Configuring Apache Ranger Authentication with UNIX, LDAP, or AD

Date published:
Date modified:

https://docs.cloudera.com/

Legal Notice

Contents

Configuring Ranger Authentication with UNIX, LDAP, AD, or PAM

This section describes how to configure the authentication method that determines who is allowed to log in to the Ranger web UI. The options are local UNIX, LDAP, AD, or PAM.

Note: In CDP Public Cloud, identity management is provided by FreeIPA, and configured using the Management Console. Therefore for CDP Public Cloud you should leave the Admin Authentication Method set to the UNIX authentication settings. For more information on FreeIPA, see Managing FreeIPA in the Identify Management documentation.

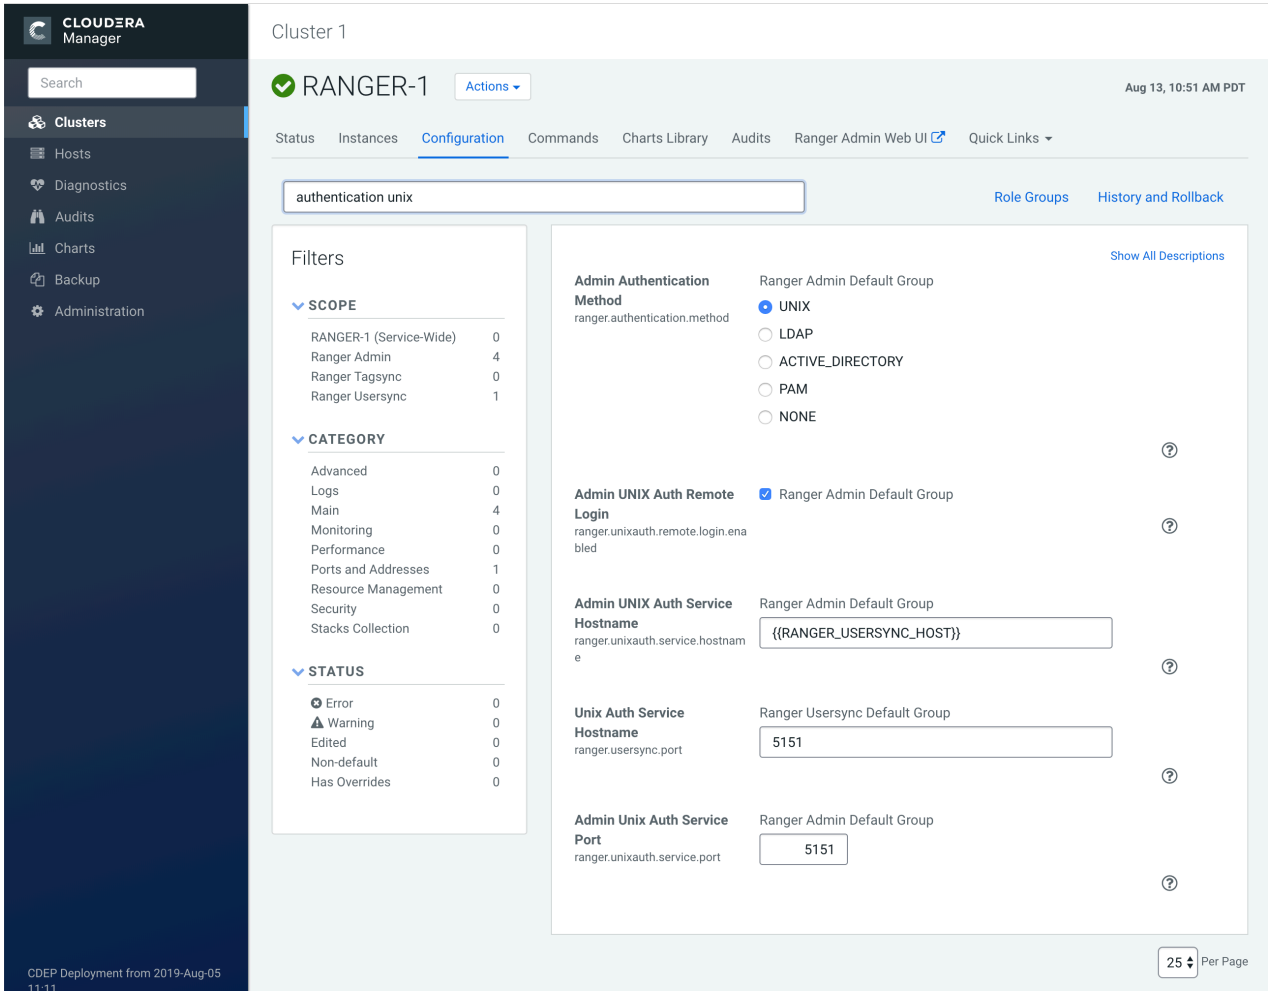

Related Information

Cloudera Management Console

CDP Cloud Management Console: Managing user access and authorization

Managing FreeIPA

Configure Ranger authentication for UNIX

How to configure Ranger to use UNIX for user authentication.

About this task

Note: In CDP Public Cloud, identity management is provided by FreeIPA, and configured using the Management Console. Therefore for CDP Public Cloud you should leave the Admin Authentication Method set to the UNIX authentication settings. For more information on FreeIPA, see Managing FreeIPA in the Identify Management documentation.

Procedure

1. In Cloudera Manager, select Ranger, then click the Configuration tab.
2. To display the UNIX authentication settings, type "authentication unix" in the Search box.

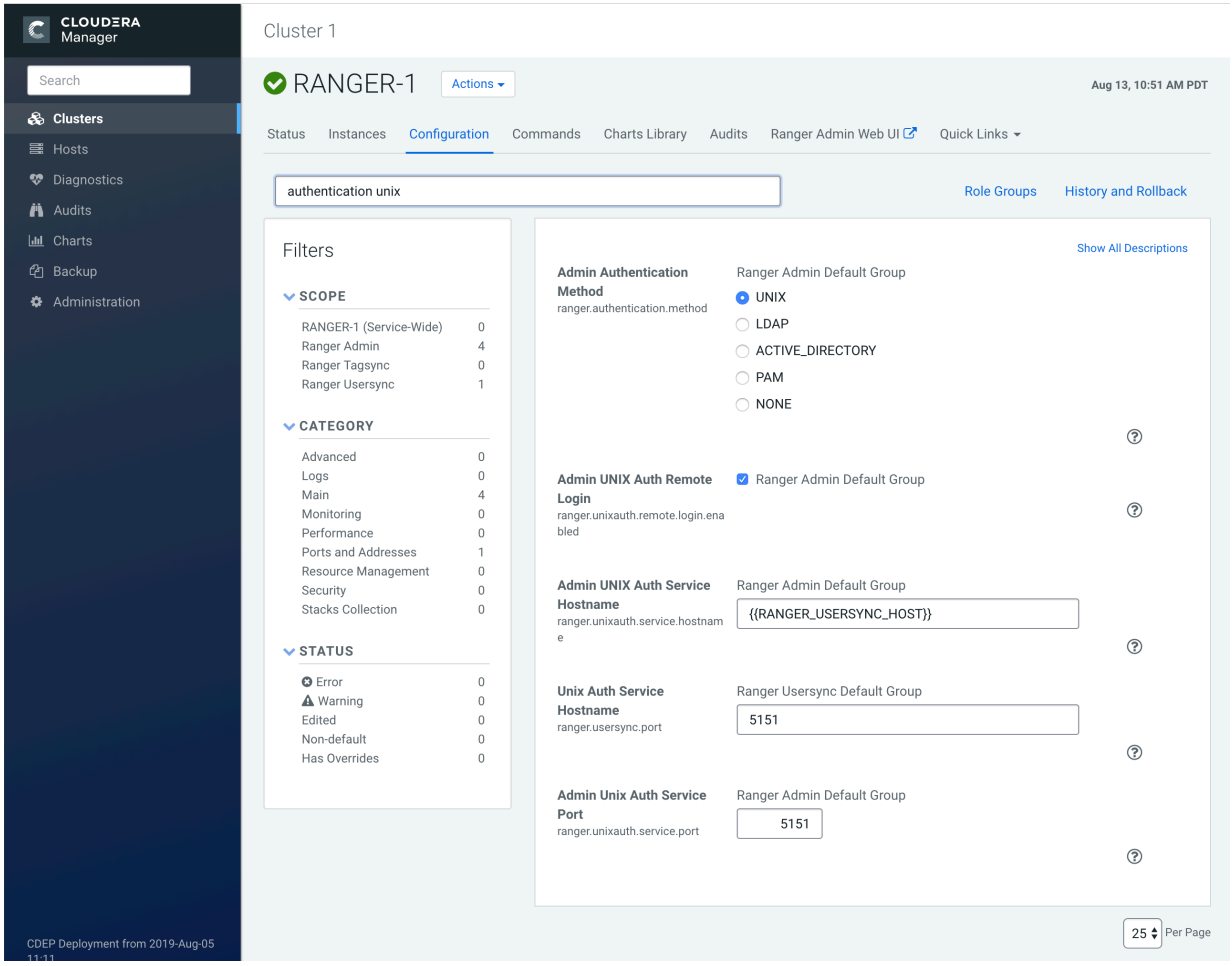

3. Configure the following settings for UNIX authentication, then click Save Changes.

Table 1: UNIX Authentication Settings

Configuration Property	Description	Default Value	Example Value	Requi
Admin Authentication Method	The Ranger authentication method.	UNIX	UNIX	Yes, t auther
Allow remote Login	Flag to enable/disable remote login. Only used if the Authentication method is UNIX.	TRUE	TRUE	No.

Configuration Property	Description	Default Value	Example Value	Requi
ranger.unixauth.service.hostname	The FQDN of the host where the UNIX authentication service is running. Only used if the Authentication method is UNIX. {{RANGER_USERSYNC_HOST}} is a placeholder value that is replaced with the host where Ranger Usersync is installed in the cluster.	localhost	myunixhost.domain.com	Yes, if selecte
ranger.unixauth.service.port	The port number where the ranger-usersync module is running the UNIX Authentication Service.	5151	5151	Yes, if selecte

Related Information

Cloudera Management Console

Configure Ranger authentication for AD

How to configure Ranger to use Active Directory (AD) for user authentication.

About this task

Note: In CDP Public Cloud, identity management is provided by FreeIPA, and configured using the Management Console. Therefore for CDP Public Cloud you should leave the Admin Authentication Method set to the UNIX authentication settings. For more information on FreeIPA, see Managing FreeIPA in the Identify Management documentation.

Procedure

1. Select Cloudera Manager Ranger Configuration , type authentication in Search.
Ranger authentication property settings display. You may need to scroll down to see the AD settings.

