CDP Private Cloud Base 7.0.3

Release Notes

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

CLOUDERA

https://docs.cloudera.com/

Legal Notice

Contents

Cloudera Manager 7.0.3 Release Notes

What's New in Cloudera Manager 7.0.3

This topic describes new features in Cloudera Manager.

Apache Ranger

Apache Ranger provides auditing, authentication, and authorization functionality for your CDP - Data Center clusters. Apache Ranger provides a centralized framework for collecting access audit history and reporting data, including filtering on various parameters. Ranger enhances audit information obtained from Hadoop components and provides insights through this centralized reporting capability.

Apache Ranger also manages access control through a user interface that ensures consistent policy administration across CDP - Data Center components. Security administrators can define security policies at the database, table, column, and file levels, and can administer permissions for specific LDAP-based groups or individual users. Rules based on dynamic conditions such as time or geolocation can also be added to an existing policy rule. The Ranger authorization model is pluggable and can be easily extended to any data source using a service-based definition. For customers familiar with Cloudera Enterprise, Apache Ranger replaces the Sentry service.

Apache Atlas

Apache Atlas now provides governance for your data. Apache Atlas serves as a common metadata store that is designed to exchange metadata both within and outside of the Hadoop stack. Close integration of Atlas with Apache Ranger enables you to define, administer, and manage security and compliance policies consistently across all components of the Hadoop stack. For customers familiar with Cloudera Enterprise, Apache Atlas replaces Cloudera Navigator and also provides the following capabilities:

- Dynamic row filtering
- Dynamic column masking
- Attribute-based access control
- SparkSQL fine-grained access control

Solr, HBase and Kudu on Compute Clusters

Creation of Solr, HBase and Kudu services on Compute Clusters is now enabled.

LDAP authentication for Kafka clients

You can now configure LDAP to allow Kafka clients to authenticate using LDAP.

OPSAPS-53093

Backup and Disaster Recovery is now called Replication Manager.

To access replication functionality in Cloudera Manager Admin Console select Replication from the left navigation menu.

Upgrade Domains

Upgrade Domains enable faster cluster restarts, faster Cloudera Runtime upgrades, and seamless OS patching & hardware upgrades across large clusters. Upgrade Domains provide an alternative to the default HDFS block placement policy, distributing data across a set of hosts (potentially larger than a single rack) that Cloudera Manager can upgrade/restart at once without compromising service and data availability. When you select Upgrade Domains as

the block placement policy, you also assign an Upgrade Domain group to each DataNode host. The NameNode uses these groups to distribute blocks when writing data, and to orchestrate rolling restarts and upgrades. This feature is useful for very large clusters, or for clusters where rolling restarts happen frequently.

Cloudera Manager Upgrade Limitations

Upgrades from Cloudera Manager 5 or 6 to Cloudera Manager 7.x are not supported and will fail when Cloudera Manager server starts.

Core Configuration Service

The Core Configuration service allows you to create more types of clusters without having to include the HDFS service. Previously, the HDFS service was required in many cases even when data was not being stored in HDFS because some services like Sentry and Spark required cluster-wide configuration files that Cloudera Manager deploys within the HDFS service. The Core Configuration service provides this configuration in a standalone fashion and thus eliminates the need for an HDFS service for certain types of clusters where no HDFS storage is required (e.g. Kudu, Kafka, or ‘Compute’ clusters using exclusively object storage like S3 or ADLS). The Core Configuration service is also useful when creating a Compute cluster that accesses data on an HDFS service located in the Base cluster.

"Impala for Compute" and "Spark for Compute" no longer require HDFS. You can define the Core Configuration Service instead.

See Core Configuration Service

Metric Filtering

Metrics Filters allow you to limit the amount of metric data sent to the Cloudera Manager Service Monitor In large clusters, some services, such as Kudu, send a high volume of non-essential metrics data to the Service Monitor, which can overload it, causing gaps in the data reported from these metrics in charts/dashboards & metrics queries, and potentially limiting the ability for Cloudera Manager to effectively monitor cluster health . To mitigate this problem, you can configure Metric Filters that limit the amount of data sent to the Service Monitor and Host Monitor. You can configure Metric Filters for any service deployed in a cluster.

See Filtering Metrics

YARN Queue Manager and Capacity Scheduler

YARN Queue Manager is the queue management graphical user interface for Apache Hadoop YARN Capacity Scheduler. You can use the YARN Queue Manager to manage your cluster capacity using queues to balance resource requirements of multiple applications from various users. Using the YARN Queue Manager, you can set scheduler level-properties and queue-level properties. You can also view, sort, search, and filter queues. Queue Manager replaces Dynamic Resource pools (as used in CDH 5 and CDH 6 clusters). Capacity Scheduler is the new default scheduler for Cloudera Runtime 7 and higher.

HTTP Strict-Transport-Security

When TLS is enabled for the Cloudera Manager Admin Console, web requests now include the HTTP Strict-Transport-Security header. For more details about this header, see Strict-Transport-Security (Mozilla).

