Cloudera Manager 7.4.3

Release Notes

Date published: 2020-11-30
Date modified: 2021-09-09

CLOUDERA

https://docs.cloudera.com/

Legal Notice

Contents

Cloudera Manager 7.4.3 Release Notes

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

Fixed Issues in Cloudera Manager 7.4.3

Fixed issues in Cloudera Manager 7.4.3

Cloudera Bug: OPSAPS-57469: Hardcoded parcel path causes initialization failure when non-standard location is used

Added support for custom parcel location in Atlas

Cloudera Bug: OPSAPS-59221: Cruise Control default number of metric windows (1) is too low

Fixes the default values of the num.broker.metrics.windows and num.partition.metrics.windows configuration properties that are needed for Cruise Control to work properly.

Cloudera Bug: OPSAPS-59264: Use kafka_principal_name variable where "kafka" principal is hardcoded in Atlas CSD

Added support to use the actual Kafka principal in Atlas.

Cloudera Bug: OPSAPS-59972: Disable the TRACE method on all HTTP ports

In Streams Messaging Manager and Schema Registry the allowed HTTP methods have been changed to GET, POST, PUT, DELETE, HEAD, and OPTIONS.

Cloudera Bug: OPSAPS-59993: Disable admin port in Schema Registry

The admin port is disabled in Schema Registry (secure port: 7791, not secure port: 7789).

Ranger will not start on Ubuntu20 & Redhat8.2

When starting Ranger you may see the error: "/usr/bin/python: No such file or directory"'.By default, Ubuntu20 and Redhat8 don’t have a default Python version configured. Instead it gives the user a choice to install, configure and run a specific Python version. so to run Python you need to explicitly type python3 or python2. User can configure the unversioned python command and set the default version. Some services in Cloudera Manager like Hue, Ranger expect to find the python command in the system’s path(/usr/bin/python). With this fix, Cloudera Manager will set python2 as the default version on Redhat8 and Ubuntu20 hosts while installing Cloudera Manager agents. No additional user actions are required.

Cloudera Bug: OPSAPS-60391: Import Deployment throws NPE on users without authRoles

When importing a deployment, you may encounter a failure, with the logs containing a NullPointerException. This is potentially caused by user entries in the deployment missing an authRole field. One reason this may happen is when external authentication and authorization (such as via LDAP or SAML) is being used. As a workaround, you can edit the JSON of the deployment to delete these users, then attempt to import the deployment again. This has now been fixed.

Cloudera Bug: OPSAPS-60514: Need to add ranger jdbc test connection parameter when user is using Oracle db

Added parameter for test query in Ranger to configure the Ranger JPA jdbc test connection according to the database type configured.

Cloudera Bug: OPSAPS-60532: Ranger plugin's audit could not authenticate to zookeeper in Streams Messaging Manager and Schema Registry

Generate JAAS config for Streams Messaging Manager and SR to SASL authenticate to zookeeper.

Cloudera Bug: OPSAPS-60537: Permission on warehouse/tablespace/external/hive after installing Impala

CDP clusters that have Impala service installed have HDFS ACLs for the external warehouse directory set to default:other::rwx. This means that any application or user which does direct

HDFS operations on external tables located in external an warehouse directory will have all the permissions to read, write or list any files.

If there is Impala service installed, a application (like spark) has read,write, execute privileges on HDFS directories for all the external tables. If Impala service is not installed the external warehouse directory permissions are only read/write enabled for users of the hive group.

If there are applications other than Impala or Hive which need HDFS access to the external tables, the HDFS permissions for such applications should be explicitly allowed by Ranger.

Additional notes: External Hive warehouse directory is created with default:other::rwx permissions when the Impala service is installed in Cloudera Manager. This can cause new subdirectories within external warehouse directory (databases and tables) to have rwx permissions on HDFS. With this patch, the external warehouse is created with default:other::--- permissions. In the case of a non-Hive or non-Impala user needing the subdirectory permissions they should be managed by Ranger.