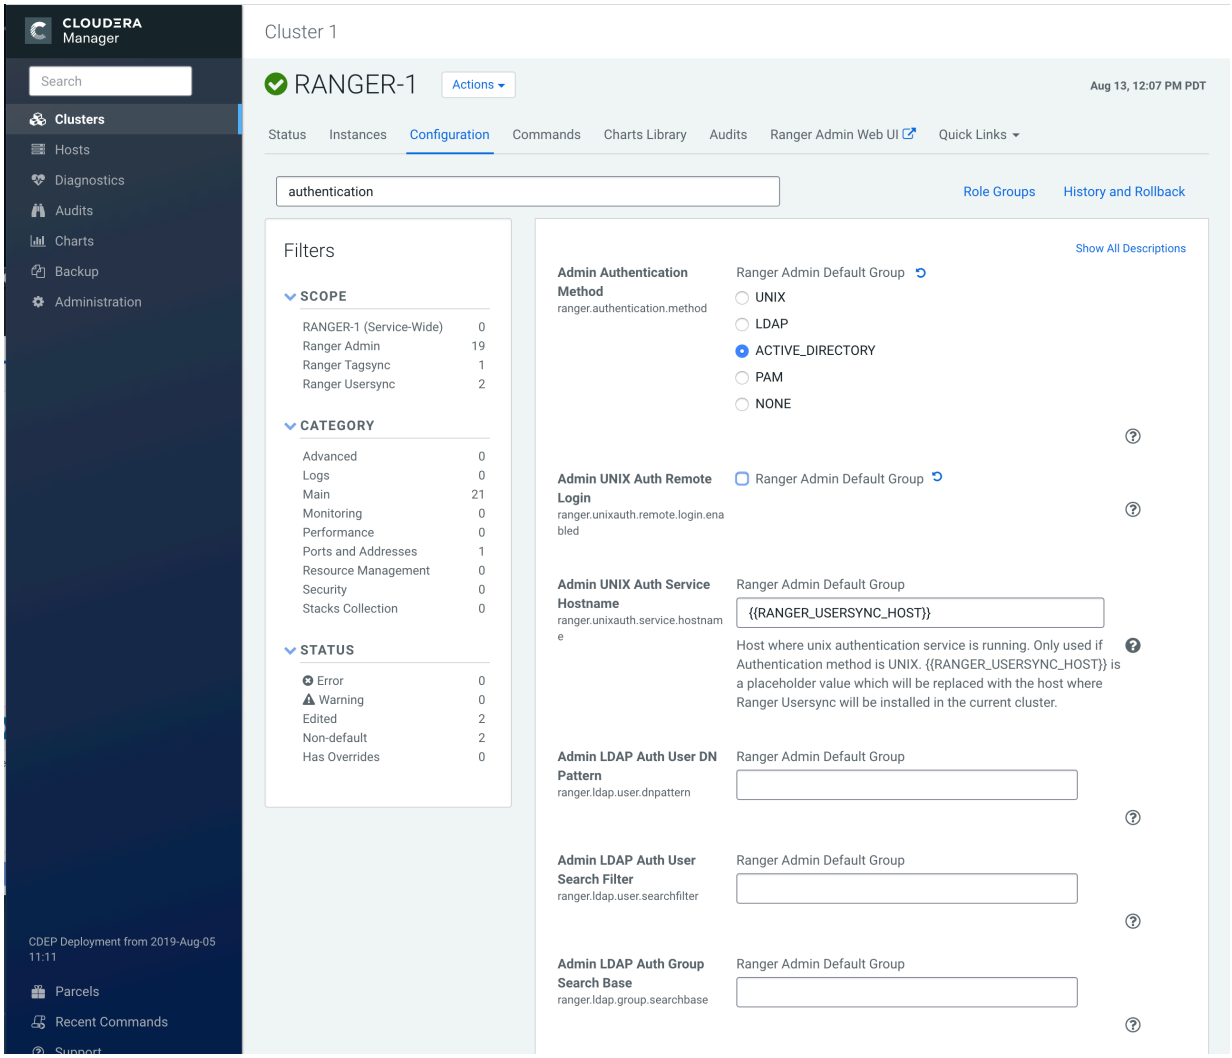

2. Configure the following settings for AD authentication, then click Save Changes.

Property	Description	Default value	Sample values
Admin Authentication Method	The Ranger authentication method.	UNIX	ACTIVE_DIRECTORY
Admin AD Auth Base DN ranger.ldap.ad.base.dn	The Distinguished Name (DN) of the starting point for directory server searches.	N/A	dc=example,dc=com
Admin AD Auth Bind DN ranger.ldap.ad.bind.dn	The full Distinguished Name (DN), including Common Name (CN) of an LDAP user account that has privileges to search for users.	N/A	cn=adadmin,cn=Users,dc=example,dc=com
Admin AD Auth Bind Password ranger.ldap.ad.bind.password	Password for the bind.dn.	N/A	Secret123!
Admin AD Auth Domain Name ranger.ldap.ad.domain	The domain name of the AD Authentication service.	N/A	example.com

Property	Description	Default value	Sample values
Admin AD Auth Referral ranger.ldap.ad.referral*	See below.	ignore	follow | ignore | throw
Admin AD Auth URL ranger.ldap.ad.url	The AD server URL, for example: ldap://<AD-Servername>Port	N/A	ldap://<AD-Servername>Port
Admin AD Auth User Search Filter ranger.ldap.ad.user.searchfilter	AD user search filter.	N/A	

* There are three possible values for ranger.ldap.ad.referral:

- follow
- throw
- ignore

The recommended setting is: follow.

When searching a directory, the server might return several search results, along with a few continuation references that show where to obtain further results. These results and references might be interleaved at the protocol level.

When ranger.ldap.ad.referral is set to follow:

The AD service provider processes all of the normal entries first, and then follows the continuation references.

When ranger.ldap.ad.referral is set to throw:

All of the normal entries are returned in the enumeration first, before theReferralException is thrown.

By contrast, a referral error response is processed immediately when this property is set to follow or throw.

When ranger.ldap.ad.referral is set to ignore:

The server should return referral entries as ordinary entries (or plain text). This might return partial results for the search. In the case of AD, a PartialResultException is returned when referrals are encountered while search results are processed.

Related Information

Cloudera Management Console

Configure Ranger authentication for LDAP

How to configure Ranger to use LDAP for user authentication.

About this task

Note: In CDP Public Cloud, identity management is provided by FreeIPA, and configured using the Management Console. Therefore for CDP Public Cloud you should leave the Admin Authentication Method set to the UNIX authentication settings. For more information on FreeIPA, see Managing FreeIPA in the Identify Management documentation.

Procedure

1. In Cloudera Manager, select Ranger, then click the Configuration tab.

2. To display the authentication settings, type "authentication" in the Search box. You may need to scroll down to see all of the LDAP settings.

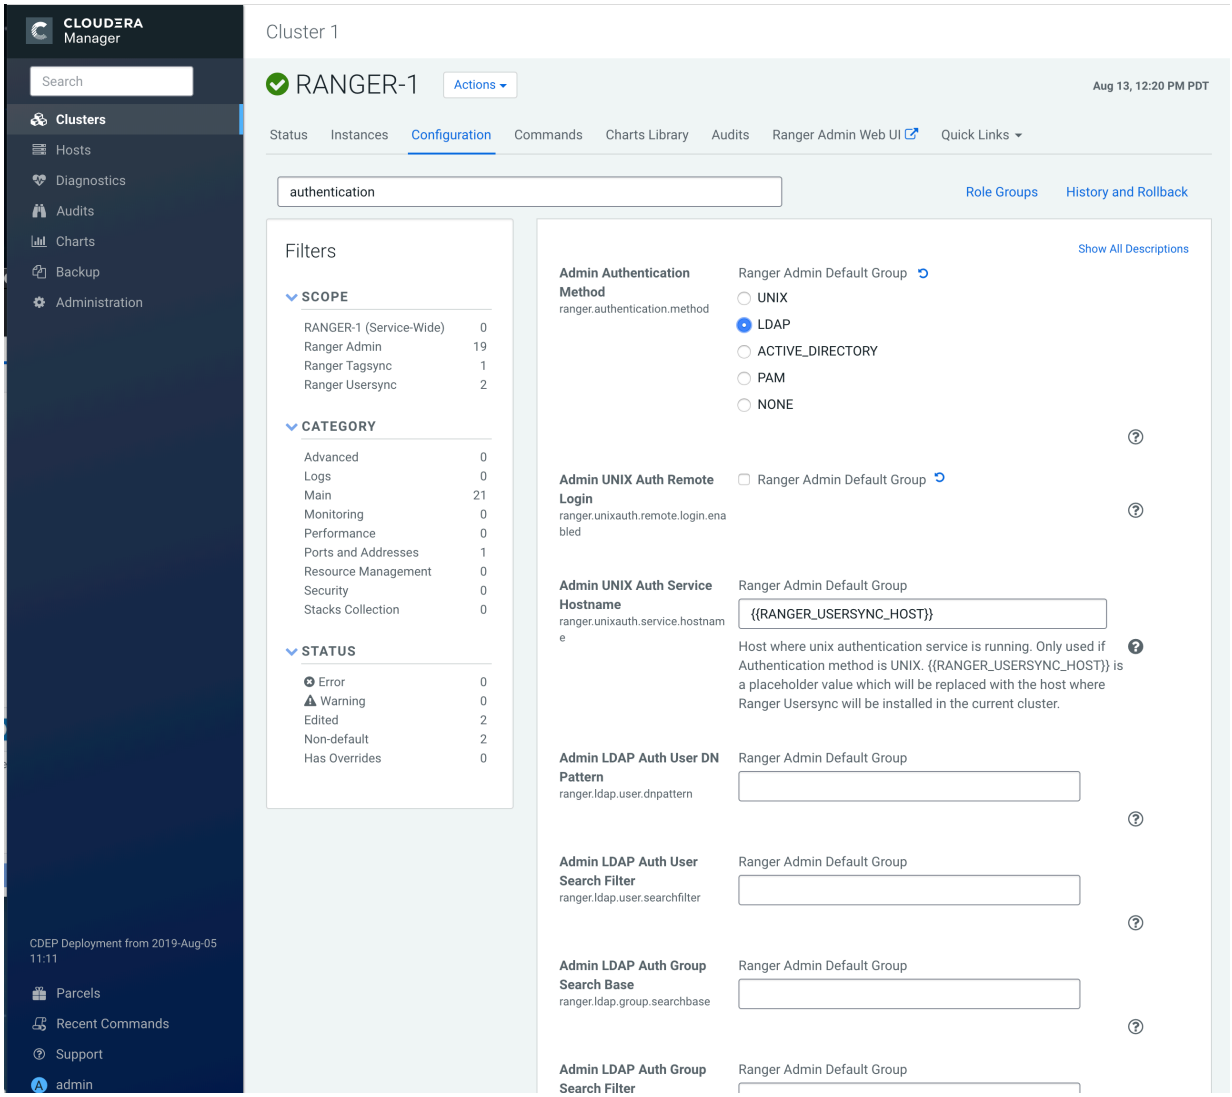

3. Configure the following settings for LDAP authentication, then click Save Changes.

Property	Required ?	Description	Default value	Sample values
Admin Authentication Method	Required	The Ranger authentication method.	UNIX	LDAP
Admin LDAP Auth Group Search Base ranger.ldap.group.searchbase	Optional	The LDAP group search base.	N/A	(|(CN=Hdp_users)(CN=Hdp_admins))
Admin LDAP Auth Group Search Filter ranger.ldap.group.searchfilter	Optional	The LDAP group search filter.	N/A	
Admin LDAP Auth URL ranger.ldap.url	Required	The LDAP server URL	N/A	ldap://localhost:389 or ldaps://localhost:636

Property	Required ?	Description	Default value	Sample values
Admin LDAP Auth Bind User ranger.ldap.bind.dn	Required	Full distinguished name (DN), including common name (CN), of an LDAP user account that has privileges to search for users. This user is used for searching the users. This could be a read-only LDAP user.	N/A	cn=admin,dc=example,dc=com
Admin LDAP Auth Bind User Password ranger.ldap.bind.password	Required	Password for the account that can search for users.	N/A	Secret123!
Admin LDAP Auth User Search Filter ranger.ldap.user.searchfilter	Required	The LDAP user search filter.	N/A	
Admin LDAP Auth Base DN ranger.ldap.base.dn	Required	The Distinguished Name (DN) of the starting point for directory server searches.	N/A	dc=example,dc=com
Admin LDAP Auth Group Role Attribute ranger.ldap.group.roleattribute	Optional	The LDAP group role attribute.	N/A	cn
Admin LDAP Auth Referral ranger.ldap.referral*	Required	See below.	ignore	follow | ignore | throw
Admin LDAP Auth User DN Pattern ranger.ldap.user.dnpattern	Required	The LDAP user DN.	N/A	uid={0},ou=users,dc=xasecure,dc=net

* There are three possible values for ranger.ldap.ad.referral: follow, throw, and ignore. The recommended setting is follow.

