

Cloudera Runtime 7.2.7

## Ranger Auditing

Date published: 2019-11-01

Date modified:

# CLOUDERA

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

# Legal Notice

© Cloudera Inc. 2025. 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

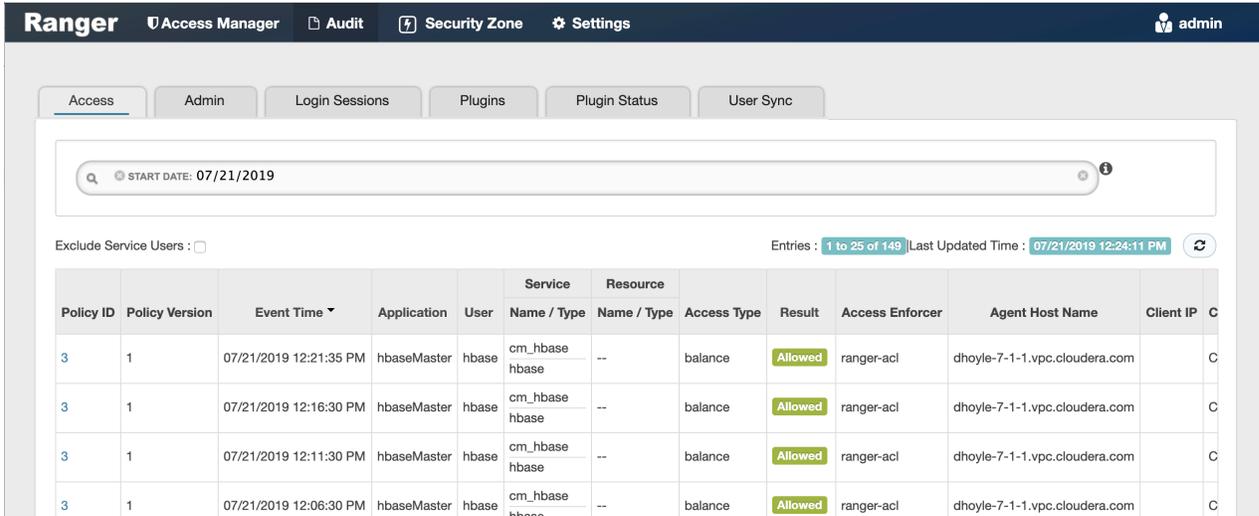
<b>Audit Overview.....</b>	<b>4</b>
<b>Managing Auditing with Ranger.....</b>	<b>4</b>
View audit details.....	4
Create a read-only Admin user (Auditor).....	7
<b>Changing Ranger audit storage location and migrating data.....</b>	<b>8</b>

## 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.

## Managing Auditing with Ranger

To explore options for auditing policies in Ranger, click Audit in the top menu.



The screenshot shows the Ranger web interface. The top navigation bar includes 'Ranger', 'Access Manager', 'Audit', 'Security Zone', and 'Settings'. The user is logged in as 'admin'. Below the navigation bar, there are tabs for 'Access', 'Admin', 'Login Sessions', 'Plugins', 'Plugin Status', and 'User Sync'. The 'Audit' tab is active, displaying a search bar with 'START DATE: 07/21/2019'. Below the search bar, there is a table of audit entries. The table has columns for Policy ID, Policy Version, Event Time, Application, User, Service Name / Type, Resource Name / Type, Access Type, Result, Access Enforcer, Agent Host Name, Client IP, and a 'C' column. The table shows four entries, all with a result of 'Allowed'.

Policy ID	Policy Version	Event Time	Application	User	Service Name / Type	Resource Name / Type	Access Type	Result	Access Enforcer	Agent Host Name	Client IP	C
3	1	07/21/2019 12:21:35 PM	hbaseMaster	hbase	cm_hbase hbase	--	balance	Allowed	ranger-acl	dhoyle-7-1-1.vpc.cloudera.com		C
3	1	07/21/2019 12:16:30 PM	hbaseMaster	hbase	cm_hbase hbase	--	balance	Allowed	ranger-acl	dhoyle-7-1-1.vpc.cloudera.com		C
3	1	07/21/2019 12:11:30 PM	hbaseMaster	hbase	cm_hbase hbase	--	balance	Allowed	ranger-acl	dhoyle-7-1-1.vpc.cloudera.com		C
3	1	07/21/2019 12:06:30 PM	hbaseMaster	hbase	cm_hbase hbase	--	balance	Allowed	ranger-acl	dhoyle-7-1-1.vpc.cloudera.com		C

There are six tabs on the Audit page:

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

## View audit details

How to view operation details in Ranger audits.

