Cloudera Manager 7.6.0

Release Notes

Date published: 2020-11-30
Date modified: 2022-02-24

CLOUDERA

https://docs.cloudera.com/

Legal Notice

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

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA 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.

Contents

Cloudera Manager 7.6.0 Release Notes

Known issues, fixed issues and new features for Cloudera Manager 7.6.0.

What's New in Cloudera Manager 7.6.0

New features and changed behavior for Cloudera Manager 7.6.0.

Track swap rate vs. total swap used

Added a new health test that tracks swap rate, to reveal information that swap usage alone does not convey. The health test can be enabled by configuring Swap Memory Rate Thresholds.

Zeppelin support for Livy 3 / Spark3 interpreter

Added an additional Livy 3 (for spark3) interpreter. There was a need for an interpreter to run Spark 3 jobs. If a cluster had both Spark2 and Spark3 configured, Zeppelin user could use only one.

Toggle metrics collection without a restart

For most roles, changes to configuration parameters "Enable Metric Collection" and "Metric Filter" can now also be applied with a configuration refresh, without restarting the role.

Support metrics collection from secure endpoints

Custom Service Descriptors can specify that the Cloudera Manager Agent host certificate is to be used for the purpose of TLS client verification when collecting metrics. Existing Custom Service Descriptors are unaffected by the change.

Streams Messaging Manager should authenticate to Streams Replication Manager Service

Streams Messaging Manager now automatically configures Basic Authentication when connecting to Streams Replication Manager and the service dependency based auto-configuration is in use. For manual Streams Replication Manager connectivity configurations, Basic Auth configurations were added (Streams Replication Manager Basic Authentication, Streams Replication Manager Basic Authentication Username, Streams Replication Manager Basic Authentication Password).

Agent should run SS command only over IPv4

Cloudera Manager Agent uses 'ss' command to detect port conflicts. This command fails with a segmentation fault when IPv6 is disabled on the host and causes error messages to flood in the logs. This fix resolves the issue by limiting 'ss' to obtain only IPv4 info. for port detection.

Add support for JSON schema type in the registry configuration template

Schema Registry now supports Avro and JSON schemas.

Make Streams Messaging Manager Cache-Control part of default Streams Messaging Manager REST Server API's responses' headers

New Streams Messaging Manager configuration property has been added: "cache.control.http.response.header.value". This configuration allows you to configure the Cache-Control header's value for certain endpoints. Configure it in the following key-value like fashion:

- The key is the path prefix to the endpoints where the Cache-Control header should be added.
- The value is the value of Cache-Control header. In order to turn off functionalities provided by the Cache-Control header just delete the entries, or set the value to ""no-store"". To disable caching, set the value of the above mentioned configurations to "no-store"

Co-located Kafka configuration uses full Streams Replication Manager principal name instead of short

When the Streams Replication Manager Co-located Kafka Cluster Alias configuration is used to auto-configure the connection to the co-located Kafka cluster, and Kerberos is enabled, the JAAS configuration is dynamically generated on each host. As a result, you can now use the service dependency method to define a Kerberos enabled co-located cluster.

Create Streams Replication Manager Service Basic Auth service user automatically, export it with a dependency extension

Streams Replication Manager automatically creates a Basic Authentication credential for co-located services (users can change the credentials using Streams Replication Manager Service Co-Located Service Username and Streams Replication Manager Service Co-Located Service User Password). When Basic Authentication is enabled, this user is automatically accepted by Streams Replication Manager Service. For more information, see Configuring Basic Authentication for the Streams Replication Manager Service.

Create default Ranger repo for both DataHub and on-prem

A new service type was added to Ranger: kafka-connect. When Ranger is installed it creates the default policies for all services, so it also needs to create the default cm_kafka_connect policy which grants the default access.

Kafka Connect Ranger plugin setup and integration in CSD

If Ranger is enabled, the kafka-connect ranger plugin will be enabled and authorization will work through Ranger policies.

Add the emit.hearbeats.enabled config to Streams Replication Manager Driver

As a result of the rebase to Kafka 2.8 (KAFKA-10710), an improvement is introduced in connection with heartbeat emission. From now on you can fine tune your deployment and fully deactivate any unnecessary replications that are set up by default by configuring heartbeat emission. This can help with minimizing any performance overhead caused by unnecessary replications. To support this change, an improvement was made for the Streams Replication Manager service in Cloudera Manager. A dedicated configuration property, Enable Heartbeats, is introduced. You can use this property to configure emit.heartbeats.enabled on a global level directly in Cloudera Manager. Replication level overrides are still supported. This can be done by adding emit.heartbeats.enabled with a valid replication prefix to Streams Replication Manager's Replication Configs. For more information on configuring heartbeat emission, see Configuring Streams Replication Manager Driver heartbeat emission.

Implement Update Certificate Screen

A new Update Auto-TLS Truststore Certificate dialog box has been added to the Cloudera Manager Security page. You can use that dialog box to replace certificates in the truststore when Auto-TLS is enabled.