When searching a directory, the server might return several search results, along with a few continuation references that show where to obtain further results. These results and references might be interleaved at the protocol level.

- When this property is set to follow, the AD service provider processes all of the normal entries first, and then follows the continuation references.
- When this property is set to throw, all of the normal entries are returned in the enumeration first, before the ReferralException is thrown. By contrast, a "referral" error response is processed immediately when this property is set to follow or throw.
- When this property is set to ignore, it indicates that the server should return referral entries as ordinary entries (or plain text). This might return partial results for the search. In the case of AD, a PartialResultException is returned when referrals are encountered while search results are processed.

Related Information

Cloudera Management Console

Configure Ranger authentication for PAM

How to configure Ranger to use PAM for user authentication.

About this task

Note: In CDP Public Cloud, identity management is provided by FreeIPA, and configured using the Management Console. Therefore for CDP Public Cloud you should leave the Admin Authentication Method set to the UNIX authentication settings. For more information on FreeIPA, see Managing FreeIPA in the Identify Management documentation.

Procedure

1. In Cloudera Manager, select Ranger, then click the Configuration tab.
2. Under Admin Authentication Method, select PAM, then click Save Changes.

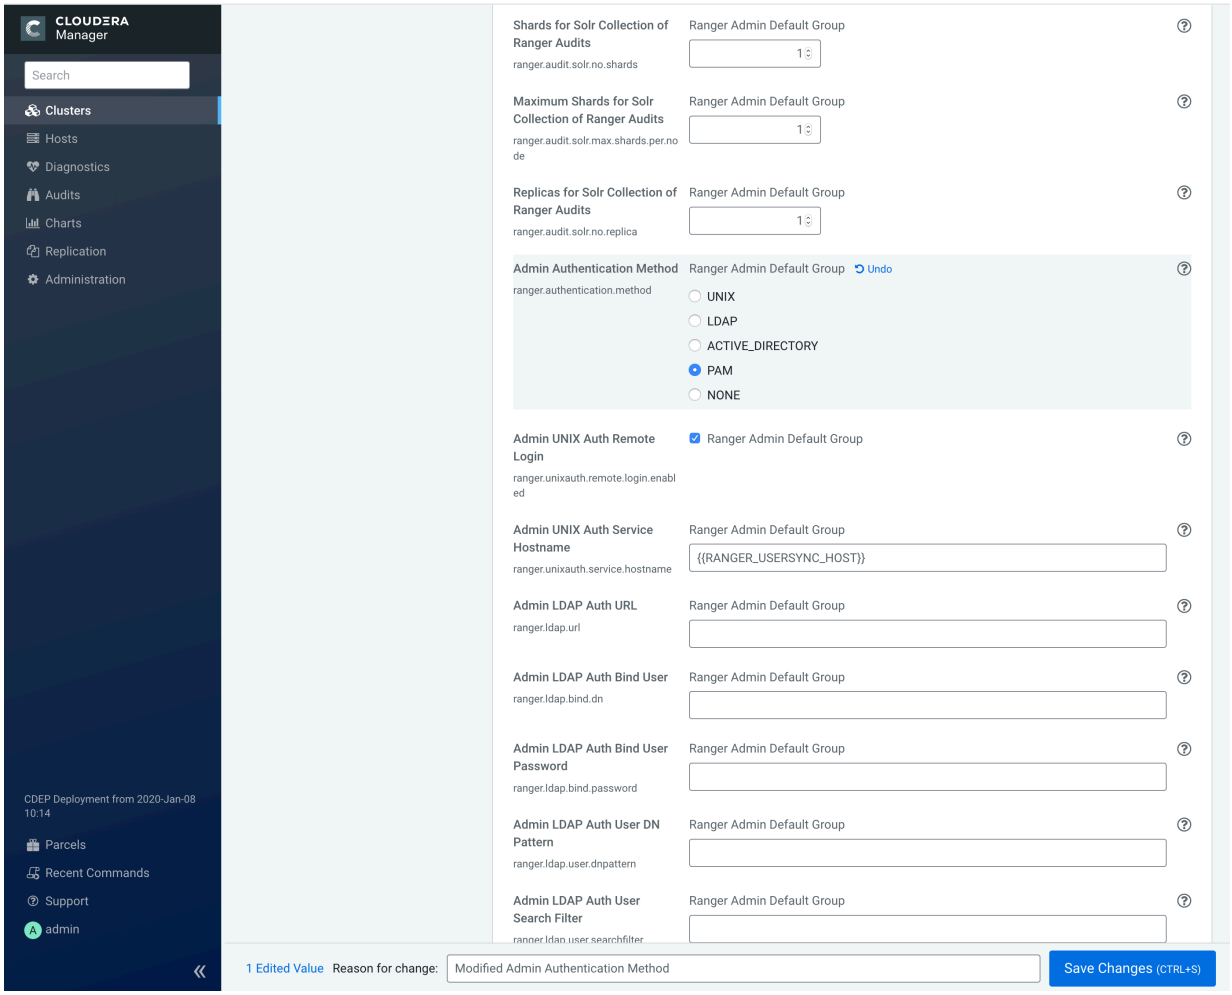

3. Create the following two PAM files:
   - /etc/pam.d/ranger-admin with the following content:

```
#%PAM-1.0
auth sufficient pam_unix.so
auth sufficient pam_sss.so
account sufficient pam_unix.so
account sufficient pam_sss.so
```

   - /etc/pam.d/ranger-remote with the following content:

```
#%PAM-1.0
auth sufficient pam_unix.so
auth sufficient pam_sss.so
account sufficient pam_unix.so
```

```
account sufficient pam_sss.so
```

4. Confirm that the /etc/shadow file has 444 permissions.
5. Select Actions > Restart to restart Ranger.

Ranger AD Integration

A conceptual overview of Ranger-AD integration architecture.

Ranger AD Integration: Architecture Overview

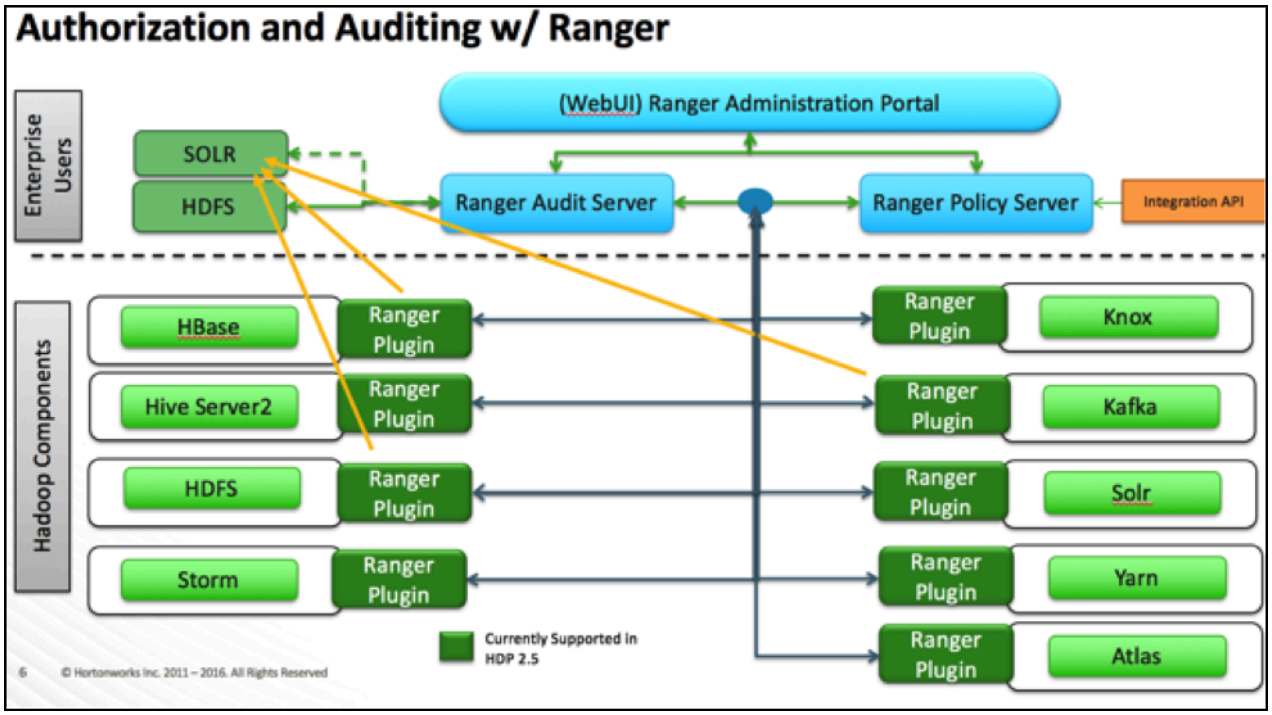

When a Ranger plugin for a component (such as HBase or HDFS) is activated, Ranger is in full control of any access. There is two-way communication between the Ranger plugin and the Ranger (Admin) Policy Server (RPS):

1. Plugins to RPS: Ranger plugins regularly call the RPS to see if new policies were defined in the Ranger Administration Portal (RAP). Generally it takes approximately 30 seconds for a policy to be updated.
2. RPS to components: The RPS queries the component for meta objects that live on the component to base policies upon (this provides the autocomplete and drop-down list when defining policies).

The first communication channel (Plugin to RPS) is essential for the plugin to function, whereas the second (RPS to components) is optional. It would still be possible to define and enforce policies without the second channel, but you would not have autocomplete during policy definition.

Configuration details on both communication channels are configured in both Cloudera Manager and in the Ranger Administration Portal.

