Cloudera Manager 7.2.1

Release Notes

Date published: 2020-05-28
Date modified: 2020-05-29

https://docs.cloudera.com/

Legal Notice

Contents

Cloudera Manager 7.2.1 Release Notes

Known Issues, Fixed Issues and New features for Cloudera Manager and CDP Data Center.

What's New in Cloudera Manager 7.2.1

New features and changed behavior in Cloudera Manager. (CDP Public Cloud)

API for regenerating host certificates can now rotate certificates without using SSH

Using the Cloudera Manager API, host certificates can now be rotated without providing SSH credentials, provided that the host has a healthy heartbeat.

The endpoint for generating host certificates is:

```
/hosts/{hostId}/commands/generateHostCerts
```

Changed Behavior

New default configuration values in Kafka

The default value of the Kafka zookeeper.session.timeout.ms configuration property for the has been changed from 6000 to 18000 and the default value for replica.lag.time.max.ms has been changed from 10000 to 30000.

Cruise Control now infers Kerberos and SSL settings

The security.protocol configuration property of the Cruise Control service has been removed, and the value is now inferred from the Kafka broker configuration.

The Hue load balancer now rebalances users to the least used Hue instance

The Hue load balancer now generates a new cookie value when the Load Balancer role is restarted.

Admin port has been removed from Solr configurations

The Admin port has been removed from Solr configurations for clusters running CDH 6.x or Cloudera Runtime 7.x because it is no longer used.

Changing port numbers to non-ephemeral ports

Kafka Connect default ports are now non-ephemeral ports.

The default wal provider has been changed from AsyncFSWal to Filesystem

A new configuration property, hbase.wal.regiongrouping.delegate.provider has been added to the HBase configuration properties.

Hue in Data Hub has the Solr Data Lake as a dependency instead of the Data Hub

It is now possible to configure Hue's Solr dependency correctly to support building Solr applications in Data Hub clusters using the Hue interface.

Fixed Issues in Cloudera Manager 7.2.1

This topic lists the issues that have been fixed in Cloudera Manager since the previous release of Cloudera Manager.

Cloudera Bug: OPSAPS-56286: Schema Registry Health Check broken with multiple instances

Health State is now fixed when multiple instances of the Schema Registry role are deployed in the cluster.

Cloudera Bug: OPSAPS-56335: Add HBase to list of dependencies of HueServiceDependencyValidator for Hue-HBase

Currently, the Hue service when deployed in a compute cluster requires Impala or HiveServer2 to be present in the compute cluster. This hindered installation of HBase with Hue in HBase compute clusters . This fix adds HBase to the list and enables installation of Hue for HBase without the need to install Impala and HiveServer 2.

Cloudera Bug: OPSAPS-57317: Apache Ranger user cannot send requests to SchemaRegistry

Added the user 'rangerlookup' to the default list of users when creating the schemaregistry policy in Apache Ranger.

Cloudera Bug: OPSAPS-57113: ssl.principal.mapping.rules property configured in the Cloudera Manager Admin Console is not correctly propagated to Kafka brokers

The Kafka SSL Advanced Configuration Snippet now propagates configuration values containing dollar signs correctly.

Cloudera Bug: OPSAPS-57067: Yarn stuck in stale configuration on Data Hub cluster after deployment

YARN no longer reports the following configuration as stale when running Data Hub on Azure: yarn.cluster.scaling.recommendation.enable.

Cloudera Bug: OPSAPS-54954: Cloudera Manager - Streams Replication Manager's Replication Configurations wording goes off page

Streams Replication Manager's Replication Configuration help text now displays correctly.

Cloudera Bug: OPSAPS-57014: Log rotation does not remove old logs

Hive now correctly rotate its logs (by discarding the older logs).

TSB-431: Cloudera Manager 6.x issue with the service role Resume

If a selected service role on a node is restarted and fails, and the customer clicks the "Resume" button in Cloudera Manager, the service role on all of the nodes will be restarted concurrently.

Workaround:

- Instead of performing a restart we recommend performing a stop/start of the services.
- The issue is addressed in Cloudera Manager 7.2.1 and higher versions

For more information about this issue, see the corresponding Knowledge article:Cloudera Customer Advisory: Cloudera Manager 6.x issue with service role Resume

Known Issues in Cloudera Manager 7.2.1

Learn about the known issues in Cloudera Manager 7.2.1, the impact or changes to the functionality, and the workaround.

