

Cloudera Runtime 7.2.13

## Ranger Auditing

Date published: 2020-07-28

Date modified: 2021-12-13

# 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

<b>Audit Overview.....</b>	<b>4</b>
<b>Managing Auditing with Ranger.....</b>	<b>4</b>
Viewing audit details.....	5
Viewing audit metrics.....	7
Creating a read-only Admin user (Auditor).....	10
Configuring Ranger audit properties for Solr.....	11
Configuring Ranger audit properties for HDFS.....	12
Triggering HDFS audit files rollover.....	13
<b>Ranger Audit Filters.....</b>	<b>14</b>
Default Ranger audit filters.....	15
Configuring a Ranger audit filter policy.....	18
How to set audit filters in Ranger Admin Web UI.....	21
Filter service access logs from Ranger UI.....	22
<b>Excluding audits for specific users, groups, and roles.....</b>	<b>24</b>
<b>Configuring Ranger audits to show actual client IP address.....</b>	<b>25</b>

## Audit Overview

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.

Ranger plugins support storing audit data to multiple audit destinations.

### Solr

The Solr audit destination is a short term audit destination (with a default TTL of 90 days) managed by Solr which can be configured by a Ranger Admin user. The Ranger Admin Web UI displays the access audit data from the audit data stored in Solr.

### HDFS

The HDFS audit destination is a long term audit destination for archival/compliance purposes. The HDFS audit destination has no default retention/purge period. A customer must manage the storage/retention/purge/archival of audit data stored in HDFS manually.

### Related Information

[Configuring Ranger audit properties for Solr](#)

[Configuring Ranger audit properties for HDFS](#)

## Managing Auditing with Ranger

You can manage auditing using the Audit page in the Ranger Admin Web UI.

To explore options for auditing policies, click Audit in the top menu of the Ranger Admin Web UI.

Policy ID	Policy Version	Event Time	Application	User	Service (Name / Type)	Resource (Name / Type)	Access Type	Permission	Result	Access Enforce
25	1	11/16/2022 04:09:52 PM	kafka	streamsmgmr	cm_kafka kafka	__smm-app consumergroup	consume	consume	Allowed	ranger-acl
26	1	11/16/2022 04:09:51 PM	kafka	streamsmgmr	cm_kafka kafka	__KafkaCruiseControl... topic	describe_configs	describe_configs	Allowed	ranger-acl
26	1	11/16/2022 04:09:51 PM	kafka	streamsmgmr	cm_kafka kafka	__smm-app-smm-co... topic	describe_configs	describe_configs	Allowed	ranger-acl
26	1	11/16/2022 04:09:51 PM	kafka	streamsmgmr	cm_kafka kafka	__smm-app-smm-pro... topic	describe_configs	describe_configs	Allowed	ranger-acl
26	1	11/16/2022 04:09:51 PM	kafka	streamsmgmr	cm_kafka kafka	__smm-app-smm-pro... topic	describe_configs	describe_configs	Allowed	ranger-acl
26	1	11/16/2022 04:09:51 PM	kafka	streamsmgmr	cm_kafka kafka	__smm-app-smm-co... topic	describe_configs	describe_configs	Allowed	ranger-acl
36	1	11/16/2022 04:09:51 PM	kafka	streamsmgmr	cm_kafka kafka	connect-offsets topic	describe_configs	describe_configs	Allowed	ranger-acl
26	1	11/16/2022 04:09:51 PM	kafka	streamsmgmr	cm_kafka kafka	__CruiseControlMetrics topic	describe_configs	describe_configs	Allowed	ranger-acl
26	1	11/16/2022 04:09:51 PM	kafka	streamsmgmr	cm_kafka kafka	__smm_producer_me... topic	describe_configs	describe_configs	Allowed	ranger-acl

Seven tabs sub-divide the Audit page:

- Access
- Admin
- Login sessions
- Plugins
- Plugin Status
- User Sync
- Metrics

## Viewing audit details

How to view policy and audit log details in Ranger audits.

### Procedure

To view policy details for a specific audit log, click Access Policy ID .

### Audit > Access: hbasemaster

The screenshot displays the Ranger Access Manager interface. The top navigation bar includes 'Ranger', 'Access Manager', 'Audit', 'Security Zone', and 'Settings'. The user 'admin' is logged in. The main content area shows a table of audit logs with columns for Policy ID, Policy Version, Event Time, Application, User, Service (Name, Type), and Resource (Name). A blue arrow points from the 'Access Policy ID' column of the first row to the 'Policy Details' modal window on the right.

The 'Policy Details' modal window shows the following information:

- Service Name: cm\_hbase
- Service Type: hbase
- Policy Type: Access
- Policy ID: 5
- Version: 1
- Policy Name: all - table, column-family, column
- Policy Labels: HBase Table, HBase Column-family, HBase Column, Description, Audit Logging
- Allow Condition: Select Role, Select Group, Select User, Permissions, Delegate Admin
- Exclude from Allow Conditions: Select Role, Select Group, Select User, Permissions, Delegate Admin
- Deny All Other Accesses: FALSE
- Deny Condition: Select Role, Select Group, Select User, Permissions, Delegate Admin
- Exclude from Deny Conditions: Select Role, Select Group, Select User, Permissions, Delegate Admin
- Updated By: Admin, Updated On: 11/16/2022 09:11 AM, Created By: Admin, Created On: 11/14/2022 09:11 AM

### Audit > Access: HadoopSQL



**Note:** The Hive plugin audit handler now logs UPDATE operations as INSERT, UPDATE, DELETE, and TRUNCATE specifically.

Ranger Access Manager Audit Security Zone Settings admin

Access Admin Login Sessions Plugins Plugin Status User Sync Metrics

SEARCH SERVICE NAME: Hadoop SQL

Exclude Service Users  Last Updated Time: 08/02/2022 10:28:59 AM Entries: 1 to 14 of 14 Columns