Example for HDFS plugin on a kerberized cluster:

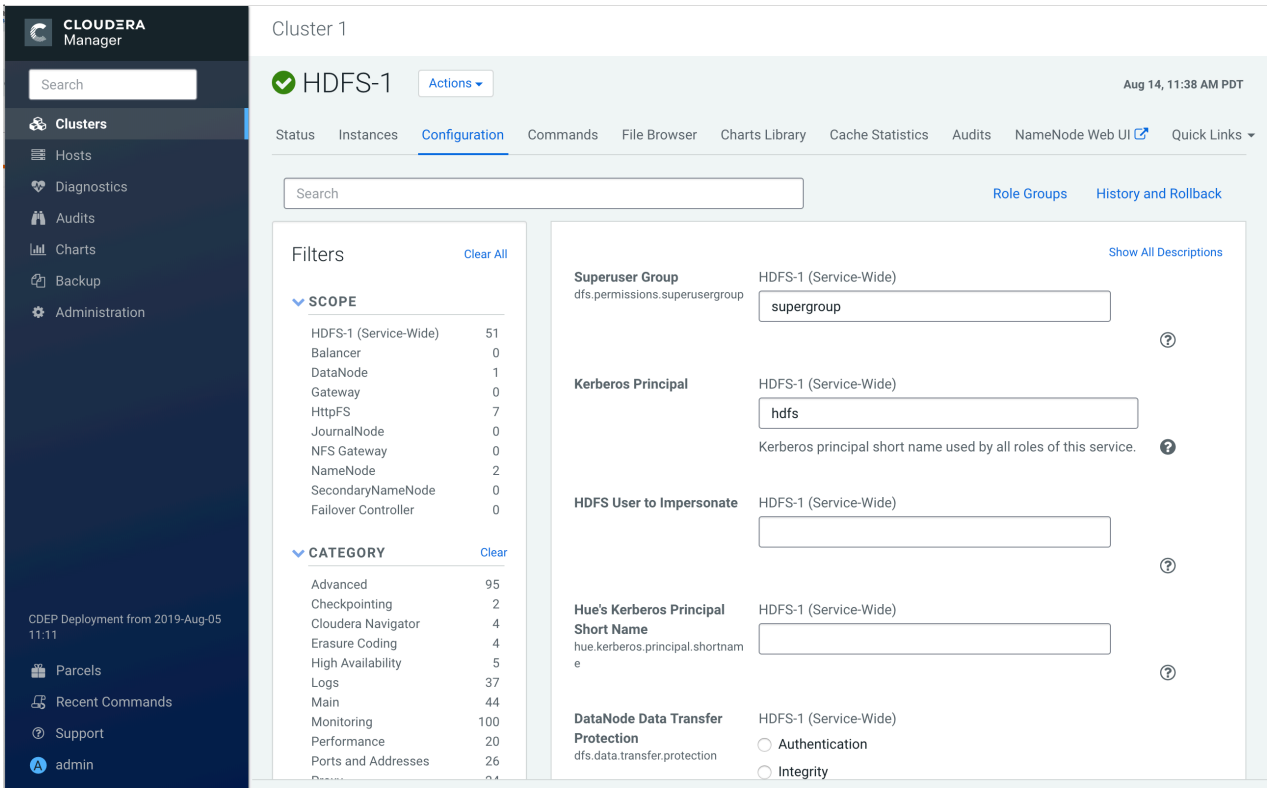

The Kerberos principal short name for the HDFS service,"hdfs", is the one that is involved the second communication channel (RPS to components) for getting metadata from HDFS (such as HDFS folders) across. The settings on the HDFS configuration must match those set in Ranger (by selecting Access > Manager > Resource Based Policies, then selecting the Edit icon for the HDFS service:

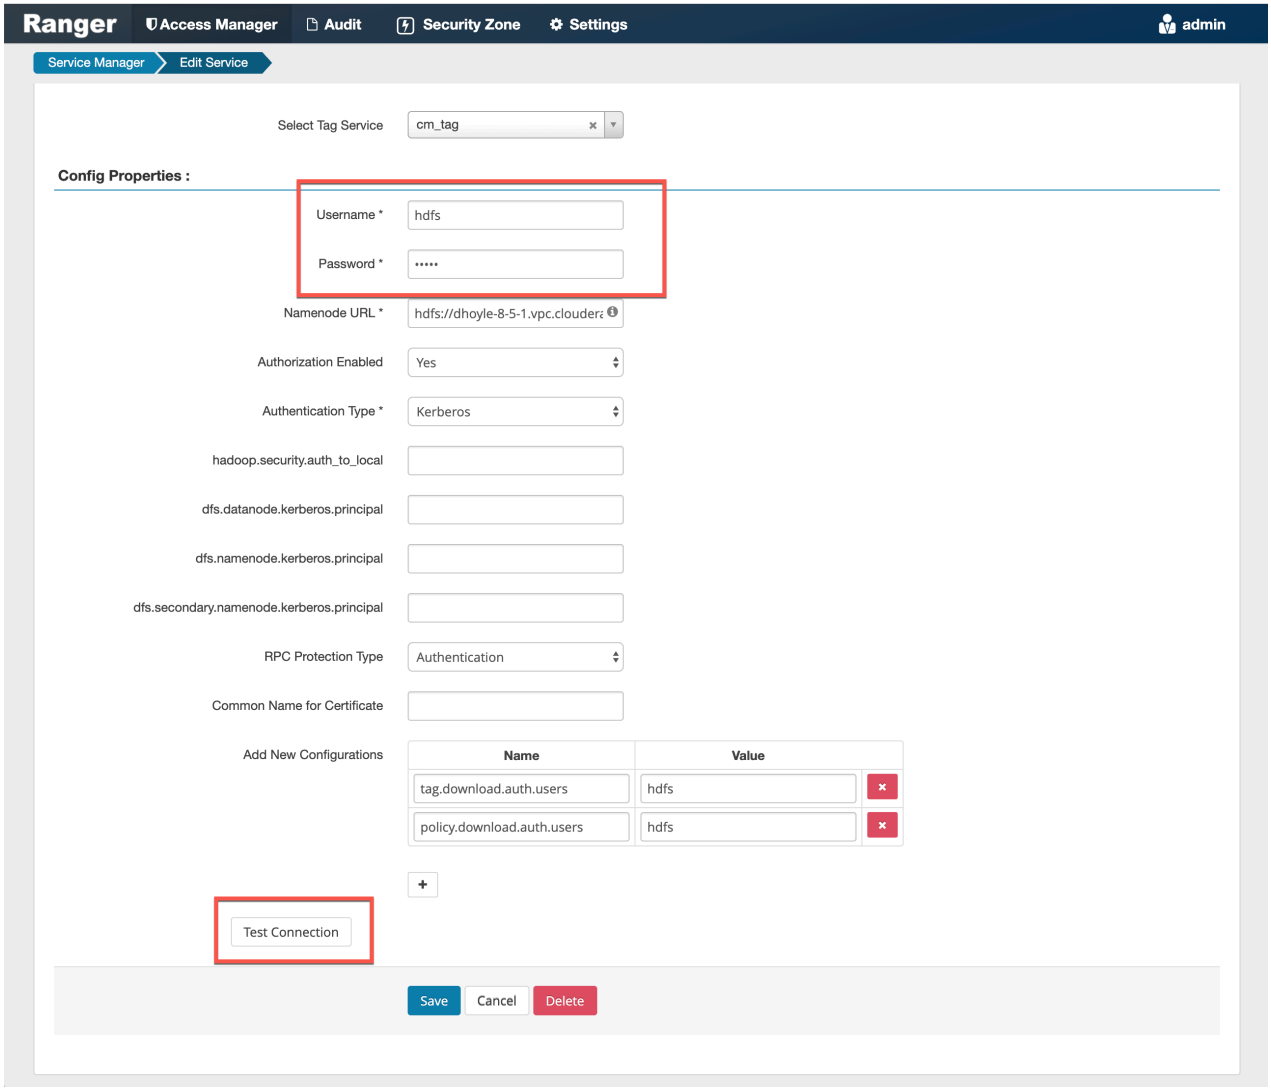

To verify the second communication channel (RPS to components) click Test Connection for the applicable service (as shown above for the HDFS service). A confirmation message appears if the connection works successfully.

To verify if the paramount first communication channel (Plugins to RPS) works, select Audit > Plugins in Ranger:

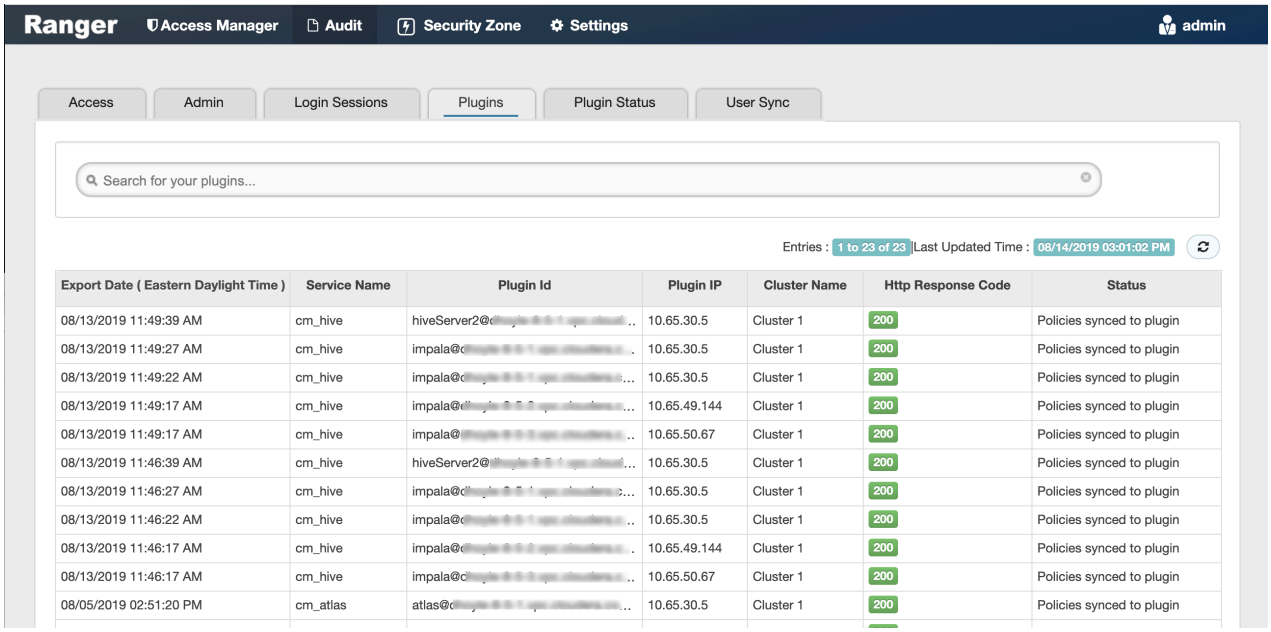

Ranger AD Integration: Ranger Audit

Ranger plugins furthermore send their audit event (whether access was granted or not and based on which policy) directly to the configured sink for audits, which can be HDFS, Solr or both. This is indicated by the yellow arrows in the architectural graph.

The audit access tab on the RAP (Audit > Access) is only populated if Solr is used as the sink.