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"
```

Technical Service Bulletins (TSB)

TSB 2021-472: Customer Advisory for Navigator Metadata Server startup issue

If the Navigator Metadata Server is executing purge, and the clean up process is interrupted, the Navigator Metadata Server will not be able to restart.

Impact

Navigator Metadata Server cannot be restarted if the process is killed or crashes during executing a purge. Error message:

[Update NAV_EXTRACTOR_STATUS set ENABLED_FOR_NEXT_EXTRACTION = 'true']; SQL state [72000]; error code [12899]; ORA-12899: value too large for column "NAVMS"."NAV_EXTRACTOR_STATUS"."ENABLED_FOR_NEXT_EXTRACTION" (actual: 4, maximum: 1; nested exception is java.sql.SQLException: ORA-12899: value too large for column "NAVMS"."NAV_EXTRACTOR_STATUS"."ENABLED_FOR_NEXT_EXTRACTION" (actual: 4, maximum: 1)

Action required

- Upgrade:
  - Cloudera Manager 6.3.4: Request a patch (PATCH-4489).
  - Cloudera Manager 7.2.1, 7.2.2, 7.2.3, 7.2.4, 7.2.5, 7.2.6 and 7.3.0: Upgrade to a Cloudera Manager version containing the fix.
- Workaround:
  1. Log in to the Navigator Metadata Server database.
  2. Update NAV_MAINTENANCE_HISTORY set STATUS = "INCOMPLETE" where STATUS like 'IN_PROGRESS'.
  3. Update NAV_EXTRACTOR_STATUS set ENABLED_FOR_NEXT_EXTRACTION = 1 where ENABLED_FOR_NEXT_EXTRACTION = 0.
  4. NMS is able to start and extractors are enabled.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article:

Cloudera Customer Advisory-472: Navigator Metadata Server startup issue

TSB 2021-481: Lineage is not extracted with Cloudera Manager 7.2.x and 7.3.1 managing CDH6 or CDH5

Cloudera Manager - Upgrade to Guava 28.1 to avoid CVE-2018-10237 triggered a Guava method version mismatch causing an exception in Navigator Metadata Server. As a result no new lineage and metadata is extracted with Cloudera Manager 7.2.4 and later with CDH6 and CDH5.

Impact

Lineage and metadata are no longer updated in Cloudera Navigator after upgrading to Cloudera Manager 7.2.x or Cloudera Manager 7.3.1 when managing CDH5 or CDH6.

Action required

Upgrade to the patched release of CM 7.3.1 available as PATCH-4822, or to an upcoming version later than 7.3.1. After upgrade, existing entities will have metadata extracted when extraction resumes and no lineage will be permanently lost.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article:

Cloudera Customer Advisory-481: Lineage is not extracted with Cloudera Manager 7.2.x and 7.3.1 managing CDH 6 or CDH 5

TSB 2021-491: Authorization Bypass in Cloudera Manager (CVE-2021-30132/CVE-2021-32483

Cloudera Manager (CM) 7.4.0 and earlier versions have incorrect Access Control in place for certain endpoints. A user who has a knowledge to the direct path of a resource or a URL to call a particular function, can access it without having the proper role granted. The vulnerable endpoints were CVE-2021-30132 /cmf/alerts/config?task= and CVE-2021-32483 /cmf/views/view?viewName=.

CVE

- CVE-2021-30132
  - Alerts config - 4.3 (Medium)
  - CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:L/I:N/A:N
- CVE-2021-32483
  - Views - 4.3 (Medium)
  - CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:L/I:N/A:N

Impact

A user with read only privilege is able to see configuration information in the UI.

Action required

Upgrade to a version containing the fix.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: TSB 2021-491: Authorization Bypass in Cloudera Manager (CVE-2021-30132 / CVE-2021-32483)

TSB 2022-507 Certificate expiry issue in CDP

The Transport Layer Security (TLS) keystore needs to be manually rotated due to an issue with certificate rotation.

The Root Cause Analysis is that the keystore path of the Cloudera Manager (CM) server is set to a directory based on the non-FQDN (Fully Qualified Domain Name) of the CM server. However, the certificate rotation on a directory happens based on the FQDN of the CM server. This results in a situation in which the keystore of the CM server does not get updated.

Impact

The clusters could experience downtime.

Action required

- Workaround if the certificates have not yet expired:
  1. Back up the existing host keystore from the directory based on the hostname of the CM server. Example:

     ```
     cp -R /etc/cloudera-scm-server/certs/hosts-key-store/example
     -datalake-1-master0/ /etc/cloudera-scm-server/certs/hosts-ke
     y-store/example-datalake-1-master0.backup
     ```

  2. Copy the keystore from a directory based on the FQDN of the CM server. Example:

     ```
     cp -Rf /etc/cloudera-scm-server/certs/hosts-key-store/exampl
     e-datalake-1-master0.domain.site/* /etc/cloudera-scm-server/
     certs/hosts-key-store/example-datalake-1-master0/
     ```

  3. Restart the CM server
  4. Confirm that OpenSSL now shows a certificate with the expected expiration time. Example:

     ```
     openssl s_client -connect $(grep "server_host" /etc/cloudera
     -scm-agent/config.ini | sed s/server_host=//):7182 </dev/nul
     l | openssl x509 -text -noout
     ```

  5. Repeat these steps after each host certificate rotation.

- Workaround if the certificates have already expired:
  1. You must run commands on each host with expired certificates to regenerate new ones.
  2. For each affected host (including the Cloudera Manager server host if necessary), let “<hos t_FQDN>” be the fully-qualified domain name of that host:
     a. Run the following command on the Cloudera Manager server host as root:

```
/opt/cloudera/cm-agent/bin/certmanager --location
/etc/cloudera-scm-server/certs gen_node_cert --rotate --o
utput=/tmp/<host_FQDN>.tar <host_FQDN>
```

     b. Copy /tmp/<host_FQDN>.tar to the affected host.
     c. Run the following commands on the affected host as root:

     - ```
       /opt/cloudera/cm-agent/bin/cm install_certs /tmp/<ho
       st_FQDN>.tar
       ```
     - ```
       chmod 755 /var/lib/cloudera-scm-agent/agent-cert/
       ```

  3. Restart Cloudera Manager by running the following command on the Cloudera Manager server host as root:

```
service cloudera-scm-server restart
```

  4. Restart the Knox service by running the following commands on the Cloudera Manager server host as any user, replacing “UpdateWithYourUser” and “UpdateWithYourCluste rName” with the workload user and cluster name, respectively:

     - ```
       WORKLOAD_USER="UpdateWithYourUser"
       ```
     - ```
       CM_SERVER="http://$(hostname -f):7180"
       ```
     - ```
       CM_API_VERSION=$(curl -s -L -k -u ${WORKLOAD_USER} -X GET
                    "${CM_SERVER}/api/version") && echo ${CM_API
       _VERSION}
       ```
     - ```
       CM_CLUSTER_NAME=<UpdateWithYourCusterName>
       ```
     - ```
       KNOX_SERVICE_NAME=$(curl -s -L -k -u ${WORKLOAD_USER} -X
        GET
                    "${CM_SERVER}/api/${CM_API_VERSION}/clust
       ers/${CM_CLUSTER_NAME}/services/" | awk -F
                    "[ |:|,]" '/name.*knox/ {print $(NF - 1 )}'
        | sed 's|"||g') && echo
                    ${KNOX_SERVICE_NAME}
       ```
     - ```
       curl -s -L -k -u ${WORKLOAD_USER} -X POST
                    "${CM_SERVER}/api/${CM_API_VERSION}/clusters
       /${CM_CLUSTER_NAME}/services/${KNOX_SERVICE_NAME}/comman
       ds/restart"
       ```

  5. Follow the steps in the above section: “Workaround if the certificates have not yet expired”

Knowledge article

For the latest update on this issue, please see the corresponding Knowledge article: TSB 2022-507: Certificate expiry issue in CDP

TSB 2021-488: Cloudera Manager is vulnerable to Cross-Site-Scripting attack (CVE-2021-29243 and CVE-2021-32482)

Cloudera Manager may be vulnerable to Cross-Site-Scripting vulnerabilities identified by CVE-2021-29243 and CVE-2021-32482. A remote attacker can exploit this vulnerability and execute malicious code in the affected application.

CVE

- CVE-2021-29243
- CVE-2021-32482

Impact

This is an XSS issue. An administrator could be tricked to click on a link that may expose certain information such as session cookies.

Action required

- **Upgrade (recommended)**

Upgrade to a version containing the fix.

- **Workaround**

None

Knowledge article

For the latest update on this issue see the corresponding Knowledge article:

TSB 2021-488: Cloudera Manager vulnerable to Cross-Site-Scripting attack (CVE-2021-29243 and CVE-2021-32482)