Policy ID	Policy Version	Event Time	Application	User	Service (Name / Type)	Resource (Name / Type)	Access Type	Permission	Result	Access Enforcer	Agent Host Name	Client IP	Cluster Name	Zone Name	Event Count	Tags
--	--	08/02/2022 12:48:02 PM	hiveServer2	hrt_1	Hadoop SQL Hadoop SQL	test_db_dkxawj/test_table... @table	INSERT	updates	Denied	ranger-acl	quasar-lowyd-1.quasar-lowyd.r...	172.27.33.69	Cluster 1	Zone Name	1	--
9	1	08/02/2022 12:47:32 PM	hiveServer2	hrt_0a	Hadoop SQL Hadoop SQL	test_db_dkxawj/test_table... @table	INSERT	updates	Allowed	ranger-acl	quasar-lowyd-2.quasar-lowyd.r...	172.27.33.69	Cluster 1	Zone Name	1	--
--	--	08/02/2022 12:47:01 PM	hiveServer2	hrt_1	Hadoop SQL Hadoop SQL	test_db_dkxawj/test_table... @table	TRUNCATE	updates	Denied	ranger-acl	quasar-lowyd-1.quasar-lowyd.r...	172.27.33.69	Cluster 1	Zone Name	1	--
9	1	08/02/2022 12:46:30 PM	hiveServer2	hrt_0a	Hadoop SQL Hadoop SQL	test_db_dkxawj/test_table... @table	TRUNCATE	updates	Allowed	ranger-acl	quasar-lowyd-2.quasar-lowyd.r...	172.27.33.69	Cluster 1	Zone Name	1	--
--	--	08/02/2022 12:46:12 PM	hiveServer2	hrt_1	Hadoop SQL Hadoop SQL	test_db_dkxawj/test_table... @column	UPDATE	updates	Denied	ranger-acl	quasar-lowyd-1.quasar-lowyd.r...	172.27.33.69	Cluster 1	Zone Name	1	--
9	1	08/02/2022 12:45:46 PM	hiveServer2	hrt_0a	Hadoop SQL Hadoop SQL	test_db_dkxawj/test_table... @column	UPDATE	updates	Allowed	ranger-acl	quasar-lowyd-2.quasar-lowyd.r...	172.27.33.69	Cluster 1	Zone Name	1	--
9	1	08/02/2022 12:45:46 PM	hiveServer2	hrt_0a	Hadoop SQL Hadoop SQL	test_db_dkxawj/test_table... @table	SELECT	select	Allowed	ranger-acl	quasar-lowyd-2.quasar-lowyd.r...	172.27.33.69	Cluster 1	Zone Name	1	--
--	--	08/02/2022 12:45:16 PM	hiveServer2	hrt_1	Hadoop SQL Hadoop SQL	test_db_dkxawj/test_table... @table	DELETE	updates	Denied	ranger-acl	quasar-lowyd-1.quasar-lowyd.r...	172.27.33.69	Cluster 1	Zone Name	1	--
9	1	08/02/2022 12:44:46 PM	hiveServer2	hrt_0a	Hadoop SQL Hadoop SQL	test_db_dkxawj/test_table... @table	DELETE	updates	Allowed	ranger-acl	quasar-lowyd-2.quasar-lowyd.r...	172.27.33.69	Cluster 1	Zone Name	1	--
9	1	08/02/2022 12:44:46 PM	hiveServer2	hrt_0a	Hadoop SQL Hadoop SQL	test_db_dkxawj/test_table... @table	SELECT	select	Allowed	ranger-acl	quasar-lowyd-2.quasar-lowyd.r...	172.27.33.69	Cluster 1	Zone Name	1	--
--	--	08/02/2022 12:44:16 PM	hiveServer2	hrt_1	Hadoop SQL Hadoop SQL	test_db_dkxawj/test_table... @table	INSERT	updates	Denied	ranger-acl	quasar-lowyd-1.quasar-lowyd.r...	172.27.33.69	Cluster 1	Zone Name	1	--
9	1	08/02/2022 12:43:35 PM	hiveServer2	hrt_0a	Hadoop SQL Hadoop SQL	test_db_dkxawj/test_table... @table	INSERT	updates	Allowed	ranger-acl	quasar-lowyd-2.quasar-lowyd.r...	172.27.33.69	Cluster 1	Zone Name	1	--

Audit > Admin: Create

Ranger Access Manager Audit Security Zone Settings admin

Access Admin Login Sessions Plugins Plugin Status User Sync Metrics

SEARCH for your access logs...

Last Updated Time: 11/16/2022 06:22:21 PM Entries: 1 to 25 of 182 Columns

Operation	Audit Type	User	Date ( Pacific Standard Time )	Actions	Session ID
Service updated cm_kms	Ranger Service		11/14/2022 09:24:26 AM	Update	
User updated om	Ranger User	rangerusersync	11/14/2022 09:22:57 AM	Update	1
User created scm	Ranger User	rangerusersync	11/14/2022 09:22:57 AM	Create	1
User updated rangertagsync	Ranger User	rangerusersync	11/14/2022 09:22:57 AM	Update	1
User profile updated rangertagsync	User Profile	rangerusersync	11/14/2022 09:22:57 AM	Update	1
User created recon	Ranger User	rangerusersync	11/14/2022 09:22:57 AM	Create	1
User created dn				Create	1
User created rangeradmin				Create	1
User created s3g				Create	1
Policy created all - schema-group, schema-m				Create	52
Policy created all - registry-service				Create	52
Policy created all - schema-group, schema-m				Create	52
Policy created all - schema-group, schema-m				Create	52
Policy created all - serde				Create	52
Policy created all - export-import				Create	52
Service created cm_schema-registry				Create	52
Policy created grant-1668446165876				Create	39
Policy created grant-1668446165585				Create	38
Policy created all - database	Ranger Policy	admin	11/14/2022 09:11:43 AM	Create	18
Policy created all - database, table, column	Ranger Policy	admin	11/14/2022 09:11:43 AM	Create	18

**Operation : create**

Name: recon  
Date: 11/14/2022 09:22:57 AM Pacific Standard Time  
Created By: rangerusersync

**User Details:**

Fields	New Value
Login ID	recon
User Role	User
Other Attributes	{"sync_source":"Unix","full_name":"recon","original_name":"recon"}
Sync Source	Unix