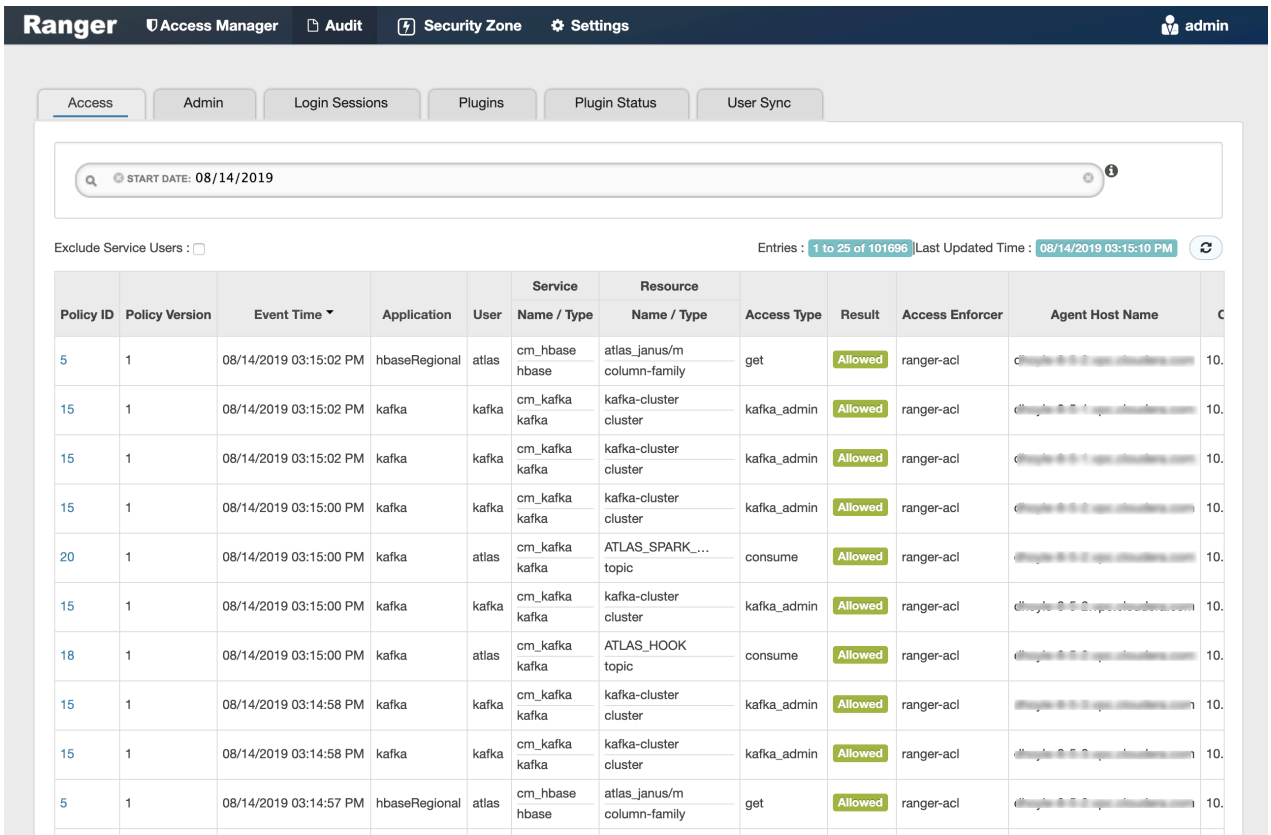

This screen points out an important Ranger feature. When the plugin is enabled AND no specific policy is in place for access to some object, the plugin will fall back to enforcing the standard component-level Access Control Lists (ACLs). For HDFS that would be the user : rwx / group : rwx / other : rwx ACLs on folders and files.

Once this defaulting to component ACLs happens, the audit events list a " - " in the Policy ID column instead of a policy number. If a Ranger policy was in control of allowing/denying access, the policy number is shown.

Ranger AD Integration: Overview

Rangers AD Integration has 2 levels:

1. Ranger UI authentication (which users can log in to Ranger itself).
2. Ranger user/group sync (which users/groups to define policies for)

Ranger UI authentication

Reference information on Ranger UI authentication, when configuring Ranger AD integration.

This is an extra AD level filter option on top of Kerberos authentication that maps to:

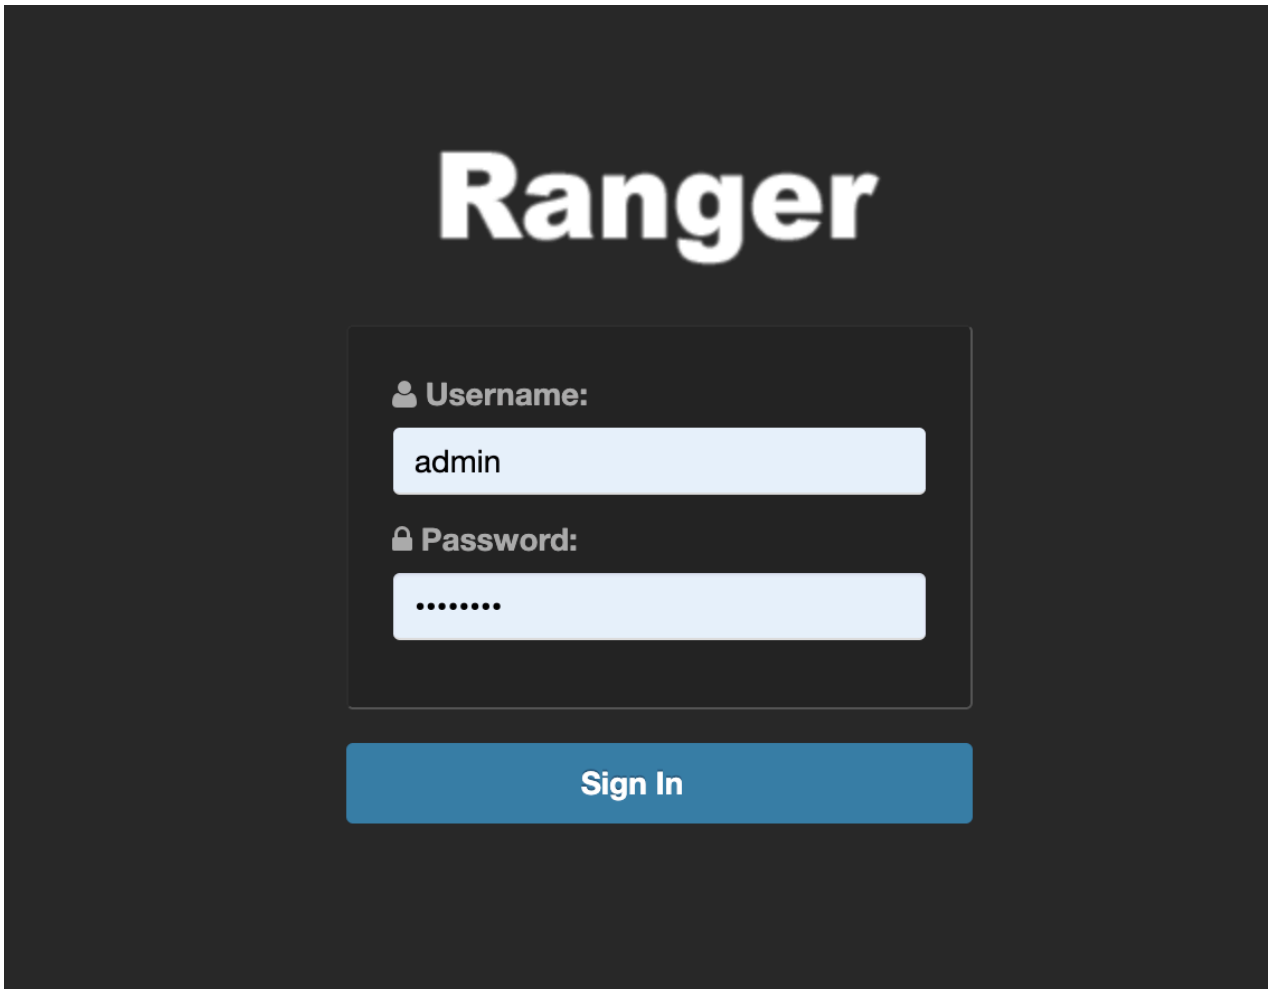

For AD there are two options for defining who can access the Ranger UI: LDAP or ACTIVE_DIRECTORY. There is not a huge amount of difference between them, but they are separate sets of properties.

ACTIVE_DIRECTORY

In Cloudera Manager, select Ranger, then click the Configuration tab.To display the authentication settings, type "authentication" in the Search box. You may need to scroll down to see the AD settings.

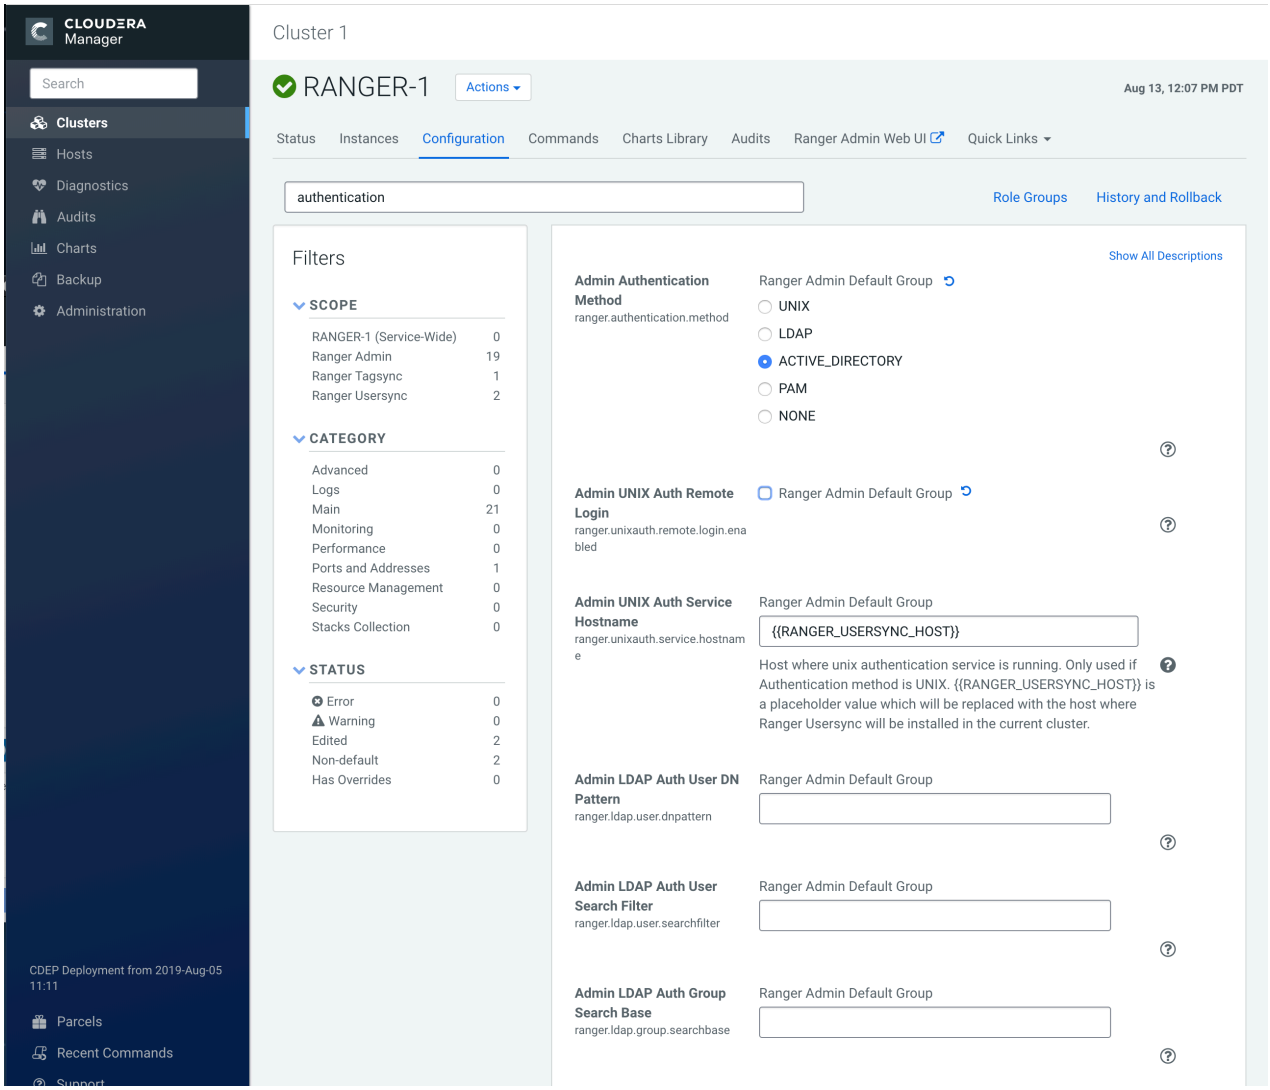

The ranger.ldap.ad.base.dn property determines the base of any search, so users not on this OU tree path can not be authenticated.

