

Cloudera Runtime 7.1.9

Configuring and Using Ranger RMS (Hive-HDFS ACL Sync)

Date published: 2020-07-28

Date modified: 2023-09-07

CLouDERA

<https://docs.cloudera.com/>

Legal Notice

© Cloudera Inc. 2024. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

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.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

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

Ranger RMS - HIVE-HDFS ACL Sync Overview.....	4
About Ranger RMS for Ozone.....	8
Ranger RMS Assumptions and Limitations.....	8
Installing/Verifying RMS for Ozone configuration.....	9
Enabling RMS for Ozone authorization.....	10
Understanding Ranger policies with RMS.....	11
How to full sync the Ranger RMS database.....	13
Configuring High Availability for Ranger RMS (Hive-HDFS ACL-Sync).....	14
Configuring Ranger RMS (Hive-HDFS / Hive-OZONE ACL Sync).....	19
Configuring HDFS plugin to view permissions through getfacl interface.....	22
Ranger RMS (Hive-HDFS ACL-Sync) Use Cases.....	23

Ranger RMS - HIVE-HDFS ACL Sync Overview

Ranger Resource Mapping Server (RMS) enables automatic translation of access policies from HIVE to HDFS.

About HIVE-HDFS ACL Sync

It is common to have different workloads use the same data – some require authorizations at the table level (Apache Hive queries) and others at the underlying files (Apache Spark jobs). Unfortunately, in such instances you would have to create and maintain separate Ranger policies for both Hive and HDFS, that correspond to each other.

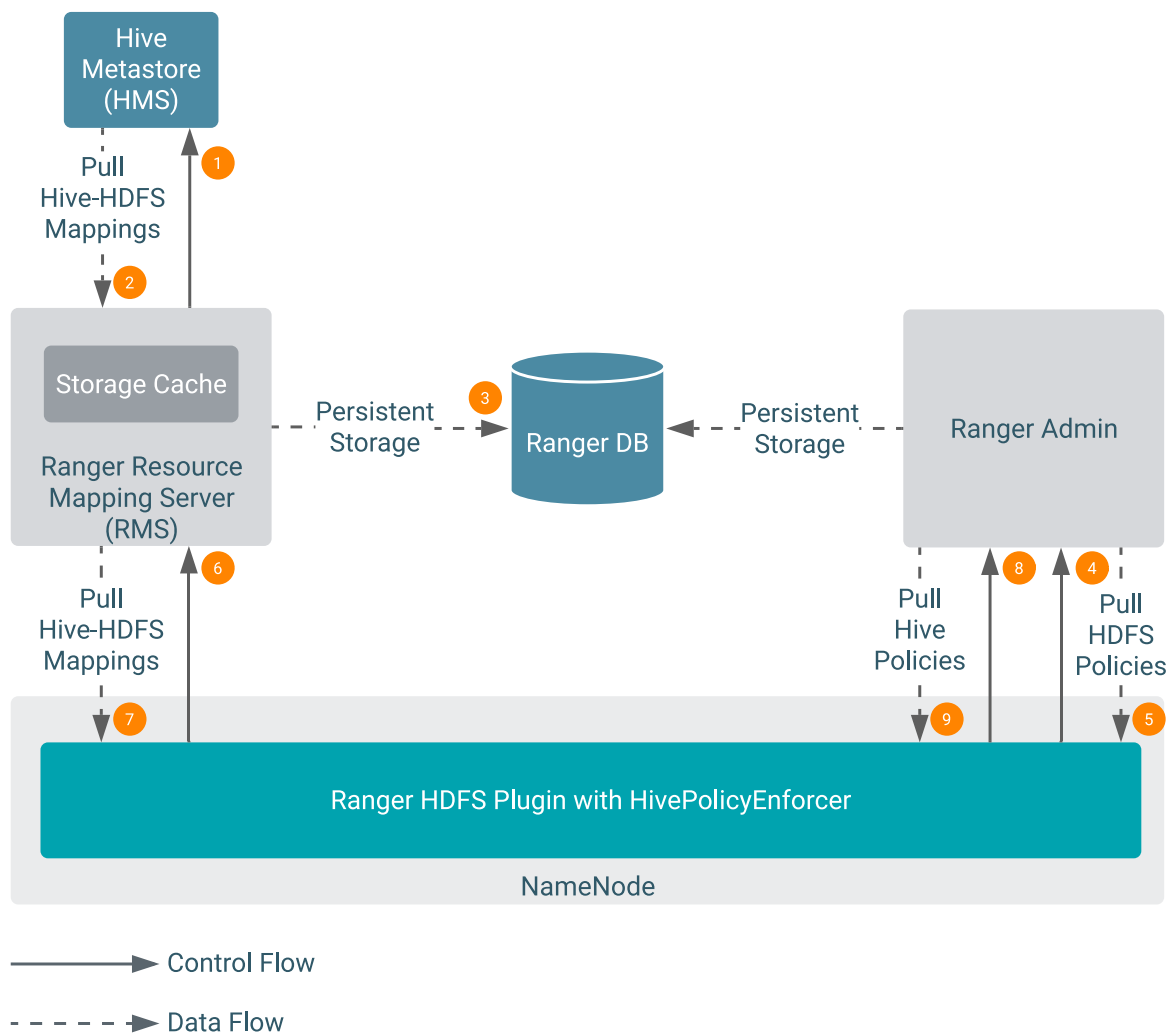
As a result, whenever a change is made on a Hive table policy, the data admin should make a consistent change in the corresponding HDFS policy. Failure to do so could result in security and/or data exposure issues. Ideally the data admin would set a single table policy, and the corresponding file access policies would automatically be kept in sync along with access audits, referring to the table policy that enforced it.

Legacy CDH users had a feature called the Hive-HDFS ACL sync which had Hive policies in Apache Sentry that automatically linked Hive permissions with HDFS ACLs. This was especially convenient for external table data used by Spark or Hive.

Prior to CDP 7.1.6, Ranger only supported manually managing Hive and HDFS policies separately. Ranger RMS (Resource Mapping Server) allows you to authorize access to HDFS directories and files using policies defined for Hive tables. RMS is the service that enables Hive-HDFS Policy Sync.

RMS periodically connects to the Hive Metastore and pulls Hive metadata (database-name, table-name) to HDFS file-name mapping. The Ranger HDFS Plugin (running in the NameNode) has been extended with an additional HivePolicyEnforcer module. The HDFS plugin downloads Hive policies from Ranger Admin, along with the mappings from Ranger RMS. HDFS access is determined by both HDFS policies and Hive policies.

Figure 1: HIVE-HDFS ACL Sync using Ranger RMS



Phase I (items 1-3 above)

Ranger RMS periodically connects to the HIVE Metastore and pulls HIVE metadata (database-name, table-name) to HDFS file-name mapping.

Phase II (items 4-9 above)

The Ranger HDFS Plugin (running in the NameNode) periodically pulls HDFS policies from Ranger Admin. With the introduction of Ranger RMS, the Ranger HDFS Plugin (running in the NameNode) that has been extended with an additional HIVEPolicyEnforcer module. It now pulls down the HIVE-HDFS mappings from RMS and HIVE Policies from Ranger Admin.