### Procedure

To view details for a particular operation, click any tab, then Policy ID, Operation name, or Session ID.

### Audit > Access: HBase Table

The screenshot shows the Ranger interface with the 'Audit' tab selected. The main view displays a table of audit entries for HBase access. A modal window titled 'Policy Details' is open, showing information for a policy named 'cm\_hbase'. The table below shows several 'Allowed' entries for 'cm\_hbase' on 'hbase' resources.

Policy ID	Policy Version	Event Time	Application	User	Service Name / Type	Resource Name / Type	Access Type	Result	Access
3	1	07/21/2019 12:51:30 PM	hbaseMaster	hbase	cm_hbase hbase	--	balance	Allowed	ranger
3	1	07/21/2019 12:46:30 PM	hbaseMaster	hbase	cm_hbase hbase	--	balance	Allowed	ranger
3	1	07/21/2019 12:41:30 PM	hbaseMaster	hbase	cm_hbase hbase	--	balance	Allowed	ranger-acl
3	1	07/21/2019 12:36:30 PM	hbaseMaster	hbase	cm_hbase hbase	--	balance	Allowed	ranger-acl
3	1	07/21/2019 12:31:31 PM	hbaseMaster	hbase	cm_hbase hbase	--	balance	Allowed	ranger-acl
3	1	07/21/2019 12:26:30 PM	hbaseMaster	hbase	cm_hbase hbase	--	balance	Allowed	ranger-acl

### Audit > Admin: Update

The screenshot shows the Ranger interface with the 'Admin' tab selected. The main view displays an audit log table. A modal window titled 'Operation: update' is open, showing details for an update operation on 'tag\_service2'.

Operation	Audit Type	User	Date (Eastern Daylight Time)	Actions	Session Id
Service updated tag_service2	Ranger Service	admin	07/21/2019 01:09:34 PM	Update	40
Group created temp_employees	Ranger Group	admin	07/20/2019 02:15:05 PM	Create	38
Group created audit	Ranger Group	admin	07/18/2019 04:18:42 PM	Create	35
Exported policies	Ranger Policy	admin	07/17/2019 03:06:22 PM	Export, Import	32
Service updated tag_service1	Ranger Service	admin	07/15/2019 04:11:25 PM	Update	
Policy created EXPIRES_ON	Ranger Policy		07/15/2019 04:11:25 PM	Create	
Service created tag_tag	Ranger Service		07/15/2019 04:11:25 PM	Create	

Audit > Admin: Create

The screenshot displays the Ranger Admin interface. At the top, there is a navigation bar with 'Ranger' and menu items: 'Access Manager', 'Audit', 'Security Zone', and 'Settings'. The user 'admin' is logged in. Below the navigation bar, there are tabs for 'Access', 'Admin', 'Login Sessions', 'Plugins', 'Plugin Status', and 'User Sync'. The 'Admin' tab is active, showing a search bar for access logs and a table of audit entries.

The audit log table has the following columns: Operation, Audit Type, User, Date (Eastern Daylight Time), Actions, and Session Id. The entry for 'Security Zone created security-zone2' is highlighted, and a blue arrow points to its 'Create' action button.

An 'Operation : create' modal window is open, showing details for the 'security-zone2' operation. The modal includes the following information:

- Name : security-zone2
- Date : 07/14/2019 05:24:36 PM Eastern Daylight Time
- Created By : admin
- Zone Details:

Fields	New Value
Zone Description	--
Zone Audit User Groups	--
Zone Audit Users	auditor1
Zone Admin User Groups	--
Zone Admin Users	admin
Zone Tag Services	cm_tag
Zone Name	security-zone2

The modal also includes a 'Zone Service Details' section with fields for 'Service Name' and 'Zone Service Resources', and an 'OK' button at the bottom right.

## Audit > User Sync: Sync details

The screenshot shows the Ranger Admin console interface. At the top, there are navigation tabs: Access, Admin, Login Sessions, Plugins, Plugin Status, and User Sync. Below the tabs is a search bar with the text "START DATE: 07/21/2019". To the right of the search bar, it says "Entries: 1 to 25 of 803" and "Last Updated Time: 07/21/2019 01:23:45 PM".

The main content is a table with the following columns: User Name, Sync Source, Number Of New (Users, Groups), Number Of Modified (Users, Groups), Event Time, and Sync Details. The table contains multiple rows for 'rangerusersync' with 'Unix' as the sync source. A blue box highlights the 'Sync Details' icon in the second row, and a blue arrow points to a modal window titled 'Sync Details'.