Add "http.metrics.reporter.filter" config to Kafka Csd

"kafka.http.metrics.reporter.exclude.filter": Complies the regex that is provided in the Config, and metrics that match the regex won't be reported by Cloudera Manager, and because of this won't be shown by Streams Messaging Manager either. The upstream-compatible JMX names are used for this filtering. Suggested default: "^kafka.log.Log.*". These metrics are not shown in Cloudera Manager or Streams Messaging Manager by default.

Enable setting offset in Schema Registry DB

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

New UpdateGlobalTruststore command in Cloudera Manager that can replace certificates in the truststore when Auto-TLS is enabled

Once Auto-TLS is enabled, Cloudera Manager lacked a way to let users add, remove, or replace certificates from the truststore. Customers may want to manage the CA certificates across all cluster nodes. You can now run a Cloudera Manager API call to upload certificates.

In order to run this command, call the new UpdateGlobalTruststore Cloudera Manager API command under ClouderaManagerResource and upload the new root certificate authority bundle. The new command updates the global truststore files and distributes them to each agent host. Ensure that you upload the entire truststore that contains all the root certificate authorities you want Cloudera Manager to be aware of.

Restart the service or cluster after running the command. The Cloudera Manager Agents restart after the command finishes running.

Support Cruise Control metric reporter in Kafka

Cruise Control introduces a new metrics reporter in addition to the existing "Cloudera Manager metrics reporter". The "Cruise Control metrics reporter" can be selected using the metric.reporter configuration property of Cruise Control. The upgraded clusters are going to use the "CM metrics reporter" by default, but the newly created ones will have the "Cruise Control metrics reporter". This parameter can be modified manually.

When Ranger authentication is enabled, then the Cruise Control metrics reporter topic and principal names are specific. If any of them changed, the Ranger policy has to be modified too. This is also needed when the policy's details are changed.

Cloudera Manager should display the Knox URL for Oozie UI when Knox is enabled

When the Knox gateway is available on the cluster and its discovery is enabled for Oozie then the Web UI link of Oozie through Knox will appear among the direct links.

Zookeeper SSL/TLS support for Cruise Control

Cruise Control introduced the ability to communicate with ZooKeeper through a secured TLS channel. Cruise Control uses secure communication with ZooKeeper automatically when TLS is enabled on the cluster.

Implement a new checkbox for Oozie to disable the Oozie UI

A new checkbox is implemented on the Oozie configuration page that can be used to turn off the Oozie UI completely. This means that none of these Oozie UI resources will be available. If you are concerned about JQuery vulnerabilities, that cannot be fixed in the short term, you can use this feature to get rid of these by not exposing the Oozie UI.

Cloudera Manager now supports rolling restarts of HA-enabled Schema Registry

The Schema Registry service can now be restarted using rolling restart in Cloudera Manager.

Include HBCK metrics for HBase

The following HBase metrics have been added: - orphan_regions_on_regionserver - orphan_regions_on_filesystem - inconsistent_regions - region_holes - regin_overlaps - unknown_server_regions - empty_region_info_regions

Allow users to access Kafka External Accounts in public cloud

The Limited Cluster Administrator role now has permission to access the External Accounts page, and manage a restricted set of external accounts (limited to the Kafka group).

Fixed Issues in Cloudera Manager 7.6.0

Fixed issues in Cloudera Manager 7.6.0

Cloudera Bug: OPSAPS-23472: Fix inaccuracies in Report Manager quota cache

Due to the inaccuracy of the HDFS quotas, the quota cache has been disabled. The modified quota will appear in the usage reports after the new HDFS fsimage was processed.

Cloudera Bug: OPSAPS-60041: Change HBase Default Compaction Events rules

Cloudera Manager was not able to show HBase Compaction events under HBase/Quicklinks because the format of the compaction logs was changed in HBase, and Cloudera Manager events rules did not account for this change.

This issue is now fixed and users can use the quick links to access compaction logs.

Cloudera Bug: OPSAPS-60066: Zeppelin fails on first run.

Zeppelin fails on first run. If Zeppelin's home directory '/var/lib/zeppelin/' is deleted, new Zeppelin service creation fails because the created directory has root as its owner. This change configures eppelin to create the directory during service creation with the correct permissions.

Cloudera Bug: OPSAPS-61178: Knox does not discover Cruise Control REST Endpoint

If CruiseControl is deployed in a cluster, Knox could not add that service to the cdp-proxy-api topology. This has been fixed

Cloudera Bug: OPSAPS-61209: Alert publisher keeps logging "Connection Refused" for smtp server

Disabled email alerts by default in Alert Publisher, preventing exceptions from being logged when default mail server settings are unsuitable for the deployment. Email alerts, if desired, have to be enabled manually.

Cloudera Bug: OPSAPS-61233: MultiException after generating Knox SSO Token with empty TTL value

Knox's SSO token integration features require a secret to be generated and stored in the appropriate credential store for Knox. Due to previous upgrade-related issues, the 'knox_token_mac_key' Cloudera Manager parameter had to be removed, and that secret had to be generated/saved manually using the KnoxCLI. Without those manual steps, the 'Token Integration' feature did not work. This fix made the entire process automated.