After phase II completes, the requested HDFS access is determined in the NameNode by the HDFS and HIVE policies defined by the Ranger Administrator.

Currently, Ranger performs authorization of the HDFS commands which require access to the hierarchy of files/directory rooted at the argument passed to the HDFS command as described below. Some examples of such commands are :

```
hdfs dfs -count -q -h -v <directory>; hdfs dfs -R <directory>
```

HDFS Authorization Interface

When these commands are invoked, HDFS Namenode builds a tree of i-nodes corresponding to <directory>, and passes it to the authorizer with a flag indicating that subAccess (access to the directory hierarchy rooted at <directory>) is to be checked.

About database-level grants feature

Legacy CDH users used HIVE policies in Apache Sentry that automatically linked HIVE permissions with HDFS ACLs. This was especially convenient for external table data used by Spark or HIVE. Specifically, using Sentry, you could make grants at the HIVE database level and HDFS permissions would propagate to the database directory, and to all tables and partitions under it.

Previously, Ranger only supported managing HIVE and HDFS policies separately. Ranger Resource Mapping Server (RMS) now allows you to create a database level policy in HIVE and have these permissions propagate to the HDFS locations and all tables under it. RMS is the service that enables HIVE-HDFS ACL Sync.

RMS captures database metadata from the HIVE Meta Store (HMS). After the first, full-synchronization run, RMS downloads mappings for tables and databases present in the HMS.

Whenever you create a new database, RMS synchronizes metadata information from HMS and uses it to update the resource mapping file linking HIVE database resources to their corresponding HDFS location. Any user with access permissions on a HIVE database automatically receives similar HDFS file-level access permissions on the database's data files. Select/ Read access for any user in the database location is allowed through default HIVE policy for all-databases. This behavior is treated as `_any` access, which is similar to the HIVE command `show tables`. If a user has no HIVE policy which allows access on the database, then the user is denied access to the corresponding HDFS location of that database. Previously, users were not allowed to access the HDFS location of a database even if the user had permission to access the database through a HIVE policy. The HDFS to HIVE access type mappings follow:

Access Type mapping for HDFS to HIVE for Database:

- `_any=[_any]`
- `read=[_any]`
- `write=[create, drop, alter]`
- `execute=[_any]`

Access Type mapping for HDFS to HIVE for Table:

- `_any=[_any]`
- `read=[select]`
- `write=[update, alter]`
- `execute=[_any]`

If you create tables under a database but the HDFS location of the corresponding table does not reside under the HDFS location of that database (for example: table locations are external locations), the HIVE policies (database-name, table = *, column= *) translate into HDFS access rules and allow the HDFS NameNode to enforce them. If the policy is created only for the database resource, the same access translates to the HDFS location of that database only; not for the tables residing under that database.

Ranger RMS Assumptions and Limitations

- All partitions of a table are assumed to be under the location specified for the table. Therefore, table permissions will not authorize access to partitions that store data outside the location specified for the table. For example, if a table is located in a `/warehouse/foo` HDFS directory, all partitions of the table must have locations that are under the `/warehouse/foo` directory.
- The Ranger RMS service is not set up automatically when a CDP Private Cloud Base cluster is deployed. You must install and configure Ranger RMS separately.
- Ranger policies should be configured (with `ranger_rms` user access) before RMS is started and runs the first sync from the HIVE Metastore (HMS).
- The Ranger RMS ACL-sync feature supports a single logical HMS, to evaluate HDFS access via HIVE permissions. This is aligned with the Sentry implementation in CDH.

- Permissions granted on views (traditional and materialized) do not extend to HDFS access. This is aligned with the Sentry implementation in CDH.
- If a Private Cloud Base deployment supports multiple logical HMS with a single Ranger, Ranger RMS (Hive-HDFS ACL-Sync) works with only one logical HMS. Permissions granted on databases/tables in other logical HMS instances will not be considered to authorize HDFS access.
- Namenode memory requirements must be increased, based on the number of HIVE table mappings downloaded to HDFS Ranger plugin. Additionally, maintaining HIVE policies in memory cache will also require additional memory.
- Expect Namenode CPU load to increase, due to additional access evaluation performed to enforce HIVE policies and periodic downloading and processing of the HIVE table mappings. The latter increase is proportional to the number of table mappings downloaded to HDFS Ranger plugin.
- When multiple databases are mapped to single HDFS location, and if a HIVE policy allows a user to access one database. Then, users will be able to access its HDFS location and all other files & directories under it. This may include table or database directories of other databases and tables. However, users will not be able to access other databases or tables under it through Hive queries.

For example,

music_a, music_b, music_c are created at HDFS path '/data'.

Policy-A to allow 'sam' user 'all' access on resource = {database=music_a; table= * ; column= * ; }

Now, 'sam' user will get all access on hdfs path /data and files, directories under it. Therefore, 'sam' user will be able to access hdfs location of tables under music_b and music_c databases as long as those locations reside under /data directory.

However, 'sam' user will not be able to access music_b and music_c databases or any tables under these databases through Hive queries.



Note: This is the expected behaviour when the hdfs location for different database/tables are designed in such a way that these entities share the same hdfs path or HIVE database/table locations fall under another database hdfs location.

Comparison with Sentry HDFS ACL sync

The Ranger RMS (Hive-HDFS ACL-Sync) feature resembles the Sentry HDFS ACL Sync feature in the way it downloads and keeps track of the HIVE table to HDFS location mapping.

It differs from Sentry in the way it completely and transparently supports all features that Ranger policies express. Therefore, support for tag-based policies, security-zones, masking and row-filtering and audit logging is included with this implementation.

Also, the feature is enabled or disabled by a simple configuration on the HDFS side, allowing each installation the option of turning this feature on or off.

Considerations when upgrading Ranger RMS

Prior to CDP-7.1.7, Ranger RMS synchronizes only table metadata. Database metadata mappings were not downloaded from HMS. After migrating from CDP versions <7.1.7 to CDP 7.1.8+, only pre-existing table mappings will be available. Any newly created tables or database mappings will be synchronized in RMS. Any pre-existing database mappings will not be present. To ensure the pre-existing databases are mapped correctly, you must perform a full-sync after completing the upgrade.

Related Information

[Installing Ranger RMS](#)

About Ranger RMS for Ozone

In CDP 7.1.9, Ranger RMS will support authorization for Ozone storage locations. RMS for Ozone will co-exist with Hive-HDFS ACL sync and provide authorization for both HDFS and Ozone file systems.

RMS periodically connects to the Hive Metastore (HMS), pulls the following Hive metadata

- database-name
- table-name
- location

and stores it into the Ranger database. The Ranger Ozone plugin (running in the OzoneManager) has been extended with an additional RangerOzoneHiveChainedPlugin module. After enabling RMS for Ozone authorization, it downloads Hive policies, Tags and Roles from Ranger Admin, along with the resource-mappings from Ranger RMS. Both Ozone policies and Hive policies determine Ozone access.