Cloudera Bug: OPSAPS-60559: Fix omid options

Incorrect -Xmx settings when Cloudera Manager starts Omid has been fixed.

Cloudera Bug: OPSAPS-60562: Add Raz claim to the DT tokens issued by Knox

Additional audience claim "raz" is now added to tokens issues by IDBroker.

Cloudera Bug: OPSAPS-60601: Improper parameter passing in Streams Replication Manager CSD

The generated configuration for Streams Replication Manager could become corrupted, either when the security configuration uses JAAS properties AND at least one of the JAAS secrets defined in a Kafka External Account contained a space. in this case, the Streams Replication Manager client won't be able to connect to the respective cluster - or at least one replication flow's target is an unsecured cluster; in this case, replication won't start. The issue has been fixed. JAAS secrets can be used in Kafka External Accounts, and an unsecure cluster can be the destination of a replication flow.

Cloudera Bug: OPSAPS-60630: Extension of OPSAPS-59969 - Add role status and health summary to api/v43/hosts

The HostResource/{hosts} API call of the Cloudera Manager API now contains each role's health and status if the view is set to FULL or FULL_WITH_HEALTH_CHECK_EXPLANATION.

Cloudera Bug: OPSAPS-60648: Cruise control access log is in the process's folder1 instead of /var/log/cruisecontrol

The access.log of Cruise Control has been moved to the log directory where other logs of the service can be found.

Cloudera Bug: OPSAPS-60660: Expose Knox Token TTL for token generation in the Cloudera Manager Admin Console

End-users could not configure a Knox token TTL in the Cloudera Manager Admin Console.

Cloudera Bug: OPSAPS-60663: Kudu version is incorrect in Cloudera Manager

Fixed the Kudu version property in KUDU CSD.

Cloudera Bug: OPSAPS-60685: HBase health check problem caused by missing opentelemetry

Included opentelemetry jars in Cloudera Runtime libraries to fix an HBase healtcheck problem.

Cloudera Bug: OPSAPS-60738: Error message dis;lays when trying enable/disable HDFS HA or add an HDFS nameservice

Adding a Nameservice or Enabling /Disabling HDFS HA failed when Cloudera Manager is configured using Knox. This issue is fixed.

Cloudera Bug: OPSAPS-60766: Stack Area charts don't put the layers on top of each other

Fixed an issue where stack charts are now stacked on top of each other rather than placed in front of each other.

Cloudera Bug: OPSAPS-60775: Streams Replication Manager does not generate external account configurations for the Streams Replication Manager Service

The Streams Replication Manager Service configuration now contains the Kafka External Accounts configuration, enabling Streams Replication Manager Service to access Kafka clusters defined through External Accounts.

Cloudera Bug: OPSAPS-60803: Use the security protocol from Kafka dependency extension in CruiseControl

The security.protocol can be overridden using the Cruise Control Server Advanced Configuration Snippet (Safety Valve) for cruisecontrol.properties

Cloudera Bug: OPSAPS-61010:Cloudera Manager Redhat8.2 IBM PowerPC agent installation failure on libboost_python3

Added boost-python3 as a dependency to the Cloudera Manager Agent RPM

Cloudera Bug: OPSAPS-61141: CDP environment fails with Failed to create HDFS directory on Azure with RAZ enabled

after creating new RangerRaz-identity

```
<cdp-env-name>-RangerRazIdentity
          Storage Account<cdp-env-name>
          Storage Blob Data Owner
          Storage Blob Delegator
          we could able to bring up RAZ on Azure
```

Cloudera Bug: OPSAPS-61215: Add support for telemetry publisher to read logs from GCS

Telemetry publisher is now able to read the logs from GCS storage..

Cloudera Bug: OPSAPS-61289: Fix minimumMemory requirement for OMID TS during upgrade to 7.2.11

The minimum memory requirement for Omid TSO is now updated during upgrade.