Cloudera Bug: OPSAPS-61296: Solr Data folders should be readable only by the Solr user

Solr data directory permissions have changed so they are only readable by the owner and its group.

Cloudera Bug: OPSAPS-61323: HDFS Log Rotation Issues

Enabled purging older HDFS garbage collection log files with a maximum retention of 10 files and a maximum file size of 200MB.

Cloudera Bug: OPSAPS-61708: Dag scanner in Telemetry publisher can run into java.lang.NullPointerException under some race conditions.

The DAG scanner looks for files that are not processed and could be processed. If it finds multiple files during a particular run (it runs every minute), that could trigger a race condition leading to this issue. This is fixed.

Cloudera Bug: OPSAPS-61792: Cloudera Manager runs bundle collection process on only one Knox instance

Fixed an issue where diagnostic bundle only included Knox logs of one instance in the case of a cluster with an HA Knox configuration.

Cloudera Bug: OPSAPS-61799: Script to generate host certificates fails during 'Add Hosts' step when the JDK version has changed since the initial setup

A JDK upgrade causes the keytool binary to change locations. This is because the path to keytool is hard coded in frozen_config.ini when enabling Auto-TLS. This results in a failure to regenerate certificates. For the same reason, adding hosts to the cluster would fail. This change allows the keytool binary path to be verified upon reading and if it is not valid, it attempts to use the keytool located under the current JAVA_HOME.

Cloudera Bug: OPSAPS-61803: Remove Jetty version from error page for additional security

Removed the "Powered by Jetty" message displayed by the Cloudera Manager Event Server's Jetty error page for additional security.

Cloudera Bug: OPSAPS-61824: Updated the Atlas CSD to have ATLAS_HOOK created with `delete` cleanup policy

The default ATLAS_HOOK topic now has the "delete" cleanup policy.

Cloudera Bug: OPSAPS-61835: Improve handling of empty command arguments

Corruption of the arguments of one command in the Cloudera Manager database could prevent all running and future commands from progressing. Now, the corrupted command errors out if necessary, instead of being re-tried, and other commands are not affected.

Cloudera Bug: OPSAPS-61846: Restrict krb5.conf path only when managed by Cloudera Manager

Prior to this fix, when upgrading Cloudera Manager, in situations where a custom path for krb5 .conf is set via Advanced Configuration Snippets, Cloudera Manager would return an error in the

Cloudera Manager Admin Console and cause the upgrade to fail. The workaround was to set the krb5.conf path to either /etc/krb5.conf or any path under /etc/hadoop.

Going forward, when Cloudera Manager manages the Kerberos configuration file location i.e. "Manage krb5.conf through Cloudera Manager" setting is enabled, the expected paths for the file are either at /etc/krb5.conf or any path under /etc/hadoop. When this setting is not enabled, the user can set any path for the krb5.conf file.

Cloudera Bug: OPSAPS-61876: Add NiFi service to ServiceDiagnosticsCommand class

Added NiFI to the Service Diagnostic class. Now diagnostic data is collected for NiFi and included in a support bundle.

Cloudera Bug: OPSAPS-61965: Fix SAML SSO - VelocityEngine runtime failure

A user gets the following error upon login to Cloudera Manager when using SAML SSO: org.apache.velocity.exception.ResourceNotFoundException: Unable to find resource '/templates/saml2-post-binding.vm' The issue is been fixed.

Cloudera Bug: OPSAPS-62063: Need to make max file size configurable for support bundle generation

Generating the support bundle for Atlas failed due to the file size limit. You can now set the file size limit by setting the MAX_BUNDLE_SIZE parameter with appropriate value in the Atlas Service Environment Advanced Configuration Snippet.

Cloudera Bug: OPSAPS-62200: The Zeppelin service cannot be upgraded: the Zeppelin service type is not supported for versions between CDH 7.2.3 and CDH 7.2.12

Removed the "SPARK_ON_YARN" service dependencies, and added support for Livy 3/Spark 3.

Cloudera Bug: OPSAPS-62242: spark.yarn.historyServer.address not getting updated

Do not overwrite spark.yarn.historyServer.address if specified in an Advanced Configuration Snippet.

Cloudera Bug: OPSAPS-62296: Fix label for Knox Gateway UI link

There has been an issue where the "Knox Gateway UI" link from the service page of a service with Knox SSO enabled, in the Cloudera Manager Admin Console, has redirected to the Knox service page in the Cloudera Manager Admin Console. The link now redirects to the Knox Gateway UI as expected.

Known Issues in Cloudera Manager 7.6.0

Known issues in Cloudera Manager 7.6.0

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.

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

Service Pack in Cloudera Manager 7.6.0

You can review the list of CDP Public Cloud hotfixes rolled into Cloudera Manager 7.6.0. This will help you to verify if a hotfix provided to you on a previous CDP Public Cloud release was included in this release.

- HOTREQ-796 - HOTFIX request for OPSAPS-63158
- HOTREQ-607 - IDBroker instances experiencing high CPU utilization
- HOTREQ-584 - OPSAPS-62583 : Chive: add an option to ignore certain partitionParameters when comparing and improve location comparison