Previously, Ranger only supported managing Hive and Ozone policies separately. Ranger Resource Mapping Server (RMS) now allows you to create a database/table level policy in HIVE and have these permissions propagate to the Ozone storage locations of these databases/tables and all files and directories under it. Users can use Ozone as a storage technology to create databases/tables along with the existing HDFS filesystem and RMS is the service that enables HIVE-OZONE ACL Sync along with the HIVE-HDFS ACL Sync. (Prior to CDP 7.1.9 only HIVE-HDFS ACL sync was supported).

After the first, full-synchronization run, RMS downloads mappings for tables and databases present in the HMS. These tables/databases could be located in both OZONE and HDFS file systems. RMS will map the tables and databases with their respective storage locations and store them into the Ranger database with its associated service-id (cm_ozone or cm_hdfs). When ranger-ozone-plugin or ranger-hdfs-plugin requests mappings, it only fetches the requested service mappings. In other words, ranger-ozone-plugin will download mappings only for tables/databases whose storage location is Ozone. Similarly, ranger-hdfs-plugin will download mappings only for tables/databases whose storage location is HDFS. Read access for any user in the database location is allowed through the default HIVE policy for all-databases. This behavior is treated as `_any` access, which is similar to the HIVE command `show tables`. If a user has no HIVE policy which allows access to the database, then the access is denied to the corresponding Ozone location of that database. This access evaluation aligns with the HDFS db-level grants feature.

Ranger RMS Assumptions and Limitations

- RMS for Ozone does not support semantics enforced by Sentry for HDFS ACLs. This implies that even when `ranger.plugin.ozone.mapping.hive.authorize.with.only.chained.policies` is set to true, to access ozone locations when RMS is enabled; READ access must be given for the user on bucket and volume where Hive Tables and Databases are located.
- In addition to Ozone policies, Ranger RMS will authorize access only for "key" resources which are the database/table locations in ozone. To access volumes & buckets, user should have appropriate access through policies of Ozone service configured Ranger Admin.
- The Ranger RMS ACL-sync feature supports a single logical HMS, to evaluate OZONE/HDFS access via HIVE permissions.
- RMS ACL sync is designed to work on a specific pair of HDFS<->Hive and OZONE<->Hive Ranger services. It will support only one pair of "HDFS and HIVE" services and one pair of "OZONE and HIVE" services.
- By default, all external tables are stored into HDFS. To create an external table in Ozone, please specify the location clause by providing the Ozone location at the time of creating a table or configure Hive External Warehouse Directory in HMS.
- The default Hive Warehouse directory is configured in HiveMetaStore (HMS) to store managed tables into HDFS. To store managed tables into Ozone, please configure the Hive Warehouse directory or update the database managed location as an Ozone storage location. Please check with the Ozone/Hive support members to know more about creating Hive Tables/Databases into Ozone.

- Metrics support for RMS with Ozone is not supported in the CDP-7.1.9 release.

Installing/Verifying RMS for Ozone configuration

Ranger RMS for Ozone is available for new CDH-7.1.9 installations.

About this task

Manual steps to verify RMS for Ozone setup:

Procedure

1. Verify if Ranger RMS Url is configured in Ozone ranger-ozone-security.xml file.
 - a) Go to Cloudera Manager Ozone Configuration Ozone Manager Advanced Configuration Snippet (Safety Valve) for ozone-conf/ranger-ozone-security.xml .
 - b) Add the Ranger RMS url to the following property:

ranger.plugin.ozone.mapping.source.url =

<rms-url>

After configuring this property with appropriate Ranger RMS Url, ranger-ozone-plugin will be able to communicate with RMS (Resource Mapping Server) for downloading files.

2. Verify if the Ozone service name is correctly configured.
 - a) Go to Cloudera Manager Ranger RMS Configuration , and search for the property ranger_rms_hms_source_service_name_ozone .
 - b) Verify the following property and value:

ranger-rms.HMS.source.service.name.ozone =

cm_ozone



Note: The value cm_ozone is the name of the OZONE service configured in the Ranger admin. If this configuration is changed after Ranger RMS is started and has synchronized with Hive Metastore, the only way to have Ranger RMS use a new configuration is by performing a RMS Full Sync. Please refer to the related topic, "[How to full sync the Ranger RMS database](#)" for specific steps.

3. Verify whether ozone schema's are configured.
 - a) Go to Cloudera Manager Ranger RMS Configuration Ranger RMS Supported Uri Scheme .
 - b) Verify the following property and value:

ranger-rms.supported.uri.scheme =

hdfs,o3fs,ofs

4. Log in to Ranger Admin Web UI.
5. On Service Manager, in Hadoop SQLservice, click Edit.
6. Verify that "om" user has been added to the tag.download.auth.users and policy.download.auth.users configurations.

Otherwise, add the following properties and values to the cm_hive service configuration:

- a) tag.download.auth.users = hive,hdfs,impala,om
- b) policy.download.auth.users = hive,hdfs,impala,om



Note: This ozone service user can also be found in the ranger-admin-site.xml file with the following property key: ranger.plugins.ozone.serviceuser=om or in the service-usernames.csv file generated by Cloudera Manager.

Related Tasks

[How to full sync the Ranger RMS database](#)

Enabling RMS for Ozone authorization

You must grant necessary permissions and update policies to access Ozone.

About this task

When Ranger is enabled in the cluster, any user other than the default admin user, om requires the necessary Ranger permissions and policy updates to access the Ozone file system. To create a Hive external table that points to the Ozone file system, the hive user should have the required permissions in Ranger. Give all accesses to hive user by updating default policy all - volume, bucket, key in cm_ozone service.

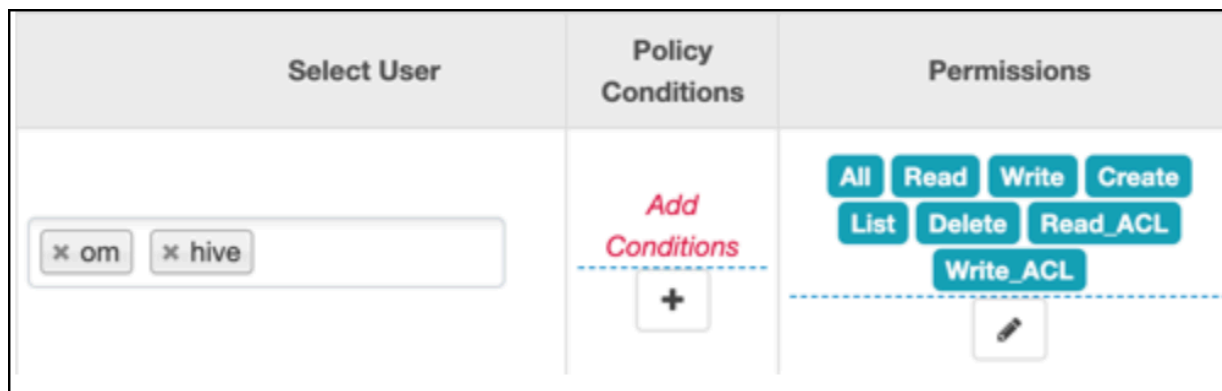
Procedure

1. Login to Ranger Admin Web UI.
 - a) Go to Cloudera Manager Clusters Ranger Ranger Admin Web UI .
 - b) Type your username and password.
 - c) Click Sign In.

The Service Manager for Resource-based Policies page displays.

