srm-control tool with the deployed client configurations. Secure storage is now working in FIPS enabled clusters.

Cloudera Bug: OPSAPS-62087: Upgrade ttorrent-core

The ttorrent-core dependency has been removed due to CVE issues. This fixes the following CVEs: CVE-2008-0071, CVE-2008-0364, CVE-2008-4434, CVE-2008-7166, CVE-2014-8515, CVE-2015-5474

Cloudera Bug: OPSAPS-62548: TopicMetrics get deleted from Cloudera Manager during restart or Kafka partition reassignment

Fixed an issue where KafkaTopicMetrics were accidentally deleted from ServiceMonitor's Timeseries database during a Kafka restart or partition leader change.

Cloudera Bug: OPSAPS-62581: Address log4j CVE-2021-44228

This patch addresses CVE-2021-44228 for log4j.

Cloudera Bug: OPSAPS-62588: Cloudera Manager: Upgrade Logredactor to version 2.0.11 to remediate CVE-2021-44228

This patch addresses CVE-2021-44228 affecting log4j library.

Cloudera Bug: OPSAPS-62670: Upgrade the transitive dependency of Log4j2 used by Hive in Cloudera Manager to 2.17.1

In some scenarios, a role may already exist in Ranger, In that case Authzmigrator tool throws an error because the role already exists. This has been fixed.

Cloudera Bug: OPSAPS-62760: Jetty version shown in HTTP header

The Jetty error page has been modified to hide the version and Jetty server name.

Cloudera Bug: OPSAPS-62770: Cloudera Manager Server setting HBaseCmdOpts.hbaseClusterKey incorrectly.

When using the Cloudera Manager API, when ApiHBaseReplicationArguments.hbaseClusterKey was left null/empty, the Cloudera Manager server supplied an incorrect default value for the cluster key. With this fix, if the value is null or empty, the correct cluster key is set.

The cluster key generated is hbase.zookeeper.quorum:hbase.zookeeper.property.clientPort:zookeeper.znode.parent.

Cloudera Bug: OPSAPS-62812: HostMonitor: Missing HSTS Header

Fixed an issue where the Service Monitor and Host Monitor only open a single port when TLS is used for increased security.

Cloudera Bug: OPSAPS-62836: WAITING_FOR_SOURCE_RESTART status when cluster setup hasn't started

Sometimes during HBase Replication First Time setup, the status may report WAITING_FOR_SOURCE_RESTART before it is really waiting for the source to restart. This behavior has been fixed to report the correct status.

Cloudera Bug: OPSAPS-62910: HBase peer is not removed if several policies are deleted simultaneously

When attempting concurrent deletion of an HBase policy, when the last policy referencing the HBase peer is deleted, the HBase peer will be deleted as well.

Cloudera Bug: OPSAPS-62948: Simultaneous HBase policy creation can fail

When attempting to create multiple HBase policies concurrently, sometimes all of the create operations did not complete. With this fix, all concurrent create operations should succeed.

Cloudera Bug: OPSAPS-62956: Failed delete of policy for which creation failed

When multiple HBase policies are created concurrently, sometimes some of them would fail, or succeed but have wrong column families saved. With this fix, they should now all succeed, and the underlying HBase peer should contain all of the tables/column families from all the policies created.

Cloudera Bug: OPSAPS-62970: Upgrade AuthzMigrator to handle Ignore "already exists" failures

This fix handles the scenario where roles already exist in Ranger.

Cloudera Bug: OPSAPS-62980: Browser crashes with out-of-memory error when Dashboards contain numerous or complex charts

Browser crashed with out-of-memory error on complex Dashboard pages due to charts resize, drag and drop functionality.

You can now use the new Edit Layout menu option on the Dashboard pages and chart sections (e.g. Cluster, Host, Service, Role status pages) for enabling resize functionality. By default, the charts position and size are static. In order to modify the chart size, position they use the Edit Layout menu option.

Cloudera Bug: OPSAPS-62998: Synchronization of HBase peer disable/enable

When concurrent HBase enable/disable policy operations are run, sometimes the final policy status was inconsistent. For instance, after the last operation that completed, the policy status was not set as desired. This situation can be sporadic. With this fix, the HBase policy status will be set reliably.

Cloudera Bug: OPSAPS-63017: Kafka Connect not appearing in Streams Messaging Manager UI

Kafka Connect was sometimes not appearing on the Streams Messaging Manager UI, even when it was properly configured.

Cloudera Bug: OPSAPS-63069: Kafka keystore and truststore type is not configured for Cruise Control metrics reporter

The keystore and truststore types are now correctly supported by the Cruise Control metrics reporter in the Kafka broker. Previously, the type would not be configured, causing issues if the store type did not match the default type. Now the type is correctly configured.