OK

## Audit > User Sync: Sync details

The screenshot shows the Ranger Admin Web UI interface. At the top, there are navigation tabs: Access Manager, Audit, Security Zone, and Settings. Below these are sub-tabs: Access, Admin, Login Sessions, Plugins, Plugin Status, User Sync, and Metrics. The 'User Sync' tab is active. A search bar contains 'START DATE: 11/16/2022'. The main content area displays a table with columns: User Name, Sync Source, Number Of New (Users, Groups), Number Of Modified (Users, Groups), Event Time, and Sync Details. The table lists multiple entries for 'rangerusersync' with 'Unix' as the sync source. A modal window titled 'Sync Details' is open, showing a table with columns 'Name' and 'Value'. The modal window contains the following data:

Name	Value
Unix	nss
File Name	/etc/passwd
Sync time	11/17/2022 01:30:49 AM
Last modified time	01/01/1970 12:00:00 AM
Minimum user id	500
Minimum group id	0
Total number of users synced	65
Total number of groups synced	96
Total number of users marked for delete	0
Total number of groups marked for delete	0

## Viewing audit metrics

How to view audit metrics information using the Ranger Admin Web UI.

### About this task

Metrics provides a high-level view of audit logs as they generate and update in Ranger. Ranger captures audit metrics throughput from the following Ranger services:

- Atlas
- HBase
- Hdfs
- Hive
- Impala
- Kafka
- Knox
- Kudu
- NiFi
- Schema-registry
- Solr
- Streams Messaging Manager
- Yarn

**Procedure**

1. To view audit metrics, in the Ranger Admin Web UI, click Audit Metrics .

The screenshot shows the Ranger Admin Web UI with the 'Metrics' tab selected. The top navigation bar includes 'Ranger', 'Access Manager', 'Audit', 'Security Zone', and 'Settings'. The user is logged in as 'admin'. Below the navigation, there are tabs for 'Access', 'Admin', 'Login Sessions', 'Plugins', 'Plugin Status', 'User Sync', and 'Metrics'. A search bar is present with the placeholder text 'Search for your user sync audits...'. The main content area displays a table of audit metrics with the following columns: Service Name, Service Type, Application Type, Cluster Name, Client IP, Service Status, Metrics Details, and Metrics Graph. The table lists 12 services, all with a status of 'Enabled'. At the bottom of the page, there is a license notice: 'Licensed under the Apache License, Version 2.0'.

Service Name	Service Type	Application Type	Cluster Name	Client IP	Service Status	Metrics Details	Metrics Graph
cm_hdfs	hdfs	hdfs	Cluster 1	172.27.13.135	Enabled	Metrics	Metrics Graph
cm_hdfs	hdfs	hdfs	Cluster 1	172.27.206.70	Enabled	Metrics	Metrics Graph
cm_hdfs	hdfs	hdfs	Cluster 1	172.27.15.128	Enabled	Metrics	Metrics Graph
cm_hbase	hbase	hbaseMaster	Cluster 1	172.27.13.135	Enabled	Metrics	Metrics Graph
cm_hbase	hbase	hbaseRegional	Cluster 1	172.27.206.70	Enabled	Metrics	Metrics Graph
cm_hbase	hbase	hbaseRegional	Cluster 1	172.27.15.128	Enabled	Metrics	Metrics Graph
cm_hbase	hbase	hbaseRegional	Cluster 1	172.27.13.135	Enabled	Metrics	Metrics Graph
cm_yarn	yarn	yarn	Cluster 1	172.27.13.135	Enabled	Metrics	Metrics Graph
cm_hive	hive	hiveServer2	Cluster 1	172.27.13.135	Enabled	Metrics	Metrics Graph
cm_kafka	kafka	kafka	Cluster 1	172.27.13.135	Enabled	Metrics	Metrics Graph
cm_kafka	kafka	kafka	Cluster 1	172.27.15.128	Enabled	Metrics	Metrics Graph



2. To view metrics details for a specific service, click Metrics.

The screenshot shows the Ranger Admin console interface. At the top, there are navigation tabs: Access, Admin, Login Sessions, Plugins, Plugin Status, User Sync, and Metrics. Below the tabs is a search bar and a table of services. The table has columns: Service Name, Service Type, Application Type, Cluster Name, Client IP, Service Status, Metrics Details, and Metrics Graph. A modal window titled "Metrics Text" is open, showing a table with the following data:

Name	Value
metrics	{*PER MINUTE*:*3*}

The "Metrics" link in the "Metrics Details" column of the first row in the main table is highlighted with an orange box, and a blue arrow points to it.

Licensed under the Apache License, Version 2.0

- To view hourly or daily metrics as a graphic for a specific service, click Metrics Graph.

The screenshot shows the Ranger Admin console interface. At the top, there are navigation tabs: Access, Admin, Login Sessions, Plugins, Plugin Status, User Sync, and Metrics. Below the tabs is a search bar with the placeholder text "Search for your user sync audits...". The main content area displays a table of services with columns: Service Name, Service Type, Application Type, Cluster Name, Client IP, Service Status, Metrics Details, and Metrics Graph. The table lists various services like cm\_hdfs, cm\_hbase, cm\_yarn, cm\_hive, and cm\_kafka. A modal window titled "Metric Details" is open, showing a bar chart titled "Audit Metrics By Day". The chart has a Y-axis ranging from 0 to 4000 and an X-axis showing dates from 2022-11-14 to 2022-11-17. The chart shows audit metrics for each day, with a peak around 3800 on 2022-11-15 and 2022-11-16. The modal window also includes a close button (X) and an OK button.