2. In Service Manager cm_ozone service, click Edit.
3. In cm_ozone policies all - volume, bucket, key policy, click Edit.
4. In Allow Conditions, add the hive user, choose necessary permissions, then click Save, as shown in:

Figure 2: Editing Ozone plugin permissions for hive user



5. Grant Read access to users only on volumes and buckets where Hive tables and databases are located. (Recommended)
 - a) create a new policy in Ozone service (cm_ozone) for volumes and buckets where Hive tables/databases are located.
 - b) In Allow Conditions, add the users, groups, and roles and give Read permission.
 or else, to grant Read access to everyone on all volumes and buckets, in Ozone service (cm_ozone) policies page,
 - a) edit the all - volume, bucket policy
 - b) add the public group to the group list
 - c) give Read permission



Note: For all databases and tables whose storage location is in Ozone; these locations will be mapped to ozone resource type Key and after enabling RMS for OZONE authorization; ranger-ozone-plugin will evaluate access on a Key based on Hive policies. If no applicable Hive policy exists, then the Ozone policies will determine access. Otherwise, access is denied.

6. Add chained properties to the ranger-ozone-security.xml file.
 - a) Go to Cloudera Manager Ozone Configuration Ozone Manager Advanced Configuration Snippet (Safety Valve) for ozone-conf/ranger-ozone-security.xml .
 - b) Click +Add to add the following properties and values:

ranger.plugin.ozone.chained.services =

cm_hive

ranger.plugin.ozone.chained.services.cm_hive.impl =

org.apache.ranger.chainedplugin.ozone.hive.RangerOzoneHiveChainedPlugin



Note: The value cm_hive is the name of the Hive/Hadoop SQL service configured in the Ranger admin.

7. Restart the Ozone service.

Results

If everything is correctly configured as explained above, the ranger-ozone-plugin automatically communicates with RMS, downloads Hive-Ozone mappings and stores them into policy-cache directory as a ozone_cm_hive_resource_mapping.json file. It will also communicate with the Ranger admin and download Hive policies, Hive tags and Hive roles.

The following files download into the policy-cache directory

(default configured value: /var/lib/ranger/ozone/policy-cache/)

after enabling RMS for Ozone authorization:

- ozone_cm_hive_resource_mapping.json
- ozone_cm_hive.json
- ozone_cm_hive_tag.json
- ozone_cm_hive_roles.json

Understanding Ranger policies with RMS

Ranger RMS for HDFS access evaluation workflow

At a high level, the Ranger RMS for HDFS access evaluation workflow is as follows:

- Ranger policies for the HDFS service are evaluated. If any policy explicitly denies access, access is denied.
- Ranger checks to see if the accessed location maps to a Hive table.
- If it does, Hive policies are evaluated for the mapped Hive table. If there is an HDFS policy allowing access, access is allowed. Otherwise, the default HDFS ACLs determine the access.

Requested HDFS permission is mapped to Hive permissions as follows:

HDFS 'read' ==> Hive 'select'

HDFS 'write' ==> Hive 'update' or 'alter'

HDFS 'execute' ==> Any Hive permission

- If there is no Hive policy that explicitly allows access to the mapped table, access is denied, otherwise access is allowed.

Appropriate tag policies are considered both during HDFS access evaluation and if needed, during Hive access evaluation phases. Also, one or more log records are generated to indicate which policy, if any, made the access decision.

The following scenarios illustrate how the access permissions are determined. All scenarios assume that the HDFS location is NOT explicitly denied access by a Ranger HDFS policy.

Location does not correspond to a Hive table.

In this case, access will be granted only if a Ranger HDFS policy allows access or HDFS native ACLs allow access. The audit log will show which policy (or Hadoop-acl) made the decision.

Location corresponds to a Hive table.

A Ranger Hive policy explicitly denied access to the mapped table for any of the accesses derived from the original HDFS request.

- Access will be denied by Hive policy.

There is no matching Ranger Hive policy.

- Access will be denied. Audit log will not specify the policy.

Ranger policy masks some columns in the mapped table.

- Access will be denied. Audit log will show Hive masking policy.

Mapped Hive table has a row-filter policy

- Access will be denied. Audit log will show Hive Row-filter policy.

A Ranger Hive policy allows access to the mapped table for the access derived from the original HDFS access request.

- Access will be granted. If the access was originally granted by HDFS policy, the audit log will show HDFS policy.

Ranger RMS for Ozone access evaluation workflow

At a high level, the Ranger RMS for Ozone access evaluation workflow is as follows:

- For every “key” access request made to ranger-ozone-plugin, there are three separate requests are received for authorization as follows:
 - To check READ access on volume
 - To check READ access on bucket
 - To check access made originally for a key
- First two accesses on volume and buckets are determined by Ozone policies in Ranger Admin. RMS will only evaluate access for the key resource.
- Further for the key access evaluation, Ranger policies for the Ozone service are evaluated. If any policy explicitly denies access, access is denied.
- Ranger checks to see if the accessed location (key) maps to a Hive table.
- If it does, Hive policies are evaluated for the mapped Hive table. If there is an OZONE policy allowing access, access is allowed. Otherwise, the access will be denied.
- Requested OZONE permission is mapped to Hive permissions as follows:

Ozone ACL	Should translate into Hive ACL For Table	Should translate into Hive ACL For Database
read	select	_any
write	update	update
create	create	create
list	select	select
delete	drop	drop
read_acl	select	select
write_acl	update	update

- If there is no Hive policy that explicitly allows access to the mapped table, access is denied, otherwise access is allowed.
- Appropriate tag policies are considered both during OZONE access evaluation and if needed, during Hive access evaluation phases.
- One or more log records are generated to indicate which policy, if any, made the access decision.

How to full sync the Ranger RMS database

You must perform a full-sync of the Ranger RMS database after upgrading RMS to v7.1.8 or above.

About this task

Performing a full-sync in Ranger RMS truncates all the existing data from RMS tables in the Ranger database. The sync process initiates a fresh synchronization from the Hive Metastore (HMS). All available tables and database metadata download from HMS and persist into the Ranger database.

Prior to CDP-7.1.7, Ranger RMS synchronizes only table metadata. Database metadata mappings were not downloaded from HMS. After migrating from CDP versions before CDP-7.1.7 to CDP 7.1.8+, only pre-existing table mappings will be available. Any newly created tables or database mappings will be synchronized in RMS. Any pre-existing database mappings will not be present.



Important: To ensure the pre-existing databases are mapped correctly, you must perform a full-sync after upgrading RMS to 7.1.8 or above versions.

Before you begin

Have a (7.1.8) cluster with Ranger and Ranger RMS service up and running.