Cloudera Bug: OPSAPS-63077: Fix decommisioning/recommissioning nodemanager failure in YARN

If Zookeeper config store is set in YARN, and QueueManager is not used every YARN node decommission caused an exception, because it called the refreshQueues command (which is disallowed if a mutable configuration store is used). This has been fixed

Cloudera Bug: OPSAPS-63104: Streams Replication Manager Service Co-Located Service password default is invalid

The Streams Replication Manager Service Basic Authentication would not work with the default, random generated password. Streams Replication Manager Service Basic Authentication default password is identical on all Streams Replication Manager Service role instances without any extra actions.

Cloudera Bug: OPSAPS-63138: HBase first time setup is green although remote command failed

When Cloudera Manager does HBase replication first time setup between the given source and destination and the HBASE_REPLICATION_AUXILIARY_INFO parameter was unintentionally left there from an older HBase replication first time setup, then the 'Admin setup remote command' step failed and was ignored on the source. Therefore the HBase replication first time setup command was still displayed as successful on the destination. This has been fixed

Cloudera Bug: OPSAPS-63158: Cruise Control may not able to start after upgrade, when the initial cluster version is 7.2.14 or lower and the post-upgrade version is 7.2.14.

This issue occured when any of Cruise Control's goals lists (default-, supported-, hard-, self-healing-, anomaly detection goals) contain com.linkedin.kafka.cruisecontrol.analyzer.goals.RackAwareGoal.

Cruise Control now automatically overrides all of the deprecated RackAwareGoal occurrences with RackAwareDistributionGoal, in its configurations during a parcel upgrade.

Cloudera Bug: OPSAPS-63419: Yarn MR Aggregation job fails with pt_PT locale

Fixed an issue where the Yarn MR Aggregation job got stuck with certain locale settings.

Cloudera Bug: OPSAPS-62471: Revert OPSAPS-32569 Disable HBase replication monitoring when Kerberos is enabled

Fixed an issue where HBase replication monitoring was not enabled when Kerberos was enabled.

Known Issues in Cloudera Manager 7.6.2

Known issues in Cloudera Manager 7.6.2

Cloudera bug: OPSAPS-59764: Memory leak in the Cloudera Manager agent while downloading the parcels.

When using the M2Crpyto library in the Cloudera Manager agent to download parcels causes a memory leak.

The Cloudera Manager server requires parcels to install a cluster. If any of the URLs of parcels are modified, then the server provides information to all the Cloudera Manager agent processes that are installed on each cluster host.

The Cloudera Manager agent then starts checking for updates regularly by downloading the manifest file that is available under each of the URLs. However, if the URL is invalid or not reachable to download the parcel, then the Cloudera Manager agent shows a 404 error message and the memory of the Cloudera Manager agent process increases due to a memory leak in the file downloader code of the agent.

To prevent this memory leak, ensure all URLs of parcels in Cloudera Manager are reachable. To achieve this, delete all unused and unreachable parcels from the Cloudera Manager parcels page.

OPSAPS-63640: Monitoring a high number of Kafka producers might cause Cloudera Manager to slow down and run out of memory

This issue has two workarounds. You can either configure a Kafka producer metric allow list or completely disable producer metrics.

- Configure a Kafka producer metric allow list:

  A producer metric allow list can be configured by adding the following properties to Kafka Broker Advanced Configuration Snippet (Safety Valve) for kafka.properties.

  ```
  producer.metrics.whitelist.enabled=true
  producer.metrics.whitelist=[***ALLOW LIST REGEX***]
  ```

  Replace *[***ALLOW LIST REGEX***]* with a regular expression matching the client.id of the producers that you want to add to the allow list. This regular expression uses the java.util.regex. Pattern class to compile the regular expression, and uses the match() method on the client.id to determine whether it fits the regular expression.

  Once configured, the metrics of producers whose client.id does not match the regular expression provided in producer.metrics.whitelist are filtered.Kafka no longer reports these metrics through the HTTP metrics endpoint. Additionally, existing metrics of the producers whose client.id does not match the regular expression are deleted.

  Because the allow list filters metrics based on the client.id of the producers, you must ensure that the client.id property is specified in each producer's configuration. Automatically generated client IDs might cause the number of unnecessary metrics to increase even if an allow list is configured.
- Completely disable producer metrics:

  Producer metrics can be completely disabled by unchecking the Enable Producer Metrics Kafka service property.