Cloudera Bug: OPSAPS-61362: Atlas import-kafka.sh fails with Failed to create new KafkaAdminClient

Atlas auth-to-local rules generation has been hanced to handle escaping a comma from the rules.

Cloudera Bug: OPSAPS-61474: Knox token API call on homepage topology failed with 404

Knox's data/applications folder gets recreated every time Knox starts.

What's New in Cloudera Manager 7.4.3

New features and changed behavior for Cloudera Manager 7.4.3.

Miscellaneous

Add "krb5.conf" location configuration into Cloudera Manager

Cloudera Manager now allows a user to set a path for the Kerberos Configuration file, krb5.conf. The user can use the AdminstrationSettingskrb5.conf file path configuration field in Cloudera Manager to set the path. The default is set to /etc/krb5.conf . For now, the valid paths are limited to /etc/hadoop/* or /etc/krb5.conf only. Note: If the configuration to 'Manage Krb5.conf using Cloudera Manager' property is also set by the Cloudera Manager admin user, it will cause Kerberos Staleness because Cloudera Manager is responsible for handling Kerberos configurations.

Expose extra command line arguments for Auto-TLS

The GenerateCMCA API now provides access to an additional arguments:

- additionalArguments: This parameter can be provided to pass additional parameters for internal CA certificates.
- subjectaltnames: Using this parameter, a list of Subject Alt Names can be provided for each host during certificate generation.

For more information, please refer to the Cloudera Manager API documentation: https://archive.cloudera.com/cm-public/7.4.3/generic/jar/cm_api/apidocs/json_ApiGenerateCmcaArguments.html

Upgrade Jetty to 9.4.latest

Jetty Server version has been upgraded to 9.4.35.v20201120, which fixes numerous security vulnerabilities .

Enable mTLS for Ratis in Ozone

New configuration parameters have been added for Ozone to separate Data Node Ratis admin/server traffic from clients:

- dfs.container.ratis.admin.port, defaults to 9857
- dfs.container.ratis.server.port, defaults to 9856

A new TLS configuration for Ratis in Data Node and OM has also been added :

- hdds.grpc.tls.enabled, defaults to false

Make Control settings for Data Hub part of the CSD

The Cruise-Control service has new configurable properties related to self-healing.

Add support for custom ZooKeeper principal in Ranger

Added support in Ranger to use custom zookeeper principal for communication with Solr service.

Add option to CollectDiagnosticDataArguments API to force diagnostic bundle upload

This is a new feature in the Cloudera Manager API. The feature adds a new parameter to ApiC ollectDiagnosticDataArguments which adds the ability to force the generated diagnostic bundle to be uploaded to Cloudera. With the previous behaviour, this upload feature was controlled by the PHONE_HOME parameter of Cloudera Manager which is still in use, but can be overridden by this new parameter. The feature is backward compatible. The parameter is optional and-if the parameter is missing, the old behaviour takes place according to the PHONE_HOME setting.

Add config 'ozone.scm.ratis.storage.dir' to Ozone

A new configuration property for Ozone, ozone.scm.ratis.storage.dir, has been added.

Add Hive configuration to set cipher suites for Hive WebUi and HS2

With the fix, Hive WebUI SSL Cipher Suites can be configured, allowing the Web UI and HiveServer2 to work with TLS security on a FIPS-enabled operating system.

Atlas

New feature to enable Atlas Hook Spooling

Atlas hook spooling feature is now available and configurable which can be enabled / disabled from configurations.

Atlas : Add hadoop-metrics2.properties in conf directory

Added configuration of hadoop-metrics2.properties in Atlas.

Add support for custom Zookeeper principal in Atlas

Added support to use actual zookeeper principal in Atlas.

Kafka

New configuration parameter for Kafka Connect role

A new property called include.connector.context is added for the KafkaConnect role, that is enabled by default. If it is enabled, additional connector context information is added to Kafka connect file logs.