The ranger.ldap.ad.user.searchfilter poperty is a dynamic filter that maps the user name in the Ranger web UI login screen to sAMAccountName. For example, the AD sAMAccountName property has example values like k.reshi and d.alora so make sure to enter a matching value for ‘Username’ in the logon dialogue.

LDAP

The LDAP properties allow for more fine tuning.

In Cloudera Manager, select Ranger, then click the Configuration tab.To display the authentication settings, type "authentication" in the Search box. You may need to scroll down to see all of the LDAP settings.

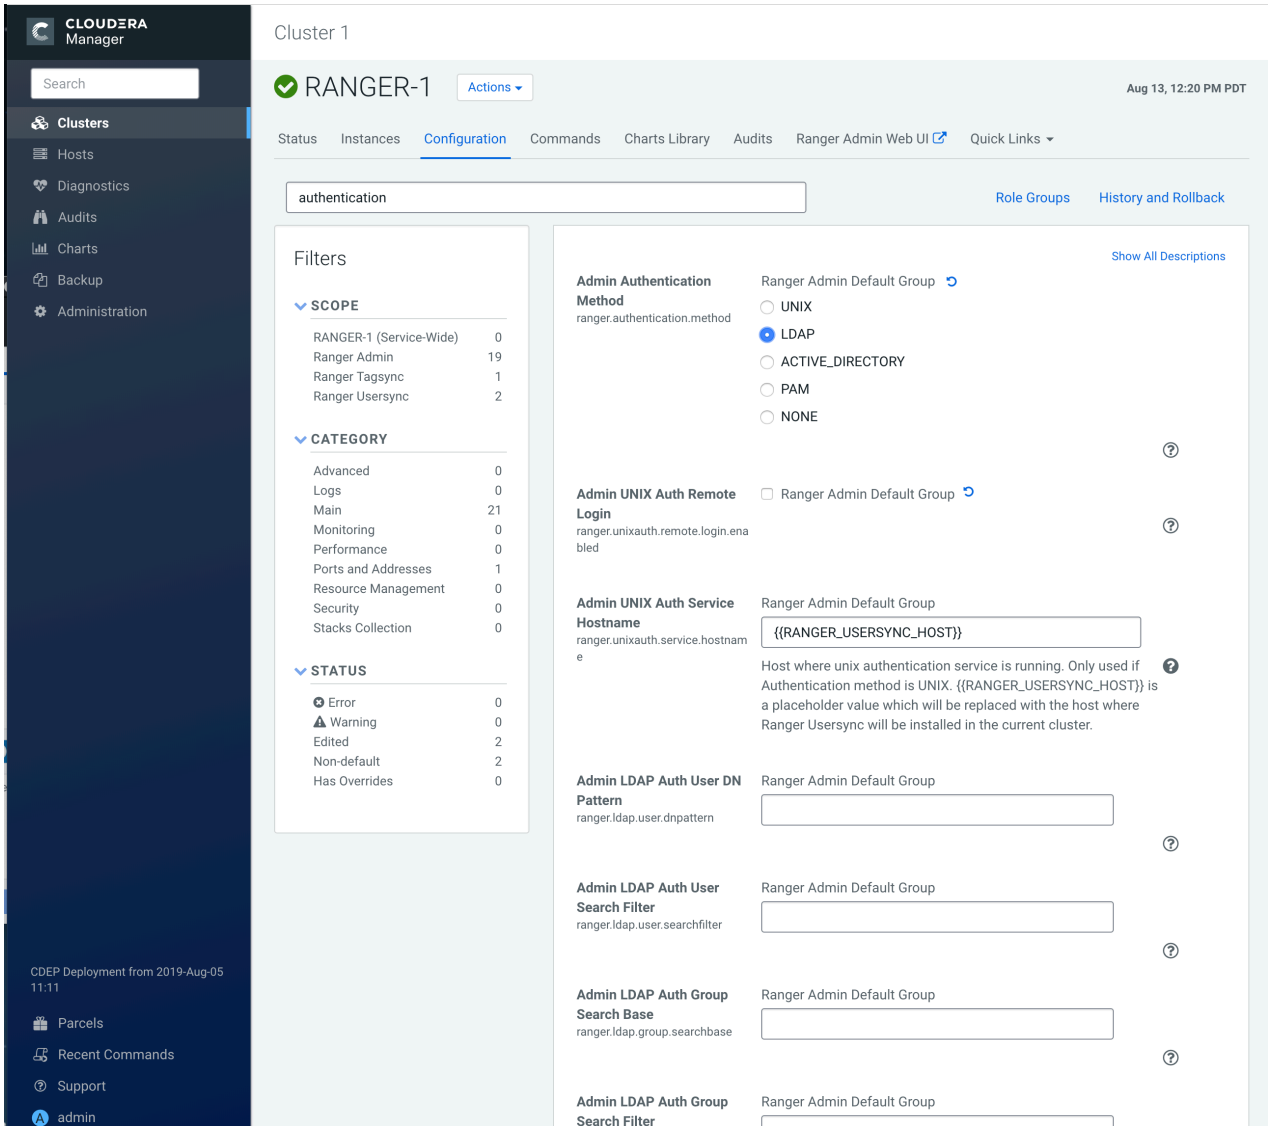

There is one catch: the ranger.ldap.user.dnpattern is evaluated first. Consider the following example value:

CN={0},OU=London,OU=Company,OU=User Accounts,OU=CorpUsers,DC=field,DC=hortonworks,DC=com

This would work, but has two side effects:

- Users would have to log on with their ‘long username’ (like ‘Kvothe Reshi / Denna Alora'), which would also mean that policies would have to be updated using that long name instead of the k.reshi short name variant.
- Traversing AD by DN patterns does not allow for applying group filters at all. In the syntax above, only users directly in OU=London would be able to log on.

This adverse behavior can be avoided by intentionally putting a DN pattern (DC=intentionally,DC=wrong) in the rang er.ldap.user.dnpattern property, AND a valid filter in User Search Filter:

(&(objectclass=user)(memberOf=CN=Hdp_admins,OU=Company,OU=User Accounts,OU=CorpUsers,DC=fiel d,DC=hortonworks,DC=com)(sAMAccountName={0}))

This works because the filter is only applied after the DN pattern query on AD does not return anything. If it does, the User Search Filter is not applied.

Ranger has a very simple approach to the internal user list that is kept in a relational schema. This list contains all users that were synced with AD ever, and all those users can potentially log in to the Ranger UI. But only Admin users can really do any policy-related things in the Ranger UI (see next section).

Be aware that all of this is only about authentication to Ranger. Someone from the ‘Hdp_admins’ group would still not have a Ranger admin role.

Related Information

Configure Ranger authentication for LDAP

Ranger UI authorization

Reference information on Ranger UI authorization, when configuring Ranger AD integration.

To configure the users, groups, and roles that can access the Ranger portal or its services, select Settings > Users/Groups/Roles in the top menu.

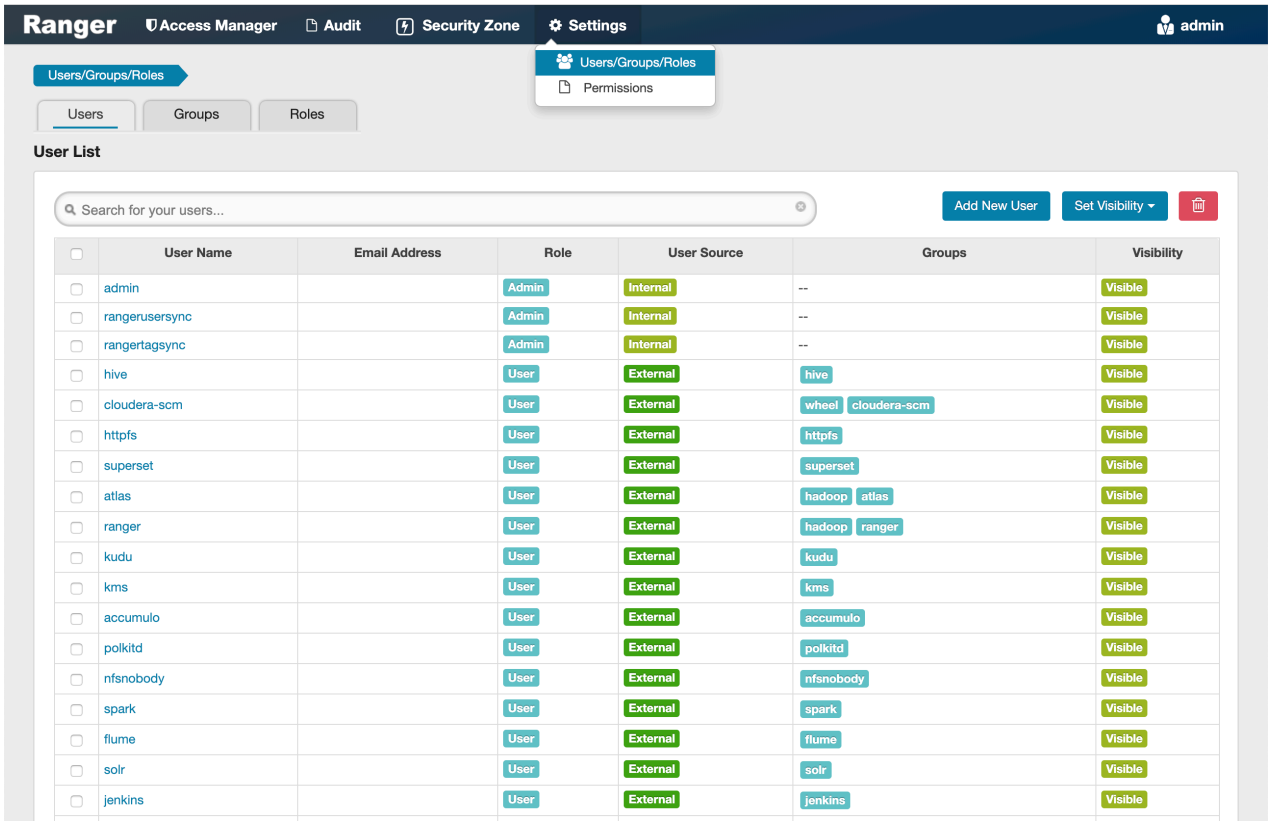

A user can be a User, Admin, or Auditor:

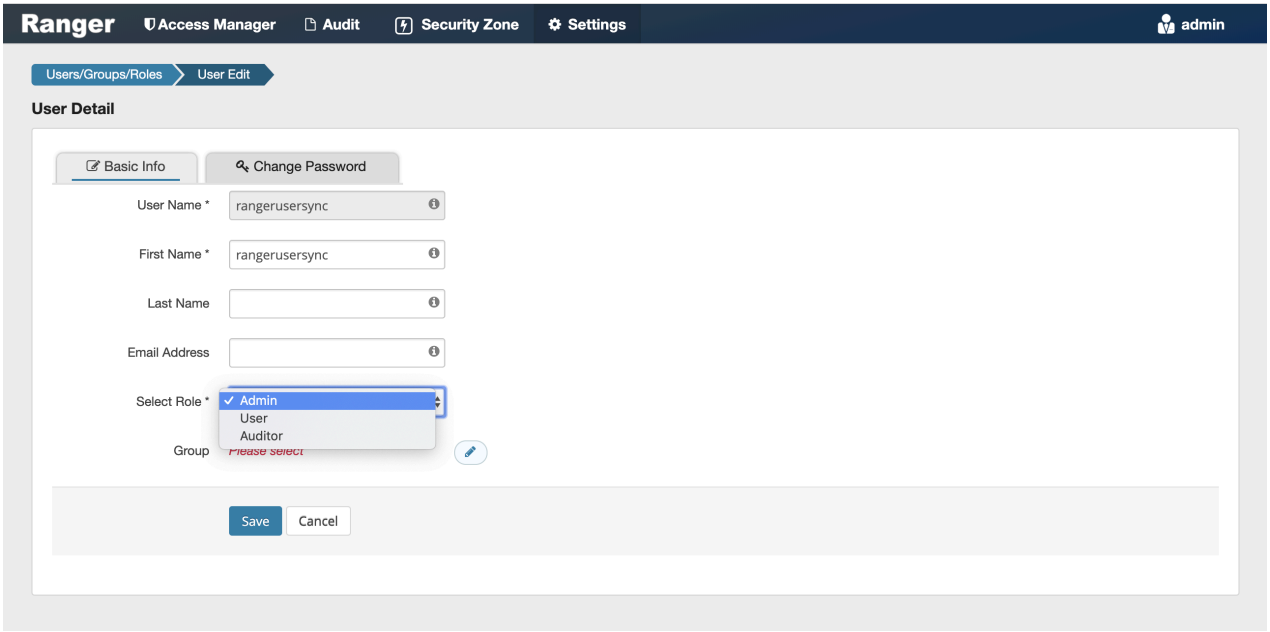

Only users with the Admin role can edit Ranger policies.

Configure Ranger Usersync for Deleted Users and Groups

How to configure Ranger Usersync for users and groups that have been deleted from the sync source.

About this task

You can configure Ranger Usersync to update Ranger when users and groups have been deleted from the sync source (UNIX, LDAP, AD or PAM). This ensures that users and groups – and their associated access permissions – do not remain in Ranger when they are deleted from sync source.

Procedure

1. In Cloudera Manager, select Ranger > Configuration, then use the Search box to search for Ranger Usersync Advanced Configuration Snippet (Safety Valve) for conf/ranger-ugsync-site.xml. Use the Add (+) icons to add the following properties, then click Save Changes.

Name	Value	Description
ranger.usersync.deletes.enabled	true	Enables deleted users and groups synchronization. The default setting is false (disabled).

Name	Value	Description
ranger.usersync.deletes.frequency	10	Sets the frequency of delete synchronization. The default setting is 10, or once every 10 Usersync cycles. Delete synchronization consumes cluster resources, so a lower (more frequent) setting may affect performance.