Procedure

1. In Cloudera Manager Ranger_RMS Actions , click Stop.
2. In Ranger_RMS Actions , click Ranger RMS Database Full Sync.
3. Confirm that the command completed successfully, as shown in the following example:

Figure 3: Confirming successful Ranger RMS database full sync

The screenshot displays a window titled "Ranger RMS Database Full Sync". At the top, it shows the status as "Finished" with a green checkmark, the context as "RANGER_RMS-1", and the execution time as "Aug 16, 8:51:53 AM" and "9.31s". Below this, it states "Command Ranger RMS Database Full Sync finished successfully on service RANGER_RMS-1." and "Completed 1 of 1 step(s)". There are three radio buttons for filtering steps: "Show All Steps" (selected), "Show Only Failed Steps", and "Show Only Running Steps". The main area shows two steps:

Step	Service	Time	Duration
Execute command Ranger RMS Database Full Sync on role Ranger RMS Server (os-mv-test-3) Command (Ranger RMS Database Full Sync (1546335788)) has completed successfully	Ranger RMS Server (os-mv-test-3)	Aug 16, 8:51:53 AM	9.31s
Ranger RMS Database Full Sync Ranger RMS Database Full Sync finished successfully on Ranger	Ranger RMS Server (os-mv-test-3)	Aug 16, 8:51:53 AM	9.26s

A "Close" button is located at the bottom right of the window.



Note: You can also confirm the text - '[I] Truncate operation executed successfully on Ranger RMS tables' from the stderr log of the command from Cloudera Manager UI.

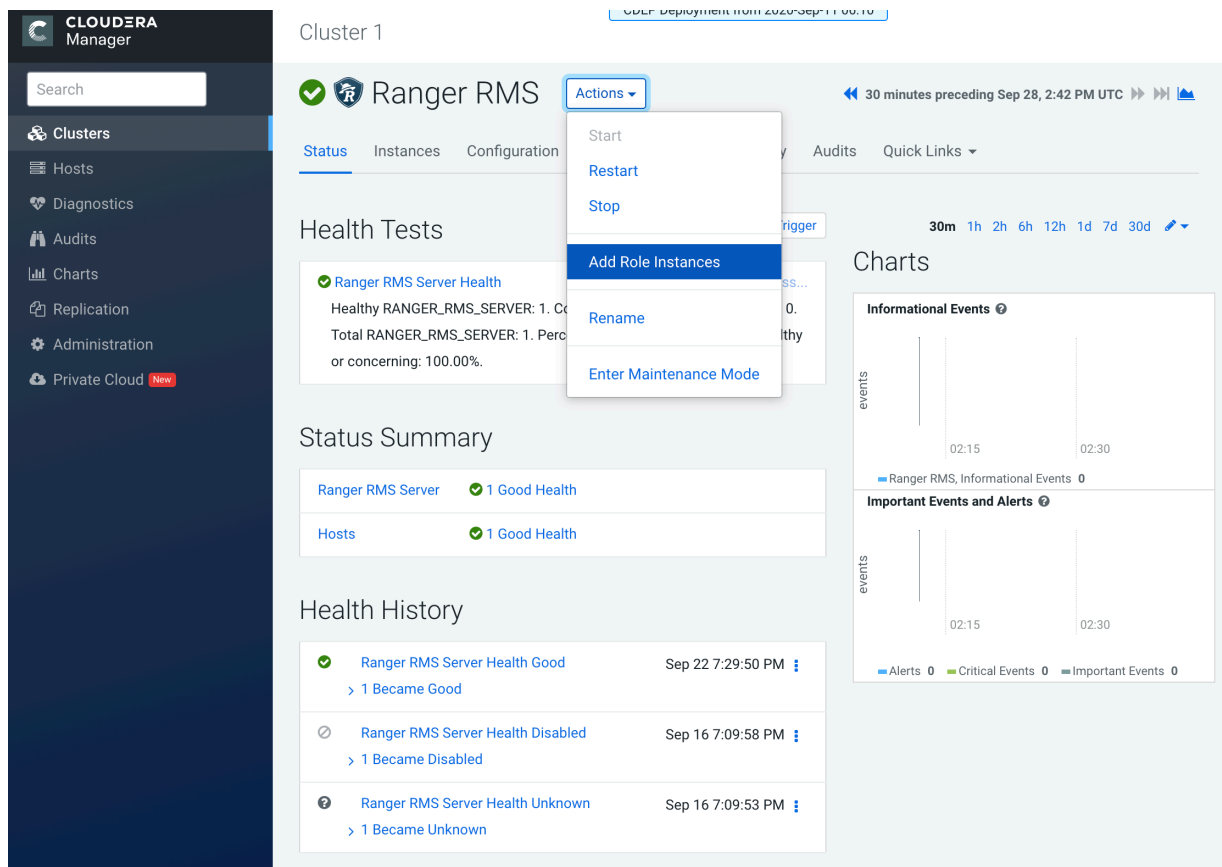
4. In Cloudera Manager Ranger_RMS , click Start.

Configuring High Availability for Ranger RMS (Hive-HDFS ACL-Sync)

Use the following steps to configure high availability for the Ranger Resource Mapping Server (RMS) and Hive-HDFS ACL Sync.

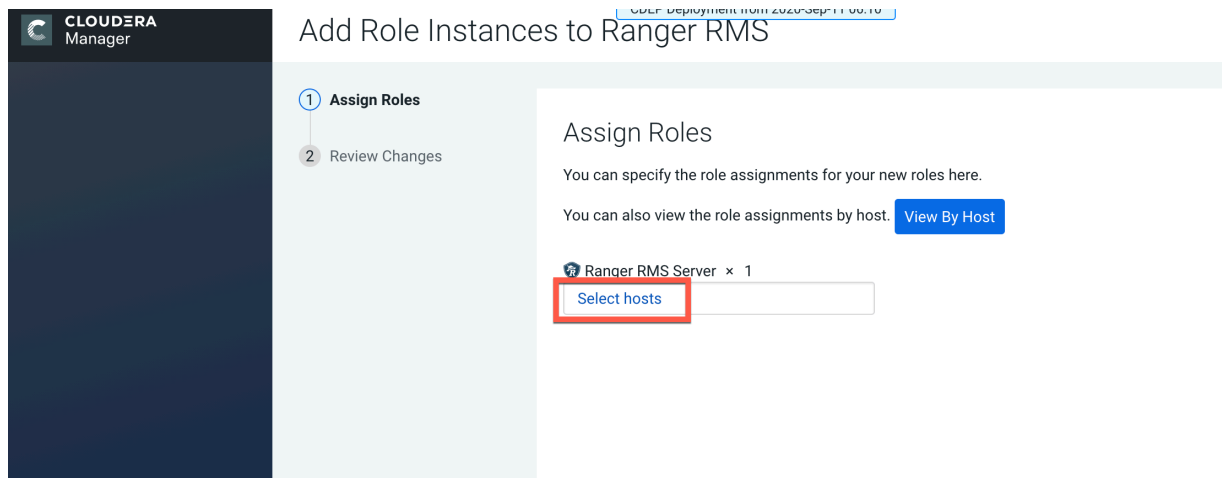
Procedure

1. In Cloudera Manager, select Ranger RMS, then select Actions > Add Role Instances.



The screenshot shows the Cloudera Manager interface for the Ranger RMS service. The left sidebar contains navigation options like Clusters, Hosts, Diagnostics, Audits, Charts, Replication, Administration, and Private Cloud. The main content area displays the Ranger RMS status, including Health Tests, Status Summary, and Health History. The 'Actions' dropdown menu is open, showing options: Start, Restart, Stop, Add Role Instances (highlighted), Rename, and Enter Maintenance Mode. The Status Summary shows 1 Good Health for both Ranger RMS Server and Hosts. The Health History shows a recent 'Good' status change on Sep 22 7:29:50 PM.