The 'Sync Details' modal window contains a table with the following data:

Name	Value
Unix	nss
File Name	/etc/passwd
Sync time	07/21/2019 10:21:48 AM
Last modified time	12/31/1969 04:00:00 PM
Minimum user id	500
Minimum group id	0
Total number of users synced	35
Total number of groups synced	39

## Create 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

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

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

The screenshot shows the Ranger web interface for creating a user. The top navigation bar includes 'Ranger', 'Access Manager', 'Audit', 'Security Zone', and 'Settings'. The user 'admin' is logged in. The breadcrumb trail is 'Users/Groups/Roles > User Create'. The 'User Detail' section contains the following fields:

- User Name \*: auditor1
- New Password \*: [masked]
- Password Confirm \*: [masked]
- First Name \*: Audrey
- Last Name: [empty]
- Email Address: [empty]
- Select Role \*: Auditor (selected)
- Group: audit

At the bottom of the form are 'Save' and 'Cancel' buttons.

4. Click Save.

## Changing Ranger audit storage location and migrating data

How to change the location of existing and future Ranger audit data collected by Solr from HDFS to local or from local to HDFS.

### Before you begin

- Stop Atlas from Cloudera Manager.
- If using Kerberos, set the SOLR\_PROCESS\_DIR environment variable.

```
# export SOLR_PROCESS_DIR=$(ls -ldtr /var/run/cloudera-scm-agent/process/
*SOLR_SERVER | tail -1)
```

### About this task

Starting with Cloudera Runtime version 7.1.4 / 7.2.2, the storage location for ranger audit data collected by Solr changed to local file system from HDFS, as was true for previous versions. The default storage location Ranger audit data storage location for Cloudera Runtime-7.1.4+ and Cloudera Runtime-7.2.2+ installations is local file system. After upgrading from an earlier Cloudera platform version, follow these steps to backup and migrate your Ranger audit data and change the location where Solr stores your future Ranger audit records.

- The default value of the index storage in the local file system is /var/lib/solr-infra. You can configure this, using Cloudera Manager Solr Configuration "Solr Data Directory" .
- The default value of the index storage in HDFS is /solr-infra. You can configure this, using Cloudera Manager Solr Configuration "HDFS Data Directory" .

## Procedure

1. Create HDFS Directory to store the collection backups.

As an HDFS super user, run the following commands to create the backup directory:

```
# hdfs dfs -mkdir /solr-backups
# hdfs dfs -chown solr:solr /solr-backups
```

2. Obtain valid kerberos ticket for Solr user.

```
# kinit -kt solr.keytab solr/$(hostname -f)
```

3. Download the configs for the collection.

```
# solrctl instancedir --get ranger_audits /tmp/ranger_audits
# solrctl instancedir --get atlas_configs /tmp/atlas_configs
```

4. Modify the solrconfig.xml for each of the configs for which data needs to be stored in HDFS.

In /tmp/<config\_name>/conf created during Step 3., edit properties in the solrconfig.xml file as follows:

- When migrating your data storage location from local file system to HDFS, replace these two lines:

```
<directoryFactory name="DirectoryFactory"
  class="{solr.directoryFactory:solr.NRTCachingDirectoryFactory}">
<lockType>${solr.lock.type:native}</lockType>
```

with

```
<directoryFactory name="DirectoryFactory"
  class="{solr.directoryFactory:org.apache.solr.core.HdfsDirectoryFactory}">
<lockType>${solr.lock.type:hdfs}</lockType>
```

- When migrating your data storage location from HDFS to local file system, replace these two lines:

```
<directoryFactory name="DirectoryFactory"
  class="{solr.directoryFactory:org.apache.solr.core.HdfsDirectoryFactory}">
<lockType>${solr.lock.type:hdfs}</lockType>
```

with

```
<directoryFactory name="DirectoryFactory"
  class="{solr.directoryFactory:solr.NRTCachingDirectoryFactory}">
<lockType>${solr.lock.type:native}</lockType>
```

5. Update the modified configs in Zookeeper.

```
# solrctl --jaas $SOLR_PROCESS_DIR/jaas.conf instancedir --update
  atlas_configs /tmp/atlas_configs
# solrctl --jaas $SOLR_PROCESS_DIR/jaas.conf instancedir --update
  ranger_audits /tmp/ranger_audits
```

6. Backup the Solr collections.

- When migrating your data storage location from local file system to HDFS, run:

```
# curl -k --negotiate -u : "https://$(hostname
  -f):8995/solr/admin/collections?action=BACKUP&name=vertex_backup&col
  lection=vertex_index&
```

```
location=hdfs://<Namenode_Hostname>:8020/solr-backups"
```

In the preceding command, the important points are name, collection, and location:

**name**

specifies the name of the backup. It should be unique per collection

**collection**

specifies the collection name for which the backup will be performed

**location**

specifies the HDFS path, where the backup will be stored

Repeat the curl command for different collections, modifying the parameters as necessary for each collection.