Service Name	Service Type	Application Type	Cluster Name	Client IP	Service Status	Metrics Details	Metrics Graph
cm_hdfs	hdfs	hdfs	Cluster 1	172.27.13.135	Enabled	Metrics	Metrics Graph
cm_hdfs	hdfs	hdfs					Metrics Graph
cm_hdfs	hdfs	hdfs					Metrics Graph
cm_hbase	hbase	hbase					Metrics Graph
cm_hbase	hbase	hbase					Metrics Graph
cm_hbase	hbase	hbase					Metrics Graph
cm_hbase	hbase	hbase					Metrics Graph
cm_hbase	hbase	hbase					Metrics Graph
cm_yarn	yarn	yarn					Metrics Graph
cm_hive	hive	hiveS					Metrics Graph
cm_kafka	kafka	kafka					Metrics Graph
cm_kafka	kafka	kafka					Metrics Graph

## Creating a read-only Admin user (Auditor)

Creating a read-only Admin user (Auditor) enables compliance activities because this user can monitor policies and audit events, but cannot make changes.

### About this task

When a user with the Auditor role logs in, they see a read-only view of Ranger policies and audit events. An Auditor can search and filter on access audit events, and access and view all tabs under Audit to understand access events. They cannot edit users or groups, export/import policies, or make changes of any kind.

### Procedure

- Select Settings > Users/Groups/Roles.
- Click Add New User.

- Complete the **User Detail** section, selecting Auditor as the role:

The screenshot shows the Ranger 'User Create' interface. The 'User Detail' section includes the following fields:

- User Name \*: auditor1
- New Password \*: [masked]
- Password Confirm \*: [masked]
- First Name \*: Audrey
- Last Name: [empty]
- Email Address: [empty]
- Select Role \*: Auditor
- Group: Please select (with a '+' button)

Below the 'Sync Details' section, there is a table with columns 'Name' and 'Value'. The table is currently empty, displaying 'No Sync Details Found!!'. At the bottom of the form are 'Save' and 'Cancel' buttons.

- Click Save.

## Configuring Ranger audit properties for Solr

How to change the default time settings that control how long Ranger keeps audit data collected by Solr.

### About this task

The Solr audit destination is intended to store short term audit records .You can configure parameters that control how much data collected by Solr that Ranger will store for auditing purposes.

**Table 1: Ranger Audit Configuration Parameters for Solr**

Parameter Name	Description	Default Setting	Units
ranger.audit.solr.config.ttl	Time To Live for Solr Collection of Ranger Audits	90	days
ranger.audit.solr.config.delete.trigger	Auto Delete Period in seconds for Solr Collection of Ranger Audits for expired documents	1	days (configurable)



**Note:** "Time To Live for Solr Collection of Ranger Audits" is also known as the Max Retention Days attribute.

### Procedure

- From Cloudera Manager choose Ranger Configuration .
- In Search, type ranger.audit.solr.config, then press Return.
- In ranger.audit.solr.config.ttl, set the the number of days to keep audit data.
- In ranger.audit.solr.config.delete.trigger set the number and units (days, minutes, hours, or seconds) to keep data for expired documents

5. Refresh the configuration, using one of the following two options:
  - a) Click Refresh Configuration, as prompted or, if Refresh Configuration does not appear,
  - b) In Actions, click Update Solr config-set for Ranger, then confirm.

## Configuring Ranger audit properties for HDFS

How to change the settings that control how Ranger writes audit records to HDFS.

### About this task

The HDFS audit destination is intended to store long-term audit records.

You can configure whether Ranger stores audit records in HDFS and at which location.

You must purge long term audit records stored in HDFS manually.

**Table 2: Ranger Audit Configuration Parameters for HDFS**

Parameter Name	Description	Default Setting	Units
ranger_plugin_hdfs_audit_enabled	controls whether Ranger writes audit records to HDFS	true	T/F
ranger_plugin_hdfs_audit_url	location at which you can access audit records written to HDFS	<hdfs.host_name> ranger/audit	string



**Note:** You can also disable storing ranger audit data to hdfs in each service specifically by setting `xasecure.audit.destination.hdfs=false` in that service.

### Procedure

1. From Cloudera Manager choose Ranger Configuration .
2. In Search, type `ranger_plugin`, then press Return.
3. In `ranger_plugin_hdfs_audit_enabled`, check/uncheck RANGER-1 (Service Wide)
4. In `ranger_plugin_hdfs_audit_url` type a valid directory on the hdfs host.
5. Refresh the configuration, using one of the following two options:
  - a) Click Refresh Configuration, as prompted or, if Refresh Configuration does not appear,
  - b) In Actions, click Update Solr config-set for Ranger, then confirm.

### What to do next

(Optional)

You may want to delete older logs from HDFS. Cloudera provides no feature to do this. You may accomplish this task manually, using a script.



**Note:** The following example script is not supported by Cloudera. It is shown for reference only. You must test this successfully in a test environment before implementing it in a production cluster.

You must specify the audit log directory by replacing the 2nd line `hdfs dfs -ls /<path_to>/<audit_logs>` in the example script.

You may also include the `-skipTrash` option, if you choose, on 7th line in the script.

```
#####
today=`date +%s`
hdfs dfs -ls /<path_to>/<audit_logs> | grep "^d" | while read line ; do
```

```

dir_date=$(echo ${line} | awk '{print $6}')
difference=$(( ( ${today} - $(date -d ${dir_date} +%s) ) / ( 24*60*60 ) ))
filePath=$(echo ${line} | awk '{print $8}')

if [ ${difference} -gt 30 ]; then
    hdfs dfs -rm -r $filePath
fi
done
#####

```

### Related Information

[How to do a cleanup of hdfs files older than a certain date using a bash script](#)

## Triggering HDFS audit files rollover

How to configure when HDFS audit files close for each service.

### About this task

By default, the Ranger Audit framework closes audit files created in HDFS or other cloud storage inline with audit event triggers. In other words, when an audit event occurs, Ranger checks the configured rollout time and then closes the file if the threshold has reached. Default audit rollout time is 24 hours. If no audit event occurs in a 24 hour period, files remain open beyond the 24 hour period. In some environments, audit log analysis that encounter an audit file open beyond the current date can cause system exceptions. If you want the files to be closed every day, so that the audit log file will have only that day's log and the next day's log will be in the next day's file, you can configure the audit framework to close files every day. To do this, you must add several configuration parameters to the ranger-`<service_name>-audit.xml` (safety valve) file for each service, using Cloudera Manager.