2. On the Add Role Instances page, click Select hosts.



The screenshot shows the 'Add Role Instances to Ranger RMS' page in Cloudera Manager. The page is titled 'Assign Roles' and includes instructions: 'You can specify the role assignments for your new roles here.' and 'You can also view the role assignments by host. View By Host'. Below this, there is a section for 'Ranger RMS Server x 1' with a 'Select hosts' button highlighted by a red box.

- On the selected hosts page, select a backup Ranger RMS host. A Ranger RM (RR) icon appears in the Added Roles column for the selected host. Click OK to continue.

2 Hosts Selected

Select hosts for a new or existing role. The host list is filtered to remove hosts that are not valid candidates; these include hosts that are unhealthy, members of other clusters, or have an incompatible version of the software installed on them.

Enter hostnames: host01, host[01-10], IP addresses or rack. [Search](#)

Tip: Click the first checkbox, hold down the Shift key and click the last checkbox to select a range.

<input type="checkbox"/>	Hostname ↑	IP Address	Rack	Cores	Physical Memory	Existing Roles	Added Roles
<input checked="" type="checkbox"/>	dhoyle714-1.dhoyle714.root.hwx.site	172.27.100.70	/default	80	251.6 GiB	AS, CCS, G, HB..., RS, DN, G, G, G, ID, KB, KC, KG, M, G, LS, RA, RT, RU, SRS, G, G, SM..., SM..., SR..., SR..., G, G, NM, ZS	RR...
<input type="checkbox"/>	dhoyle714-2.dhoyle714.root.hwx.site	172.27.100.73	/default	32	251.6 GiB	M, B, NN, NF..., SNN, G, HMS, G, HS2, LB, HS, KTR, ICS, ISS, KB, KC, LHBI, TS, G, AP, CS, LM, DM, SM	

[Cancel](#) [OK](#)

- The Add Role Instances page is redisplayed with the new backup host. Click Continue.

CLDP deployment from zuzu-sep-11 06:10

Add Role Instances to Ranger RMS

- Assign Roles
- Review Changes

Assign Roles

You can specify the role assignments for your new roles here.

You can also view the role assignments by host. [View By Host](#)

Ranger RMS Server x (1 + 1 New)

dhoyle714-1.dhoyle714.root.hwx.site

[Back](#) [Continue](#)

5. Review the settings on the Review Changes page, then click Continue.

The screenshot shows the Cloudera Manager interface for configuring Ranger RMS. The left sidebar contains navigation options: Parcels, Running Commands, Support, and the user 'admin'. The main content area is titled 'Add Role Instances to Ranger RMS' and shows a 'Review Changes' step. The settings are organized into two columns. The first column lists configuration keys, and the second column shows the default group and the current value. Each setting has a help icon (question mark) to its right.

Configuration Key	Default Group	Value
Ranger RMS Max Heapsize ranger_rms_max_heap_size	Ranger RMS Server Default Group	1 GIB
RMS HTTP Port ranger-rms.service.http.port	Ranger RMS Server Default Group	8383
RMS HTTPS Port ranger-rms.service.https.port	Ranger RMS Server Default Group	8484
Ranger RMS Server TLS/SSL Client Trust Store File ranger-rms.truststore.file	Ranger RMS Server Default Group	
Ranger RMS Server TLS/SSL Client Trust Store Password ranger-rms.truststore.password	Ranger RMS Server Default Group	
Enable TLS/SSL for Ranger RMS Server ranger-rms.service.https.attrib.ssl.enabled	<input type="checkbox"/> Ranger RMS Server Default Group	
Ranger RMS Server TLS/SSL Server JKS Keystore File Location ranger-rms.service.https.attrib.keystore.file	Ranger RMS Server Default Group	
Ranger RMS Server TLS/SSL Server JKS Keystore File Password ranger-rms.service.https.attrib.keystore.password	Ranger RMS Server Default Group	

At the bottom right of the page, there are two buttons: 'Back' and 'Continue'.

6. The new role instance appears on the Ranger RMS page.

The screenshot shows the Cloudera Manager interface for a cluster named 'Cluster 1'. The main heading is 'Ranger RMS' with an 'Actions' dropdown. Below this are tabs for 'Status', 'Instances', 'Configuration', 'Commands', 'Charts Library', 'Audits', and 'Quick Links'. A search bar and a 'Filters' toggle are present. A 'Filters' sidebar on the left lists categories like STATUS, COMMISSION STATE, MAINTENANCE MODE, RACK ID, ROLE GROUP, ROLE TYPE, STATE, and HEALTH TEST. The 'STATUS' filter shows 'Good Health' (1) and 'Stopped' (1). A table of instances is displayed with columns: Status, Role Type, State, Hostname, Commission State, and Role Group. The table contains two rows. The second row is highlighted with a red border. Below the table is a '1 - 2 of 2' indicator. Buttons for 'Add Role Instances' and 'Role Groups' are visible above the table.

Status	Role Type	State	Hostname	Commission State	Role Group
<input type="checkbox"/>	Ranger RMS Server	Started	[Redacted]	Commissioned	Ranger RMS Server Default Group
<input type="checkbox"/>	Ranger RMS Server	Stopped	[Redacted]	Commissioned	Ranger RMS Server Default Group

7. In Cloudera Manager, select Ranger RMS, then click Configuration.

- a) Select the Ranger RMS Server HA checkbox. In the Ranger RMS Server IDs box, add a comma-separated list of IDs for each RMS server.

The screenshot shows the configuration page for Ranger RMS Server HA. The 'Ranger RMS Server HA' checkbox is checked. The 'Ranger RMS Server IDs' field contains 'id1,id2'. The 'Ranger RMS Server Default Group' is set to '/ranger-rms'.

- b) Use the Add (+) icons for the Ranger RMS Server Advanced Configuration Snippet (Safety Valve) for ranger-rms-conf/ranger-rms-site.xml property to add entries for each RMS host with its corresponding server ID.

- ranger-rms.server.address.id1=<hostname of RMS server for id1>:8383
- ranger-rms.server.address.id2=<hostname of RMS server for id2>:8383



Note: If SSL is enabled, use port 8484.

The screenshot shows the configuration page for Ranger RMS Server Advanced Configuration Snippet (Safety Valve) for ranger-rms-conf/ranger-rms-site.xml. Two entries are added for id1 and id2. The 'Name' field is 'ranger-rms.server.address.id1' and the 'Value' field is 'hostname:8383'. The 'Description' field is empty. The 'Final' checkbox is unchecked.

8. Click Save Changes, then Click the Restart icon.

9. On the Stale Configurations page, click Restart Stale Services.

10. On the Restart Stale Services page, select the Re-deploy client configuration checkbox, then click Restart Now.

11. A progress indicator page appears while the services are being restarted. When the services have restarted, click Finish.

Related Information

[Installing Ranger RMS](#)

Configuring Ranger RMS (Hive-HDFS / Hive-OZONE ACL Sync)

Ranger Resource Mapping Server (RMS) should be fully configured after installation. This topic provides further information about RMS configuration settings and workflows.