The expected output would be -

```
"responseHeader": {
  "status": 0,
  "QTime": 10567},
"success": {
  "Solr_Server_Hostname:8995_solr": {
    "responseHeader": {
      "status": 0,
      "QTime": 8959}}}}
```

- When migrating your data storage location from HDFS to local file system:

Refer to Back up a Solr collection for specific steps, and make the following adjustments:

- If TLS is enabled for the Solr service, specify the trust store and password by using the ZKCLI\_JVM\_FLAGS environment variable before you begin the procedure.

```
# export ZKCLI_JVM_FLAGS="-Djavax.net.ssl.trustStore=/path/to/truststore.jks -Djavax.net.ssl.trustStorePassword="
```

- Create Snapshot

```
# solrctl --jaas $SOLR_PROCESS_DIR/jaas.conf collection --create-snapshot <snapshot_name> -c <collection_name>
```

- or use the Solr API to take the backup:

```
curl -i -k --negotiate -u : "https://(hostname -f):8995/solr/admin/collections?action=BACKUP&name=ranger_audits_bkp&collection=ranger_audits&location=/path/to/solr-backups"
```

- Export Snapshot

```
# solrctl --jaas $SOLR_PROCESS_DIR/jaas.conf collection --export-snapshot <snapshot_name> -c <collection_name> -d <destination_directory>
```



**Note:** The <destination\_directory> is a HDFS path. The ownership of this directory should be solr:solr.

## 7. Delete the collections from the original location.

All instances of Solr service should be up, running, and healthy before deleting the collections. Use Cloudera Manager to check for any alerts or warnings for any of the instances. If alerts or warnings exist, fix those before deleting the collection.

```
# solrctl collection --delete edge_index
# solrctl collection --delete vertex_index
# solrctl collection --delete fulltext_index
```

```
# solrctl collection --delete ranger_audits
```

8. Verify that the collections are deleted from the original location.

```
# solrctl collection --list
```

This will give an empty result.

9. Verify that no leftover directories for any of the collections have been deleted.

- When migrating your data storage location from local file system to HDFS:

```
# cd /var/lib/solr-infra
```

Get the value of "Solr Data Directory, using Cloudera Manager Solr Configuration .

```
# ls -ltr
```

- When migrating your data storage location from HDFS to local file system, replace these two lines:

```
# hdfs dfs -ls /solr/<collection_name>
```



**Note:** If any directory name which starts with the collection name deleted in Step 7. exists, delete/ move the directory to another path.

10. Restore the collection from backup to the new location.

Refer to Restore a Solr collection, for more specific steps.

```
# curl -k --negotiate -u : "https://$(hostname
-f):8995/solr/admin/collections?
action=RESTORE&name=<Name_of_backup>&location=hdfs:/
<<Namenode_Hostname>:8020/solr-backups&collection=<Collection_Name>"
```

```
# solrctl collection --restore ranger_audits
-l hdfs://<Namenode_Hostname>:8020/solr-backups
-b ranger_backup -i ranger1
```

The request id must be unique for each restore operation, as well as for each retry.

To check the status of restore operation:

```
# solrctl collection --request-status <requestId>
```



**Note:** If the Atlas Collections (vertex\_index, fulltext\_index and edge\_index) restore operations fail, restart the solr service and rerun the restore command. Now, the restart operations should complete successfully.

11. Verify the Atlas & Ranger functionality.

Verify that both Atlas and Ranger audits functions properly, and that you can see the latest audits in Ranger Web UI and latest lineage in Atlas.

- To verify Atlas audits, create a test table in Hive, and then query the collections to see if you are able to view the data.
- You can also query the collections every 20-30 seconds (depending on how other services utilize Atlas/Ranger), and verify if the "numDocs" value increases at every query.

```
# curl -k --negotiate -u : "https://$(hostname -f):8995/solr/edge_index/
select?q=%3A*&wt=json&ident=true&rows=0"
# curl -k --negotiate -u : "https://$(hostname -f):8995/solr/vertex_index/
select?q=%3A*&wt=json&ident=true&rows=0"
# curl -k --negotiate -u : "https://$(hostname -f):8995/solr/
fulltext_index/select?q=%3A*&wt=json&ident=true&rows=0"
# curl -k --negotiate -u : "https://$(hostname -f):8995/solr/
ranger_audits/select?q=%3A*&wt=json&ident=true&rows=0"
```