Ranger Service and Kafka

The Ranger service name for Kafka clusters is now configurable. The default (and initialized) value is cm_kafka.

Cloudera Manager Licensing

When the license key in Cloudera Manager expires, or the trial period expires, access to the Cloudera Manager Admin Console will be disabled. Cloudera Manager will still function, but users will be unable to interact with any features or their clusters from the Cloudera Manager Admin Console.

New Health Tests

- LDAP connections. The LDAP health check requires you to set a bind user to enable monitoring.
- Key Distribution Center (KDC) connections. The KDC health check requires Cloudera Manager Server to use Kerberos to enable monitoring.

New configuration parameters for Azure

Two new core-site configurations have been added to support delegation token collection on Azure cloud storage:

- fs.azure.identity.transformer.service.principal.substitution.list
- fs.azure.identity.transformer.service.principal.id

New Kafka Metric

A new metric has been added to the Kafka service for JVM Garbage Collection Rate: kafka_jvm_gc_runs.

New notification suppression parameters

Notification suppression parameters for role-level validators are now available.

Redaction in Cloudera Manager API

Previously redaction was opt-in through a JVM parameter, causing major security concerns. Customers relying on the API for backups now have a viable alternative that does not rely on exposing passwords via the API.

OPSAPS-51856, OPSAPS-52510: Single User Mode (SUM) is not supported in Cloudera Manager 7

Single User Mode is not supported for upgrades to Cloudera Manager 7.x.

Cloudera Manager User Interface Improvements

Cluster-level Configuration History

Configuration changes across all the services in a cluster are now shown in a single screen. The new configuration screen now has a search function and time-based filters.

Configuration Page Changes

- You can now toggle display of the filters on and off.
- When entering name/value pairs for environment Advanced Configuration Snippets has been enhanced with name and value fields.
- The Reason for change field is now populated automatically. You can override the field or add to the automatically-generated text.
- You can now use CNTRL + S to save configuration changes.

All Hosts Page

You can now toggle display of the filters on and off. There is a new refresh button and the page refreshes automatically every 90 seconds.

Global Search

The global search function (accessible from the left navigation menu) has been enhanced with improved sorting of results.

Enable Kerberos Wizard can now restart after errors

Previously, when user uses the Enable Kerberos wizard and there is an error with keytab retrieval, the wizard fails and there was no way to fix the problem. Cloudera Manager now allows you to resume the wizard and continue from where the error initially occurred.

Host Overrides

The HostsConfigurationpage for overriding configuration properties for selected hosts has been enhanced. See Viewing and Editing Host Overrides.

Fixed Issues in Cloudera Manager 7.0.3

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

OPSAPS-52886: When a license for Cloudera Manager expires, or the trial period expires, access to the Cloudera Manager Admin Console will be limited to only the license page until you install a new valid license

Cloudera Manager Admin Console features will no longer be disabled, but you will be unable to view or modify those features from the Cloudera Manager Admin Console.

OPSAPS-44883: Fixed installation and upgrade failures when installing the Cloudera Manager database on Maria DB 10.2.8 or higher.

This eliminates the following error message: "Key column 'REVISION_ID' doesn't exist in table"

OPSAPS-50104: Disabled support for Server Name Indication(SNI) in Cloudera Manager

SNI was causing the following browser errors:

- Chrome: This site can’t provide a secure connectionhost-10-17-100-224.coe.myco.com uses an unsupported protocol.ERR_SSL_VERSION_OR_CIPHER_MISMATCH
- Firefox:Secure Connection Failed. An error occurred during a connection to host-10-17-100-224.coe.myco.com:7183. Cannot communicate securely with peer: no common encryption algorithm(s). Error code: SSL_ERROR_NO_CYPHER_OVERLAP - curl (curl-7.29.0-46) * NSS error -12286 (SSL_ERROR_NO_CYPHER_OVERLAP) * Cannot communicate securely with peer: no common encryption algorithm(s). * Closing connection 0 curl: (35) Cannot communicate securely with peer: no common encryption algorithm(s).

OPSAPS-53041:Error generating TLS certificates

Fixed an issue that occurred when using very long host names (possibly caused by very long SDX environment names). This sometimes caused an error when generating TLS certificates.

OPSAPS-52953:Auto TLS keystore error

Fixed an issue that occurred when Auto-TLS is enabled. The following error: "no valid keystore" error sometimes occurs when starting the HBase Thrift Server.

OPSAPS-47386:

Leading and trailing white space is now trimmed from user names when creating new users either from the Cloudera Manager Admin Console the Cloudera Manager (API version v40 or higher).

OPSAPS-50447: Health Test for Hive Metastore Server Canary

The Health Test for Hive Metastore Server Canary fails to perform its task of checking HMS basic functionality (creating a database, table and partitions and then dropping them) and therefore reports bad health status in all cases.

Known Issues in Cloudera Manager 7.0.3

This topic describes known issues and workarounds for Cloudera Manager.

Installation and Upgrade Known Issues

Installation and Upgrade Limitations

Cloudera Manager 7.0.3 supports only installation of clusters running the Cloudera Runtime 7.0.3 components in this release.