Cloudera bug: OPSAPS-63881: When CDP Private Cloud Base is running on RHEL/CentOS/Oracle Linux 8.4, services fail to start because service directories under the /var/lib directory are created with 700 permission instead of 755.

Run the following command on all managed hosts to change the permissions to 755. Run the command for each directory under /var/lib:

```
chmod -R 755 [***path_to_service_dir***]
```

OPSAPS-65189: Accessing Cloudera Manager through Knox displays the following error:

Bad Message 431 reason: Request Header Fields Too Large

Workaround: Modify the Cloudera Manager Server configuration /etc/default/cloudera-scm-server file to increase the header size from 8 KB, which is the default value, to 65 KB in the Java options as shown below:

```
export CMF_JAVA_OPTS="...existing options...
-Dcom.cloudera.server.cmf.WebServerImpl.HTTP_HEADER_SIZE_BYTES=
65536
-Dcom.cloudera.server.cmf.WebServerImpl.HTTPS_HEADER_SIZE_BYTE
S=65536"
```

OPSAPS-65213: Ending the maintenance mode for a commissioned host with either an Ozone DataNode role or a Kafka Broker role running on it, might result in an error.

You may see the following error if you end the maintenance mode for Ozone and Kafka services from Cloudera Manager when the roles are not decommissioned on the host.

```
Execute command Recommission and Start on service OZONE-1
Failed to execute command Recommission and Start on service OZ
ONE-1
Recommission and Start
Command Recommission and Start is not currently available for e
xecution.
```

To resolve this issue, use the API support feature to take the host out of maintenance mode.

1. Log into Cloudera Manager as an Administrator.
2. Go to Hosts All Hosts .
3. Select the host for which you need to end the maintenance mode from the available list and click the link to open the host details page.
4. Copy the Host ID from the Details section.
5. Go to Support API Explorer .
6. Locate and click the /hosts/{hostId}/commands/exitMaintenanceMode endpoint for HostsResource API to view the API parameters.
7. Click Try it out.
8. Enter the ID of your host in the hostId field.
9. Click Execute.
10. Verify that the maintenance mode status is cleared for the host by checking the Server response code.

    The operation is successful if the API response code is 200.

If you need any guidance during this process, contact Cloudera support for further assistance.

Technical Service Bulletins

TSB 2022-597: Cloudera Manager Event server does not clean up old events

The Event Server in Cloudera Manager (CM) does not clean up old events from its index, which can fill up the disk. This leads to wrong “Event Store Size” health checks.

Component affected:

- Event Server

Products affected:

- Cloudera Data Platform (CDP) Private Cloud Base
- CDP Public Cloud

Releases affected:

- CDP Public Cloud 7.2.14 (CM 7.6.0), and 7.2.15 (CM 7.6.2)
- CDP Private Cloud Base 7.1.7 Service Pack (SP) 1 (CM 7.6.1)

Users affected:

- Users who have Event Server running

Impact:

- Event Server’s index fills up the space on the used disk eventually.

Action required

Patch: Please contact support for a patch to address this issue.

- **Workaround**
  Suggested workaround instructions:

1. Stop the Event Server.
2. Check path for Event Server's index [eventserver_index_dir] in Cloudera Manager.
3. Archive /v4 folder in this path*.
   a. Compress the v4 folder using the following command:

   ```
   tar -czvf event_archive.tar.gz ${eventserver_index_dir}/v4
   ```

   b. Copy the archived version to an external disk.
   c. Remove the ${eventserver_index_dir}/v4 folder.
4. Start the Event Server**.

   *The archived version can be restored, by archiving the current index as described above, and extracting the archived version with the following steps:

   a. Stop the Event Server.
   b. Copy event_archive.tar.gz to ${eventserver_index_dir}.
   c. Extract event_archive.tar.gz using

   ```
   tar -xvf event_archive.tar.gz
   ```

   The extracted v4 folder should be under ${eventserver_index_dir}.
   d. Start the Event Server.***

   ** After the Event Server is restarted a new index is built, which cannot be merged with the previously archived index, if that is being restored.

   *** After the archived index is restored, the Event Server will continue to build that index with the new events.
5. Delete the Event Server's index which is under /var/lib/cloudera-scm-eventserver/v4 by default, can be changed using eventserver_index_dir parameter which is without the v4 subfolder.
6. Restart the Event Server.

Monitoring:

- CM by default has thresholds to monitor the Event Server space using [*eventserver_index_directory_free_space_percentage_thresholds*] parameter.

  You can adjust these as well by following the Cloudera Manager documentation.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: TSB 2022-597: Cloudera Manager Event server does not clean up old events