### Procedure

1. From Cloudera Manager choose `<service_name>` Configuration .
2. In `<service_name>` Configuration Search , type `ranger-<service_name>`, then press Return.
3. In `<service_name>` Server Advanced Configuration Snippet (Safety Valve) for `ranger-<service_name>-audit.xml`, do the following steps:
  - a) Click + (Add).
  - b) In Name, type `xasecure.audit.destination.hdfs.file.rollover.enable.periodic.rollover`
  - c) In Value, type `true`.

When this is enabled Ranger Audit Framework will spawn a Scheduler thread which monitors the occurrence of closing threshold and closes the file. By default every night the file gets closed.

- d) Click + (Add another).
- e) In Name, type `xasecure.audit.destination.hdfs.file.rollover.sec`
- f) In Value, type an integer value in seconds.

This is the time in seconds when the file has to be closed. The default value is 86400 sec (1 day) which triggers the file to be closed at midnight and opens a new audit log for the next day. You can override the default value

can be overridden by setting this parameter. For example, if you set the value 3600 (1 hr), the file gets closed every hour.

- g) Click + (Add another).
- h) In Name, type `xasecure.audit.destination.hdfs.file.rollover.periodic.rollover.check.sec`
- i) In Value, type an integer value in seconds.

This is the time frequency of the check to be done whether the threshold time for rollover has occurred. By default the check is done every 60 secs. You can configure this parameter to delay the check time.

**Figure 1: Example: Hive service configured to trigger rollover of hdfs audit files**

The screenshot shows the Cloudera Ranger Admin web UI for the HIVE-1 service. The 'Configuration' tab is active, displaying a configuration snippet for 'ranger-audit\_safety\_valve'. The snippet contains three parameters:

- Name:** `xasecure.audit.destination.hdfs.file.rollover.enable.periodic.rollover`, **Value:** `true`
- Name:** `xasecure.audit.destination.hdfs.file.rollover.sec`, **Value:** `3600`
- Name:** `xasecure.audit.destination.hdfs.file.rollover.periodic.rollover.check.sec`, **Value:** `3600`

The interface also shows a 'Filters' sidebar on the left and a 'Save Changes (CTRL+S)' button at the bottom right.

- j) Click Save Changes (CTRL+S).

4. Repeat steps 1-3 for each service.
5. Restart the service.

## Ranger Audit Filters

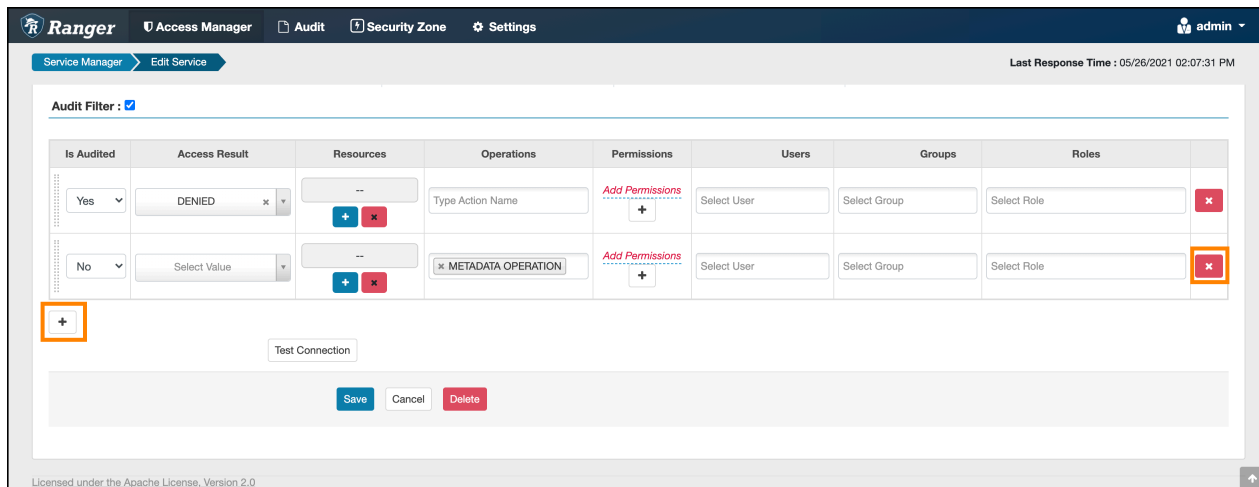
You can use Ranger audit filters to control the amount of audit log data collected and stored on your cluster.

### About Ranger audit filters

Ranger audit filters allow you to control the amount of audit log data for each Ranger service. Audit filters are defined using a JSON string that is added to each service configuration. The audit filter JSON string is a simplified form of the Ranger policy JSON. Audit filters appear as rows in the Audit Filter section of the Edit Service view for each service. The set of audit filter rows defines the audit log policy for the service. For example, the default audit log policy for the Hadoop SQL service appears in Ranger Admin web UI Service Manager Edit Service when you scroll down to Audit Filter. Audit Filter is checked (enabled) by default. In this example, the top row defines an audit filter that causes all instances of "access denied" to appear in audit logs. The lower row defines a filter that causes no

metadata operations to appear in audit logs. These two filters comprise the default audit filter policy for Hadoop SQL service.

**Figure 2: Default audit filter policy for the Hadoop SQL service**




## Default Ranger audit filters

Default audit filters for the following Ranger service appear in Edit Services and may be modified as necessary by Ranger Admin users.