Upgrades from previous versions of Cloudera Manager are not supported in this release.

Upgrades from clusters running CDH or HDP are not supported in this release.

Only Parcel installations of Cloudera Runtime are supported in this release.

Ranger Setup Issues

OPSAPS-52016: When running the Cloudera Manager Add Service wizard, passwords for the Ranger service must conform tto the following restrictions:

- Passwords must be at least 8 characters.
- Passwords must contain at least one alphabetic and one numeric character.
- The following characters cannot be used in passwords:

```
" ' \ ` ´ .
```

Other Known Issues

Connections to External Data Sources

The Microsoft ADLS connector is not supported in this release.

The Amazon S3 connector is not supported in this release.

OPSAPS-53304: During Hive replication, while importing column statistics, the system administrator must provide a valid engine type for the statistics to be usable.

Workaround: The system administrator must set the HIVE_REPL_STATS_ENGINE property in the Hive Replication Environment Advanced Configuration Snippet (Safety Valve), with the correct engine type for the column statistics to be usable. The valid values for the engine type are: hive, impala, and spark.

OPSAPS-53604 JDK Support

Only OpenJDK 8 is supported in CDP Data Center 7.0.

Oracle JDKs and OpenJDK 11 are not supported.

CDPD-5756: Hive replication fails at metastore import step with "java.net.BindException:Cannot assign requested address"

Workaround: You must run the Replication Manager service with a single thread. In the Advanced tab, the value for Number of concurrent HMS connections must be set to 0.

OPSAPS- 52546: Using Hue with HBase requires additional configurations

You must enable the following configuration parameters in the HBase service:

- Enable HBase Thrift Http Server
- Enable HBase Thrift Proxy Users

OPSAPS-53731:

When adding a new cluster, a new service, YARN Queue Manager is added by default. This service is useful for configuring the Capacity Scheduler. The Add Cluster and Add Service wizards will prompt users for the following credentials:

- Existing Cloudera Manager API Client Username
- Existing Cloudera Manager API Client Password

Enter the credentials for a user that is authorized to make configuration changes using the Cloudera Manager API.

OPSAPS-53214 HBase Hook for Atlas not enabled by default

When installing Atlas and HBase in a cluster, you must enable the HBase hook by doing the following in the Cloudera Manager Admin Console:

1. Go to the HBase service page.
2. Click the Configuration tab.
3. Search for the "Enable Atlas Hook" configuration property.
4. Select the Enable Atlas Hook for HBase.
5. Restart the HBase service.

OPSAPS-52454: Extra steps required to enable Ranger authorization in the Solr instance used by Ranger

After installing Ranger, do the following:

1. Login to Cloudera Manager.
2. Go to the Solr Service status page.
3. Click the Configuration tab.
4. Search for the Enable Ranger Authorization configuration property.
5. If the Enable Ranger Authorization property is not selected, select it.

Note: Don't select the Ranger Service dependency parameter. This is used for enabling a Solr service instance that is not used by the Ranger service.

6. Restart the Solr service.

CDPD-4139: Enabling TLS/SSL after creating a collection in Solr results in Solr not knowing that the node hosting the shard is the same.

A cluster with indices stored on HDFS created before enabling TLS1, will get the following error message after enabling TLS and starting the cluster:

"Will not load SolrCore SOLR_CORE_NAME because it has been replaced due to failover."

Workaround: Recreate the collection after enabling TLS. New collections created after enabling TLS are not affected.

OPSAPS-51224: Atlas custom properties ignored in client services

When adding a custom property for the atlas-application.properties in Atlas hook-based services such as Hive, HBase, and Impala, the custom property is not reflected in the actual configuration file that Cloudera Manager generates, causing these properties to be ignored.

Workaround: none

CDPD-6022 Accumulo not supported

Apache Accumulo is currently not supported in this version of Cloudera Runtime. Although you can access Accumulo from the command-line interface, you must not use this component in production because Cloudera does not support it.

Apache Flume is no longer supported.

Apache Pig is no longer supported

Virtual Private Clusters (VPC) are not recommended for use in production environments.

You can still create Virtual Private Clusters (Base clusters and Data contexts using Cloudera Manager, but you should only use these in development or testing environments.

OPSAPS-54299 – Installing Hive on Tez and HMS in the incorrect order causes HiveServer failure

You need to install Hive on Tez and HMS in the correct order; otherwise, HiveServer fails. You need to install additional HiveServer roles to Hive on Tez, not the Hive service; otherwise, HiveServer fails. See Installing Hive on Tez for the correct procedures.

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 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.

Knowledge article

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

TSB 2021-530: Local File Inclusion (LFI) Vulnerability in Navigator

After successful user authentication to the Navigator Metadata Server and enabling dev mode of Navigator Metadata Server, local file inclusion can be performed through the Navigator's embedded Solr web UI. All files can be accessed for reading which can be opened as cloudera-scm OS user. This is related to Apache Solr CVE-2020-13941.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: TSB 2021-530: CVE-2021-30131 - Local File Inclusion (LFI) Vulnerability in Navigator