Important configuration information - Hive Ozone

- Ranger RMS enables OZONE/HDFS access via Ranger Hive policies. Ranger RMS must be configured with the names of OZONE, HDFS and Hive services (AKA Repos). In your installation, there may be multiple Ranger services created for OZONE, HDFS and Hive. These can be seen from the Ranger Admin web UI. RMS ACL sync is designed to work on a specific pair of HDFS<->Hive and OZONE<->Hive Ranger services. Therefore, it is important to identify those service names before Ranger RMS is installed. These names should be configured during the installation of Ranger RMS. The default value for Ranger HDFS service name is `cm_hdfs`, for Ranger OZONE service name is `cm_ozone` and for the Ranger Hive service the default name is `cm_hive`.
- Before starting the Ranger RMS installation, ensure that the Hive service identified in the installation above allows the `rangerrms` user select access to all tables in all databases in default (no-zone), as well as in all security-zones for the Hive service.
- In case of custom kerberos principal/user, ensure that the Hive service identified in the installation above allows the `rangerrmsfoo0` (custom) user select access to all tables in all databases in default (no-zone), as well as in all security-zones for the Hive service.
- By default, Ranger RMS tracks only external tables in Hive. To configure Ranger RMS to also track managed Hive tables, add the following configuration setting to Ranger RMS.

```
ranger-rms.HMS.map.managed.tables=true
```



Note:

Locations behind managed tables are granted Read access only, even if the users have Write access at the Hive table level. However, this restriction is not enforced for the `hive` and `impala` users.

- To enable RMS for HDFS authorization, go to Cloudera Manager, select HDFS > Configuration > HDFS Service Advanced Configuration Snippet (Safety Valve) for `ranger-hdfs-security.xml`, then confirm the following settings:

```
ranger.plugin.hdfs.chained.services = cm_hive
ranger.plugin.hdfs.chained.services.cm_hive.impl = org.apache.ranger.chain
edplugin.hdfs.hive.RangerHdfsHiveChainedPlugin
```

- To enable RMS for OZONE authorization, go to Cloudera Manager, select OZONE > Configuration > OZONE Manager Advanced Configuration Snippet (Safety Valve) for `ozone-conf/ranger-ozone-security.xml`, then use +Add to add the following properties:

```
ranger.plugin.ozone.chained.services = cm_hive
ranger.plugin.ozone.chained.services.cm_hive.impl = org.apache.ranger.cha
inedplugin.ozone.hive.RangerOzoneHiveChainedPlugin
```



Note: If any of these configurations are changed after Ranger RMS is started and has synchronized with Hive Metastore, the only way to have Ranger RMS use a new configuration is by performing a “RMS Full Sync”. Please refer to the related topic, ["How to full sync the Ranger RMS database"](#) for specific steps.

Advanced configurations

HDFS plugin side configurations

- `ranger.plugin.hdfs.mapping.hive.authorize.with.only.chained.policies`
 - `true`: Enforce strict Sentry semantics.
 - `false`: If there is no applicable Hive policy, let HDFS determine access.
 - Default setting: `true`
- `ranger.plugin.hdfs.accesstype.mapping.read`
 - A comma-separated list of hive access types that HDFS "read" maps to.
 - Default setting: `select`
- `ranger.plugin.hdfs.accesstype.mapping.write`
 - A comma-separated list of hive access types that HDFS "write" maps to.
 - Default setting: `update,alter`
- `ranger.plugin.hdfs.accesstype.mapping.execute`
 - A comma-separated list of hive access types that HDFS "execute" maps to.
 - Default setting: `_any`
- `ranger.plugin.hdfs.mapping.source.download.interval`
 - The time in milliseconds between mappings download requests from the HDFS Ranger plugin to RMS.
 - Default setting: 30 seconds

By default, the RMS plugin checks for new mapping downloads every 30 seconds, based on this configuration. If you have mapping data (found in the `hdfs_cm_hive_resource_mapping.json` file) of approximately 360MB file size; then performing this operation every 30 seconds could cause an excessive load on the NameNode. After enabling performance logs, we can observe that `saveToCache` takes 11 seconds and `loadFromCache` operations take 7 seconds to complete. The cacheing process takes approximately 18~19 seconds to complete, as shown in the following example performance logs:

```
DEBUG org.apache.ranger.perf.resourcemapping.init: [PERF] RangerMappingR
efresher.loadFromCache(serviceName=cm_hive): 7449
DEBUG org.apache.ranger.perf.resourcemapping.init: [PERF] RangerMappin
gRefresher.saveToCache(serviceName=cm_hive): 11787
```

In this case, you should adjust the frequency of download RMS mappings to at least $18 * 2 = 36$ seconds. A more conservative value = 45 seconds. In this way, you can tune RMS configurations to optimize performance in the NameNode plugin.

Hive service configuration

- `ranger.plugin.audit.excluder.users`
 - This configuration, added in the Hive service-configs, lists the users whose access to Hive or Hive Metastore does not generate audit records. There may be a large number of audit records created when "rangerrms" makes requests to the Hive Metastore when downloading Hive table data. By adding the "rangerrms" user to the comma-separated list of users in this configuration, such audit records will not be generated.

OZONE plugin side configurations

- `ranger.plugin.ozone.mapping.source.download.interval`
 - The time in milliseconds between mappings download requests from the OZONE Ranger plugin to RMS.
 - Default setting: 30 seconds
- `ranger.plugin.ozone.privileged.user.names`
 - Default setting : `admin,dpprofiler,hue,beacon,hive,impala`

- `ranger.plugin.ozone.service.names`
 - Default setting : `hive,impala`



Note: The comma-separated lists that you define for ozone privileged user names and service names are users that, based on default hive policies, have all access for all Hive resources. Therefore, for these users, checking Hive policies when they access storage locations which map to Hive resources is unnecessary, and may cause access violations if masking/row-filtering policies are configured for public group.

OZONE to HIVE access type mapping for Table and Database

Ozone ACL	HIVE ACL For Table	HIVE ACL For Database
read	select	_any
write	update	update
create	create	create
list	select	select
delete	drop	drop
read_acl	select	select
write_acl	update	update

- `ranger.plugin.ozone.accesstype.mapping.<ozone_acl>`
 - This config is used to map Ozone access types to Hive Table access types. Modify the config `<ozone_acl>` as per the ozone access mentioned “OZONE to HIVE access type mapping for Table and Database”
 - Please refer to the “HIVE ACL for Table” column for default configuration with respect to Ozone ACL.
 - For example :


```
ranger.plugin.ozone.accesstype.mapping.read = select
```
- `ranger.plugin.ozone.db.accesstype.mapping.<ozone_acl>`
 - This config is used to map Ozone access types to Hive Database access types. Modify the config `<ozone_acl>` as per the ozone access mentioned “OZONE to HIVE access type mapping for Table and Database”
 - Please refer to the “HIVE ACL for Database” column for default configuration with respect to Ozone ACL.
 - For example :


```
ranger.plugin.ozone.db.accesstype.mapping.read = _any
```

RMS side configurations



Note: Changes to any of these requires that RMS is stopped, all rows are deleted from RMS database table `x_rms_mapping_provider` and RMS is restarted. On restart, RMS downloads complete table data from HMS, which may take a significant amount of time depending on the number of tables in HMS.