### HDFS

**Figure 3: Default audit filters for HDFS service**



### Hbase

**Figure 4: Default audit filters for the Hbase service**

Audit Filter: [🔍](#)

Is Audited	Access Result	Resources	Operations	Permissions	Users	Groups	Roles
Yes	DENIED	...	Type Action Name	Add Permissions +	Select User	Select Group	Select Role
No	Select Value	table:"ROOTS","META","_ddl","hbase:meta,hbase:acl,default,hbase	Type Action Name	Add Permissions +	in hbase	Select Group	Select Role
No	Select Value	table:atlas_janus,ATLAS_ENTITY_AUDIT_EVENTS column-family:" column:"	Type Action Name	Add Permissions +	in atlas in hbase	Select Group	Select Role
No	Select Value	...	in balance	Add Permissions +	in hbase	Select Group	Select Role

## Hadoop SQL

Figure 5: Default audit filters for the Hadoop SQL service

Audit Filter: [🔍](#)

Is Audited	Access Result	Resources	Operations	Permissions	Users	Groups	Roles
Yes	DENIED	...	Type Action Name	Add Permissions +	Select User	Select Group	Select Role
No	Select Value	...	in METADATA OPERATION	Add Permissions +	Select User	Select Group	Select Role

## Knox

Figure 6: Default audit filters for the Knox service

Audit Filter: [🔍](#)

Is Audited	Access Result	Resources	Operations	Permissions	Users	Groups	Roles
Yes	DENIED	...	Type Action Name	Add Permissions +	Select User	Select Group	Select Role
No	Select Value	...	...	Add Permissions +	in knox	Select Group	Select Role

## Solr

Figure 7: Default audit filters for the Solr service

Audit Filter: [🔍](#)

Is Audited	Access Result	Resources	Operations	Permissions	Users	Groups	Roles
Yes	DENIED	...	Type Action Name	Add Permissions +	Select User	Select Group	Select Role
No	Select Value	...	Type Action Name	Add Permissions +	in hives in hdfs in kafka in hbase in solr in rangeracl in knox in atlas	Select Group	Select Role

## Kafka

Figure 8: Default audit filters for the Kafka service



Is Audited	Access Result	Resources	Operations	Permissions	Users	Groups	Roles
Yes	DENIED	--	Type Action Name	Add Permissions +	Select User	Select Group	Select Role
No	Select Value	topic:ATLAS_ENTITIES, ATLAS_HOOK, ATLAS_SPARK_HOOK	describe, publish, consume	Add Permissions +	atlas	Select Group	Select Role
No	Select Value	topic:ATLAS_HOOK	publish, describe	Add Permissions +	hive, hbase, impala, nifi	Select Group	Select Role
No	Select Value	topic:ATLAS_ENTITIES	consume, describe	Add Permissions +	rangertagsync	Select Group	Select Role
No	Select Value	consumergroup:*	consume	Add Permissions +	atlas, rangertagsync	Select Group	Select Role
No	Select Value	--	Type Action Name	Add Permissions +	kafka	Select Group	Select Role

### Ranger KMS

Figure 9: Default audit filters for the Ranger KMS service

Is Audited	Access Result	Resources	Operations	Permissions	Users	Groups	Roles
Yes	DENIED	--	Type Action Name	Add Permissions +	Select User	Select Group	Select Role
No	Select Value	--	read	Add Permissions +	keyadmin	Select Group	Select Role

### Atlas

Figure 10: Default audit filters for the Atlas service

Is Audited	Access Result	Resources	Operations	Permissions	Users	Groups	Roles
Yes	DENIED	--	Type Action Name	Add Permissions +	Select User	Select Group	Select Role
No	Select Value	--	Type Action Name	Add Permissions +	atlas	Select Group	Select Role

### ADLS

Figure 11: Default audit filters for the ADLS service

Is Audited	Access Result	Resources	Operations	Permissions	Users	Groups	Roles
Yes	DENIED	--	Type Action Name	Add Permissions +	Select User	Select Group	Select Role
No	Select Value	--	get-status, read, list	List, Read	hive, hbase, hdfs	Select Group	Select Role

### Ozone

Figure 12: Default audit filters for the Ozone service

Is Audited	Access Result	Resources	Operations	Permissions	Users	Groups	Roles
Yes	DENIED	--	Type Action Name	Add Permissions +	Select User	Select Group	Select Role
No	Select Value	--	Type Action Name	Add Permissions +	om	Select Group	Select Role

## S3

Figure 13: Default audit filters for the S3 service

Is Audited	Access Result	Resources	Operations	Permissions	Users	Groups	Roles
Yes	DENIED	--	Type Action Name	Add Permissions +	Select User	Select Group	Select Role
No	Select Value	--	read	Add Permissions +	hive hbase hdfs yarn	Select Group	Select Role

## Tag-based services

Figure 14: Default audit filters for a tag-based service

Is Audited	Access Result	Resources	Operations	Permissions	Users	Groups	Roles
Yes	DENIED	--	Type Action Name	Add Permissions +	Select User	Select Group	Select Role



### Note:

Default audit filter policies do not exist for Yarn, NiFi, NiFi Registry, Kudu, or schema registry services.

## Configuring a Ranger audit filter policy

You can configure an audit filter as you add or edit a resource- or tag-based service.

### To configure an audit filter policy:

1. In Ranger Admin web UI Service Manager click Add or Edit for either a resource-, or tag-based service.
2. Scroll down to Audit Filter.
3. Click Audit Filter flag.

You configure a Ranger audit filter policy by adding (+), deleting (X), or modifying each audit filter row for the service.

4. Use the controls in the filter row to edit filter properties. For example, you can configure:

**Is Audited: choose Yes or No**

to include or not include a filter in the audit logs for a service

**Access Result: choose DENIED, ALLOWED, or NOT\_DETERMINED**

to include that access result in the audit log filter

**Resources: Add or Delete a resource item**

to include or remove the resource from the audit log filter

**Operations: Add or Remove an action name**

to include the action/operation in the audit log filter

(click x to remove an existing operation)

**Permissions: Add or Remove permissions**

a. Click + in Permissions to open the Add dialog.

b. Select/Unselect required permissions.

For example, in HDFS service select read, write, execute, or All permissions.