Add Kafka health test for RequestHandlerAvgIdlePercent and NetworkProcessorAvgIdlePercent

Added two new health tests for Kafka: - If NetworkProcessorAvgIdlePercent is below the threshold 0.3, we advise the user to increase num.network.threads and make the broker health concerning

- If RequestHandlerAvgIdlePercent is below the threshold 0.3, we advise 1the user to increase num.io.threads and make the broker health concerning

Knox

Add SSL support to Knox Gateway DB

SSL-related connection properties were not exposed in the Cloudera Manager Admin Console for the KNOX GATEWAY Database.

Add Database support for Knox

Previously, the KNOX_GATEWAY role lacked database support which is needed for the Knox Token generation feature (instead of storing the tokens in Zookeeper or in keystores on the local file system)

Knox autodiscovery for SQLStreamBuilder Service

SQLStreamBuilder was not auto-discovered by Knox. Users had to add the service manually into Knox topologies if they wanted to use SSB in CDP. From now on, auto service-discovery for SQL Stream Builder is available. Any previous manual configurations must be reverted.

Knox principal is not overridable in Streams Messaging Manager and Schema Registry

Custom Knox principal can be set for Schema Registry and Streams Messaging Manager by setting the knox_principal_name property in the Schema Registry Server Advanced Configuration Snippet (Safety Valve) for registry.yaml or the Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml

Ranger

Token created by user should be exchanged using token exhange API without requiring admin privileges in the environment

A new Ranger policy was created to allow public access to the new tokenexchange Knox topology.

New introduce Max Retention Days configuration parameter for Ranger audits

Users can now update the Solr document expiry ranger.audit.solr.config.ttl and ranger.audit.solr.config.delete.trigger parametersin Cloudera Manager for Ranger configurations and refresh configurations to get the Solr collection for Ranger audits updated with ttl and delete trigger.

Streams Messaging Manager

Streams Messaging Manager is now configurable to allow custom list of ciphers and SSL protocols

Streams Messaging Manager now offers the following1 configurations to customize the SSL configurations of the Streams Messaging Manager Server: streams.messaging.manager.ssl.supportedCipherSuites, streams.messaging.manager.ssl.excludedCipherSuites, streams.messaging.manager.ssl.supportedProtocols, streams.messaging.manager.ssl.excludedProtocols.

In a FIPS enabled environment, to support access from a browser, excludedCipherSuites should be updated to allow ciphers ending with "_SHA".

Streams Replication Manager

Streams Replication Manager should be configurable to allow custom list of ciphers and SSL protocols

Streams Replication Manager now offers the following configurations to customize the SSL configurations of Streams Replication Manager Service: streams.replication.manager.ssl.supportedCipherSuites, streams.replication.manager.ssl.excludedCipherSuites, streams.replication.manager.ssl.supportedProtocols, streams.replication.manager.ssl.excludedProtocols.

Introduce health test for Streams Replication Manager service

New health tests were introduced to the SRM service role which describes the state of the SRM service. With the help of these, when the Streams Application inside the SRM services goes to ERROR state or loses connectivity with the target Kafka Cluster, SRM tries to restart it, and Cloudera Manager shows that SRM Service is not functional.

Known Issues in Cloudera Manager 7.4.3

This topic describes known issues and workarounds for Cloudera Manager.

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

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-61825: Refreshing the cluster using the Refresh Cluster option fails with an error.

You may see the following error if you try to refresh your cluster from Cloudera Manager UI using the Refresh Cluster option:

```
com.cloudera.cmf.command.CmdExecException: com.cloudera.cmf.serv
ice.CommandException: No command 'UpdateSolrConfigSet' found for
 role 'DbRole{id=21, name=RANGER-RANGER_ADMIN-1, hostName=xxxxxx-
xxxxxx-1.xxxxxx-xxxxxx.root.hwx.site}'
```

Avoid using the cluster-level Refresh Cluster command from Cloudera Manager. Instead, use the role-level Refresh command for the individual roles available in the services from the Actions dropdown menu.

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.