- `ranger-rms.HMS.source.service.name`
 - The Ranger HDFS service name (default: `cm_hdfs`).
- `ranger-rms.HMS.target.service.name`
 - The Ranger Hive service name (default: `cm_hive`).
- `ranger-rms.HMS.map.managed.tables`
 - `true` – Track managed and external tables.
 - `false` – Track only external tables.
 - Default setting: `false`

- `ranger-rms.polling.notifications.frequency.ms`
 - The time in milliseconds between polls from RMS to HMS for changes to tables.
 - Default setting: 30 seconds
- `ranger-rms.HMS.source.service.name.ozone`
 - The Ranger OZONE service name
 - Default setting : `cm_ozone`
- `ranger-rms.supported.uri.scheme`
 - A comma-separated list of uri schemes supported by RMS
 - Default setting : `hdfs,ofs,o3fs`

Related Tasks

[How to full sync the Ranger RMS database](#)

Related Information

[Installing Ranger RMS](#)

Configuring HDFS plugin to view permissions through getfacl interface

You can configure the HDFS plugin to view user access permissions in a manner similar to the HDFS `getfacl` command after migrating from CDH to CDP.

Introduction

Ranger RMS (Hive-HDFS ACL Sync) feature replaces a similar feature previously supported in Sentry. Ranger RMS (Hive-HDFS ACL Sync) is implemented through RMS module which consists of RMS server and chained-plugin and corresponding storage plugin - HDFS/Ozone. Ranger RMS (Hive-HDFS ACL Sync) supports viewing access permissions on storage locations across:

- Ranger policies for storage locations
- Ranger policies for Hive entities
- any native storage ACLs

Ranger RMS (Hive-HDFS ACL Sync) is considered essential for those customers migrating from CDH to CDP clusters who also used Sentry ACL-sync feature.

Ranger RMS (Hive-HDFS ACL Sync) addresses the following customer requirements to support:

- an operational need for end-users with no access to Ranger-admin to check the permissions on a HDFS resource
- a tool-set that can verify if the Sentry access permissions are correctly migrated

Sentry implemented an HDFS command `getfacl` that provided access permissions on a given resource to the end-user in a structured, easy-to-understand form. Ranger RMS (Hive-HDFS ACL Sync) is designed to offer the same user experience to the end-users.

Design

The Ranger policy-engine supports an API that, for a resource, returns a static view of the permissions allowed for all users, groups and roles. The view of the access permissions is specific to the component's access model. In other words, for Hive resource the permissions as defined in the Hive service definition, and for HDFS resource, the view consists of read/write/execute permissions. This functionality is leveraged further, by extending it to chained plugins.

Ranger RMS (Hive-HDFS ACL Sync) internally chains HDFS plugin with Hive plugin to provide composite evaluation of access for an HDFS location that maps to a Hive entity. With this extension, the static determination of accesses for a resource considers any chained plugins (and corresponding policy-engine) in its computation. This also

requires conversion between Hive access types and HDFS access types, and combining of accesses by multiple (in this case HDFS and Hive) plugins.

The front-end for accessing this static permission map is the `getfacl` command supported by HDFS. Using this front-end obviates the need for any additional configuration set-up (other than one that enables/disables this feature). It also gives Sentry users (almost) identical experience when interrogating accesses supported for a resource.

Configuration

There are two new configuration parameters introduced for this feature. Both are specified using the safety valve mechanism on HDFS (`ranger-hdfs-security.xml`).

1. Go to Cloudera Manager HDFS Configuration .
2. In Search, type `ranger-hdfs-security.xml`.
3. In HDFS Service Advanced Configuration Snippet (Safety Valve) for `ranger-hdfs-security.xml`, add the following properties/values:

`is.facl.interception.enabled=true`

(Default:false): If true, this feature is activated.

`ranger.plugin.hdfs.policyengine.option.disable.role.resolution=true`

(Default:true) : If false, the roles are resolved to their constituent users and groups for reporting the access permissions.

4. Click Save Changes (CTRL+S).
5. Restart services with stale configurations.

Limitations

Ranger evaluates access based on static and dynamic (time-based, accessing IP-based, condition-based) policies. The static view of the access permissions given only a resource, is not a complete and always accurate picture. However, for the customers moving from Sentry to Ranger, this may not be an issue.

Ranger RMS (Hive-HDFS ACL-Sync) Use Cases

This topic presents a few common use cases for Ranger RMS (Hive-HDFS ACL-Sync).

Use Case 1: RMS Hive policies control access to a table's HDFS directories

Prerequisites:

1. Create a "Customer" Hive table under the default database.
2. Create a "unixuser1" user.
3. User "unixuser1" does not have any policy to allow it access to table "Customer".
4. User "unixuser1" tries to access the Hive data through the `hdfs` command.

Before setting up RMS:

If HDFS ACLs allow access to the location for Customer table, access will be granted to "unixuser1". The audit log will have "hadoop-acl" as the access enforcer.

After setting up RMS:

Access will not be granted to user "unixuser1". The audit log will not specify denying policy.

Use Case 2: RMS Hive policies propagate tag-based access control on tables to HDFS directories

Prerequisites:

1. Create a "Customer" Hive table under the default database.

2. Create a "unixuser1" user.
3. The tag "SPECIAL_ACCESS" is associated with the "Customer" table.
4. A policy for the tag "SPECIAL_ACCESS" provides Hive select access to "unixuser1".
5. User "unixuser1" tries to read the Hive data through the hdfs command.

Before setting up RMS:

If HDFS ACLs allow access to the location for "Customer" table, access will be granted to "unixuser1". The audit log will have "hadoop-acl" as the access enforcer.

After setting up RMS:

Access will be granted by tag-based policy for "SPECIAL_ACCESS".

Use Case 3: RMS Hive policies propagate tag-based masking on tables and denies access to HDFS directories

Prerequisites:

1. Create a "Customer" Hive table under the default database.
2. Create a "unixuser1" user.
3. The tag "SPECIAL_ACCESS" is associated with the "Customer" table.
4. A policy for the tag "SPECIAL_ACCESS" provides Hive select access to "unixuser1".
5. A masking policy for the "Customer" table is set up so that for "unixuser1" a column "SSN" is redacted.
6. User "unixuser1" tries to read the Hive data through the hdfs command.

Before setting up RMS:

If HDFS ACLs allow access to the location for Customer table, access will be granted to "unixuser1". The audit log will have "hadoop-acl" as the access enforcer.

After setting up RMS:

Access will be denied by the masking policy.

Use Case 4: RMS Hive policies take precedence over HDFS policies

Prerequisites:

1. Create a "Customer" Hive table under the default database.
2. Create a "unixuser1" user.
3. User "unixuser1" has a HDFS policy allowing read access.
4. User "unixuser1" does not have any policy to allow it access to the "Customer" table.
5. User "unixuser1" tries to access the Hive data through the hdfs command.

Before setting up RMS:

Access will be granted by the Ranger HDFS policy.

After setting up RMS:

Access will not be granted to the "unixuser1" user. The audit log will not specify a denying policy.