**Users: click Select User to see a list of defined users**

to include one or multiple users in the audit log filter

**Groups: click Select Group to see a list of defined groups**

to include one or multiple groups in the audit log filter

**Roles: click Select Role to see a list of defined roles**

to include one or multiple roles in the audit log filter

### Audit filter details

- When you save the UI selections described in the preceding list, audit filters are defined as a JSON list. Each service references a unique list.
- For example, ranger.plugin.audit.filters for the HDFS service includes:

```
[
  {
    "accessResult": "DENIED",
    "isAudited": true
  },
  {
    "users": [
      "unaudited-user1"
    ],
    "groups": [
      "unaudited-group1"
    ],
    "roles": [
      "unaudited-role1"
    ],
    "isAudited": false
  },
  {
    "actions": [
      "listStatus",
      "getFileinfo"
    ],
    "accessTypes": [
      "execute"
    ],
  },
]
```

```

    "isAudited":false
  },
  {
    "resources":{
      "path":{
        "values":[
          "/audited"
        ],
        "isRecursive":true
      }
    },
    "isAudited":true
  },
  {
    "resources":{
      "path":{
        "values":[
          "/unaudited"
        ],
        "isRecursive":true
      }
    },
    "isAudited":false
  }
]

```

- Each value in the list is an audit filter, which takes the format of a simplified Ranger policy, along with access results fields.
- Audit filters are defined with rules on Ranger policy attributes and access result attributes.
  - Policy attributes: resources, users, groups, roles, accessTypes
  - Access result attributes: isAudited, actions, accessResult
- The following audit filter specifies that accessResult=DENIED will be audited.

The isAudited flag specifies whether or not to audit.

```
{ "accessResult": "DENIED", "isAudited": true }
```

- The following audit filter specifies that “resource => /unaudited” will not be audited.

```
{ "resources": { "path": { "values": [ "/unaudited" ], "isRecursive": true } }, "isAudited": false }
```

- The following audit filter specifies that access to resource database=> sys table=> dump by user “use2” will not be audited.

```
{ "resources": { "database": { "values": [ "sys" ] }, "table": { "values": [ "dump" ] } }, "users": [ "user2" ], "isAudited": false }
```

- The following audit filter specifies that access result in actions => listStatus, getFileInfo and accessType => execute will not be audited.

```
{ "actions": [ "listStatus", "getFileinfo" ], "accessTypes": [ "execute" ], "isAudited": false }
```

- The following audit filter specifies that access by user “superuser1” and group “supergroup1” will not be audited.

```
{ "users": [ "superuser1" ], "groups": [ "supergroup1" ], "isAudited": false }
```

- The following audit filter specifies that access to any resource tagged as NO\_AUDIT will not be audited.

```
{ "resources": { "tag": { "values": [ "NO_AUDIT" ] } }, "isAudited": false }
```

## How to set audit filters in Ranger Admin Web UI

You can set specific audit filter conditions for each service, using Create/Edit Service .

### About this task

Creating audit filters for a service using the Ranger Admin Web UI can prevent audit logs from being sent to destinations like SOLR and HDFS.

### Procedure

1. In the Ranger Admin Web UI Service Manager , click Add New Service or Edit (existing service).
2. On Create/Edit Service, scroll down to Audit Filters.
  - a) Verify that Audit Filter is checked.

Optionally, define any of the following to include in the filter definition:

#### Is Audited

Defines whether audit logs are stored or not.

Is Audited=Yes: stores audit records in the defined audit destination.

Is Audited=No: do not store audit records.

#### Access Results

Denied, Allowed, or Not Determined

select to filter access=denied, access=allowed or all by selecting access=Not determined.

#### Resource

use Resource Details to include or exclude specific resources such as databases, tables, or columns.

#### Operations

select specific operations to filter

#### Permissions

select specific permissions

#### Users, Groups, Roles

select specific users, groups, and roles

- b) Click Save.

**Figure 15: Adding an audit filter that stores user systest, access=Allowed logs for Hive service**

Is Audited	Access Result	Resources	Operations	Permissions	Users	Groups	Roles	
Yes	ALLOWED	+ -	Type Action Name	Create	systest X	Select...	Select...	X

3. Test your filters to verify that defined audit filters perform as expected.

### Results

Defining specific filtering properties can prevent access logs for service users from being stored in the configured audit destination, if Is Audited = No.

## Filter service access logs from Ranger UI

You can limit display of system access/audit log records generated by service users in each service.

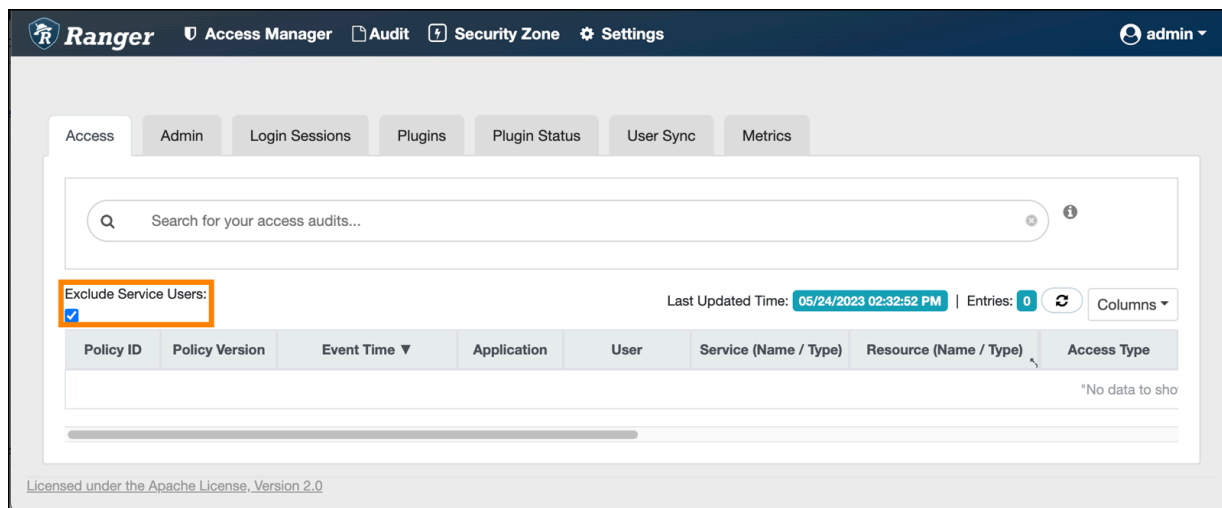
### About this task

This topic describes how to limit the display of access log records on the Access tab in the Ranger Admin Web UI.

### Procedure

1. Go to Ranger Admin Web UI Audit Access .
2. Check the Exclude Service Users box, as shown in:

**Figure 16: Setting the Exclude Service Users flag to true**



The screenshot shows the Ranger Admin Web UI interface. The top navigation bar includes the Ranger logo, 'Access Manager', 'Audit', 'Security Zone', and 'Settings' menus, along with a user profile 'admin'. Below the navigation bar, there are tabs for 'Access', 'Admin', 'Login Sessions', 'Plugins', 'Plugin Status', 'User Sync', and 'Metrics'. The 'Access' tab is selected. A search bar is present with the placeholder text 'Search for your access audits...'. Below the search bar, there is a section for 'Exclude Service Users' with a checked checkbox. To the right of this section, it shows 'Last Updated Time: 05/24/2023 02:32:52 PM' and 'Entries: 0'. Below this, there is a table with columns: 'Policy ID', 'Policy Version', 'Event Time', 'Application', 'User', 'Service (Name / Type)', 'Resource (Name / Type)', and 'Access Type'. The table is currently empty, displaying '\*No data to sho'. At the bottom left, there is a small text: 'Licensed under the Apache License, Version 2.0'.

3. Define specific component services and users for access logs to filter out, in ranger-admin-site.xml.

- a) Go to Cloudera Manager Ranger Configuration
- b) In Search, type ranger-admin-site.
- c) Define the following properties:

**Name**

ranger.plugins.<service\_name>.serviceuser

**Value**

<service\_name>

**Name**

ranger.accesslogs.exclude.users.list

**Value**

user1, user2

**Figure 17: Filtering out service and user logs for Hive service**

Ranger Admin Advanced Configuration Snippet (Safety Valve) for conf/ranger-admin-site.xml View as XML

Ranger Admin Default Group [Undo](#)

[conf/ranger-admin-site.xml\\_role\\_safety\\_valve](#)

<b>Name</b>	<input type="text" value="ranger.accesslogs.exclude.users.list"/>	🗑️ +
<b>Value</b>	<input type="text" value="test1"/>	
<b>Description</b>	<input type="text"/>	
	<input type="checkbox"/> Final	
<b>Name</b>	<input type="text" value="ranger.plugins.hive.serviceuser"/>	🗑️ +
<b>Value</b>	<input type="text" value="hive"/>	
<b>Description</b>	<input type="text"/>	
	<input type="checkbox"/> Final	

4. Click Save Changes (CTRL+S).

5. Restart the Ranger service.

### Results

Setting Exclude Service Users to true and defining specific filtering properties prevents audit logs from service users from appearing on Ranger Admin Web UI Audit Access, but does NOT prevent access logs for service users from being generated in Solr.

## Excluding audits for specific users, groups, and roles

You can exclude audit records for specific users, groups, and roles from each service from appearing in the Ranger UI.

### About this task

Ranger default log functionality creates audit log records for access and authorization requests, specifically around service accounts such as hbase, atlas and solr. Writing so much data to solr can limit the availability of Solr for further usage. This topic describes how to exclude audit records for specific users, groups, and roles from appearing in the Ranger UI. Excluding specific users, groups or roles is also known as creating a blacklist for Ranger audits.

### Procedure

1. In the Ranger Admin Web UI Service Manager , click Add New Service or Edit (existing service).
2. On Create/Edit Service, scroll down to Config Properties Add New Configurations .
3. Remove all audit filters from the existing service.



4. Click +, then type one of the following property names:

- ranger.plugin.audit.exclude.users
- ranger.plugin.audit.exclude.groups
- ranger.plugin.audit.exclude.roles

followed by one or more values.



**Note:** You can include multiple values for each exclude property using a comma-separated list.

**Figure 18: Adding an exclude users property to the HadoopSQL service**

Name	Value	
tag.download.auth.users	hive,hdfs,impala	✕
policy.download.auth.users	hive,hdfs,impala	✕
policy.grantrevoke.auth.users	hive,impala	✕
enable.hive.metastore.lookup	true	✕
default.policy.users	impala,hive,hue,beacon,admin,dp	✕
hive.site.file.path	/etc/hive/conf/hive-site.xml	✕
ranger.plugin.audit.exclude.users	testuser2	✕

After adding the above configuration; if testuser2 user performs any actions for HadoopSQL service, Audit Access logs will not appear in the Ranger UI, but are still sent to Solr.

Similarly, you can exclude (or blacklist) users belonging to a particular group or role by adding a user-specific or role-specific configuration.

## Configuring Ranger audits to show actual client IP address

How to forward the actual client IP address to audit logs generated from a Ranger plugin.

### About this task

Ranger audit logs record the IP address through which Ranger policies grant/authorize access. When Ranger is set up behind a Knox proxy server, the proxy server IP address appears in the audit logs generated for each Ranger plugin. You can configure each plugin to forward the actual client IP address on which that service runs, so that the audit logs for that service more specifically reflect access/authorization activity. You must configure each plugin individually. This topic uses the Hive (Hadoop SQL) service as an example.

## Procedure

1. From Cloudera Manager choose <service\_name> Configuration .
2. In <service\_name> Configuration Search , type ranger-plugin, then press Return.
3. In Ranger Plugin Use X-Forwarded for IP Address, check the box.
4. In Ranger Plugin Trusted Proxy IP Address, type the IP address of the Knox proxy server host.

The screenshot shows the Cloudera Manager interface for configuring the HIVE-1 service. The search bar contains 'ranger plugin'. The configuration list shows three items:

- Ranger Plugin Use X-Forwarded for IP Address**: Checked (HIVE-1 (Service-Wide)).
- Ranger Plugin Trusted Proxy IP Address**: HIVE-1 (Service-Wide). The input field contains 'KnoxServerHost.IP.address'.
- Ranger Plugin URL Auth Filesystem Schemes**: HIVE-1 (Service-Wide). The input field contains 'hdfs;file;wasb;adl:'.

## Results

Hive audit logs will now show the IP address of the host on which Hive service runs.