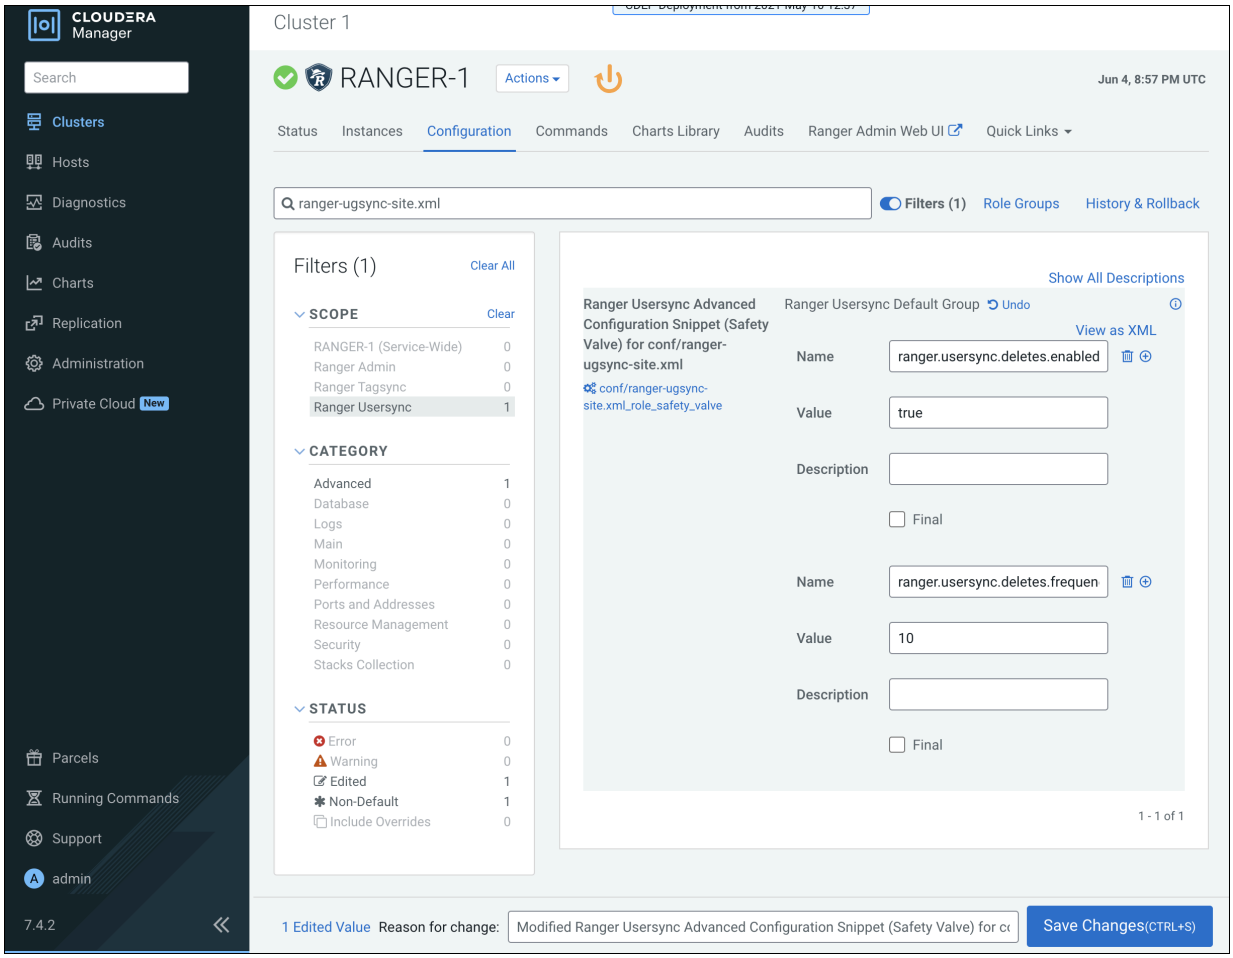

2. Click the Ranger Restart icon.

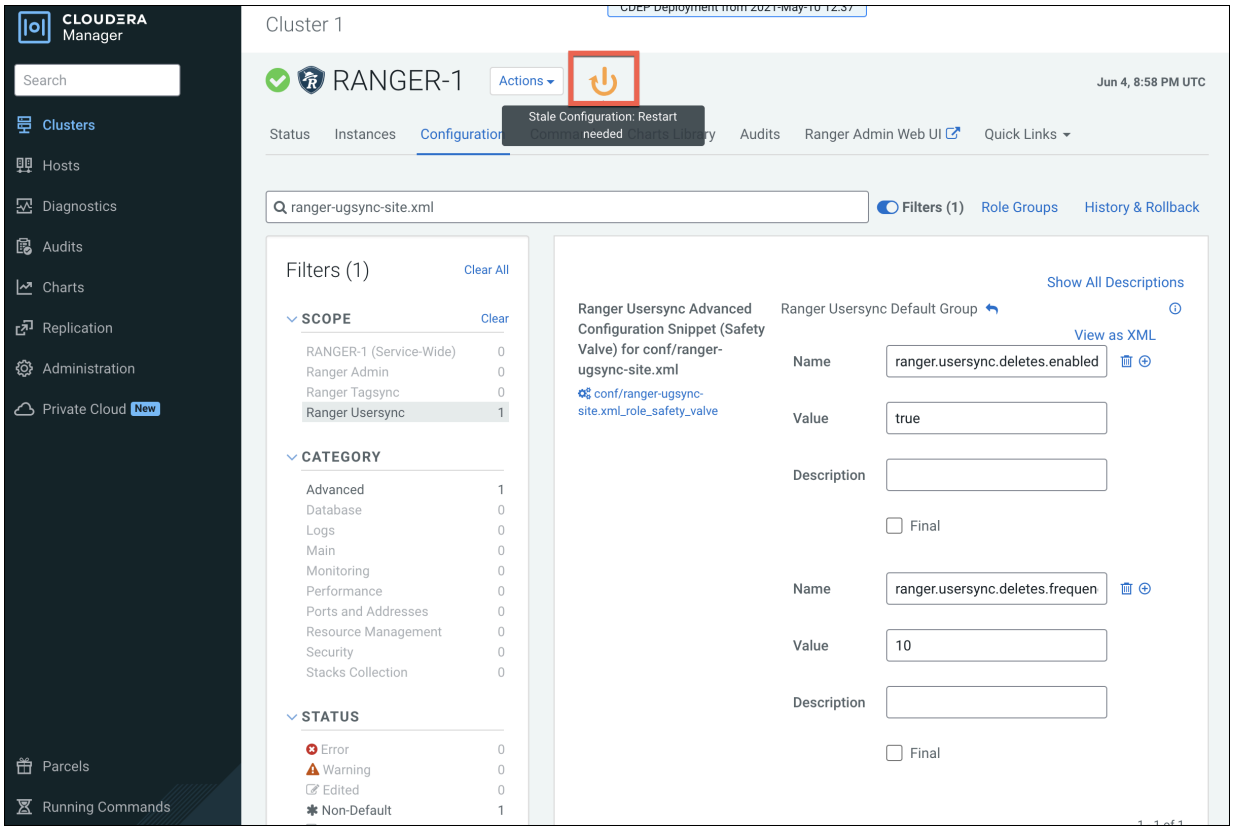

3. On the Stale Configurations page, click Restart Stale Services.

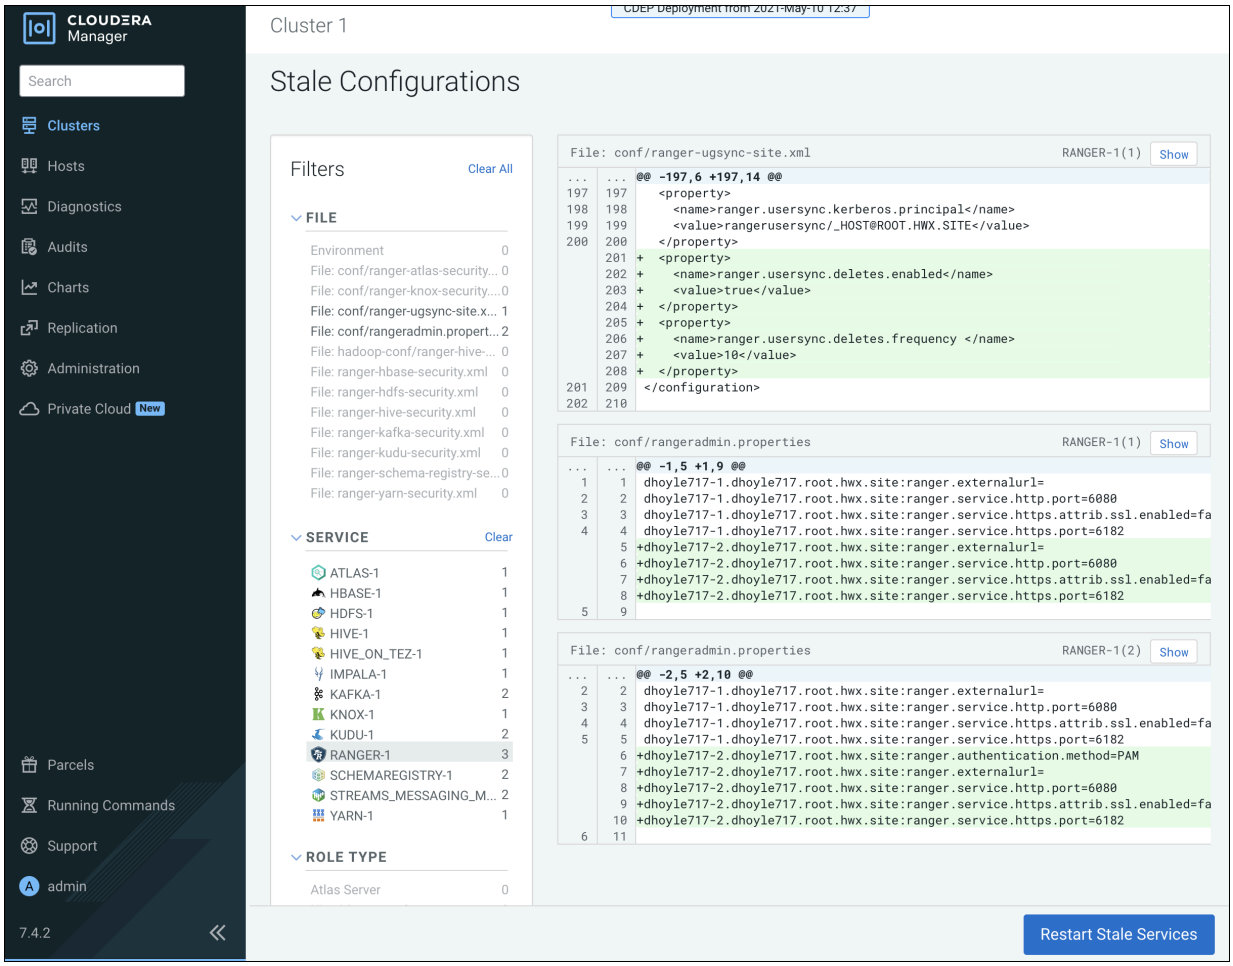

4. On the Restart Stale Services page, select the Re-deploy client configuration check box, then click Restart Now.

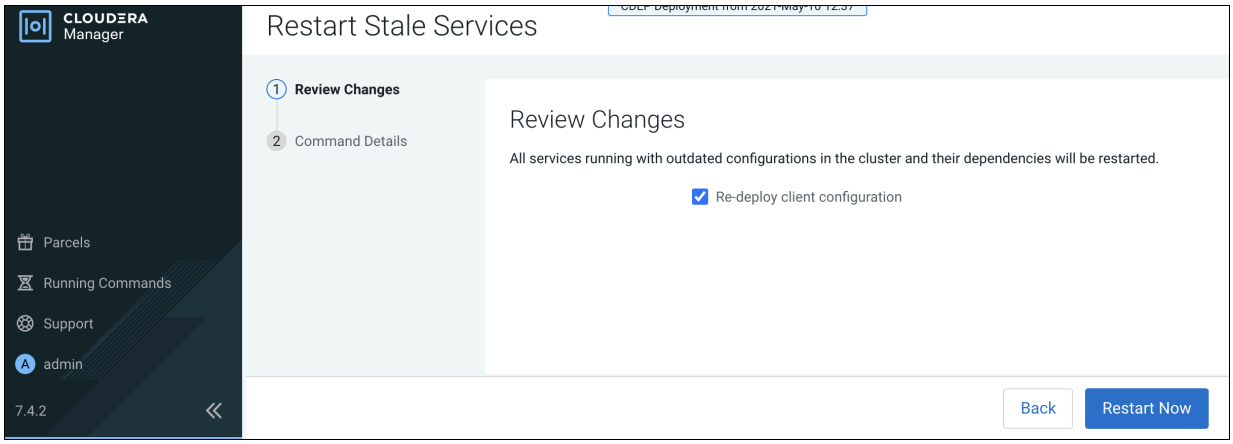

5. A progress indicator page appears while the services are being restarted. When the services have restarted, click Continue.

6. Users that have been deleted in sync source are not automatically deleted in Ranger – they are marked as Hidden and must be manually deleted by the Ranger Admin user, and then Ranger Usersync must be restarted.

   In the Ranger Admin Web UI, select Settings > Users/Groups/Roles. Click in the User List text box, then select Visibility > Hidden.

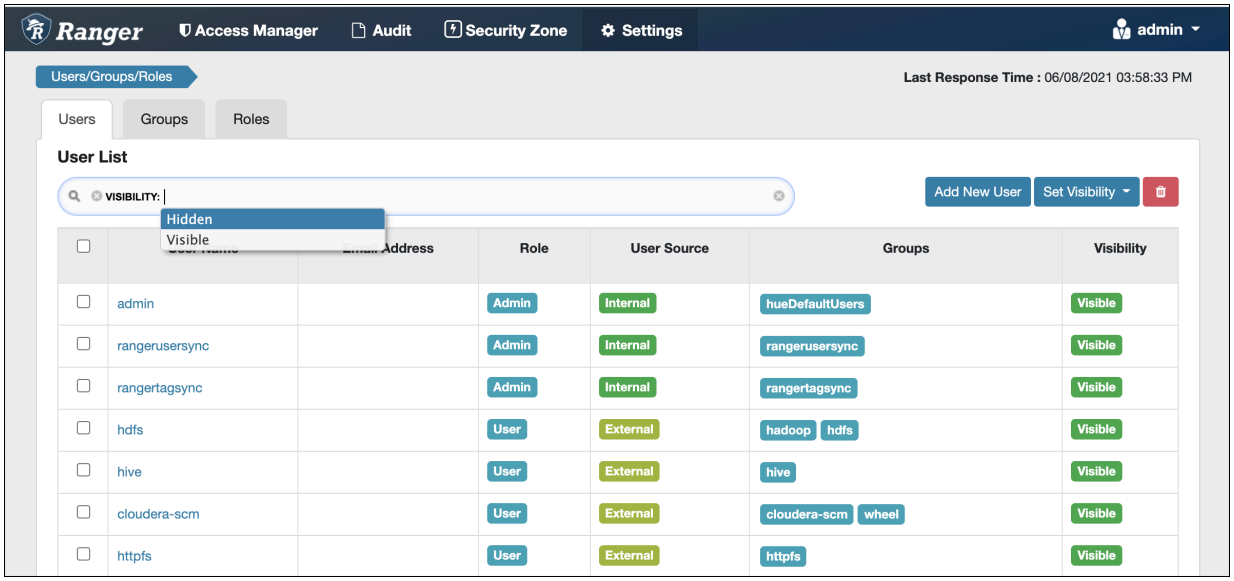

7. To delete a hidden user or group manually, select the applicable check boxes, then click the red Delete icon, as shown in the following example.

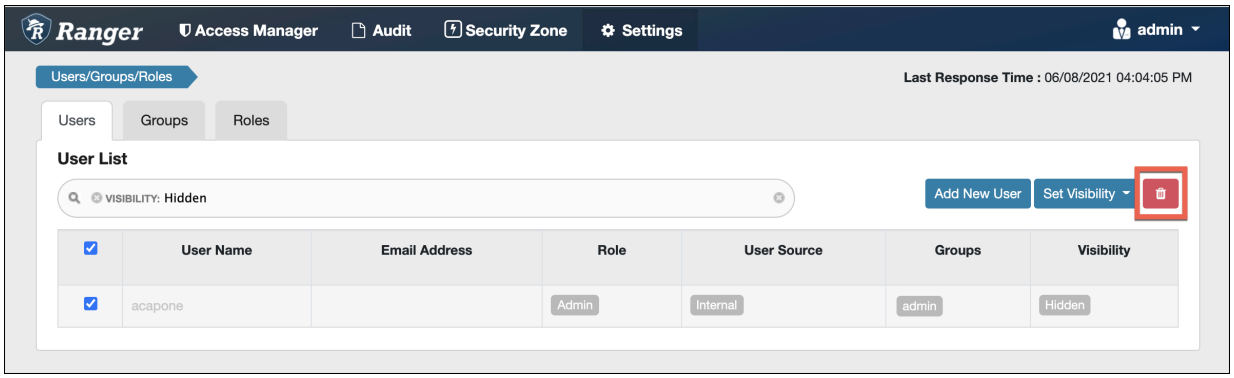

   You can delete multiple users or groups by running a "delete" script on the command line interface.

   For example:

```
Sample command to delete users:
python deleteUserGroupUtil.py -users <user file path> -admin <ranger admin
 user> -url <rangerhosturl> [-force] [-sslCertPath <cert path>] [-debug]

Sample command to delete groups:
python deleteUserGroupUtil.py -groups <group file path> -admin <ranger
admin user> -url <rangerhosturl> [-force] [-sslCertPath <cert path>] [-d
ebug]
```

Note: The deleteUserGroupUtil.py script installs as part of the Ranger installation on the node where Ranger Admin runs, in the following location: /opt/cloudera/parcels/CDH/lib/ranger-admin/ .

8. In Cloudera Manager, select Ranger > Ranger Usersync, then select Actions > Restart this Ranger Usersync.

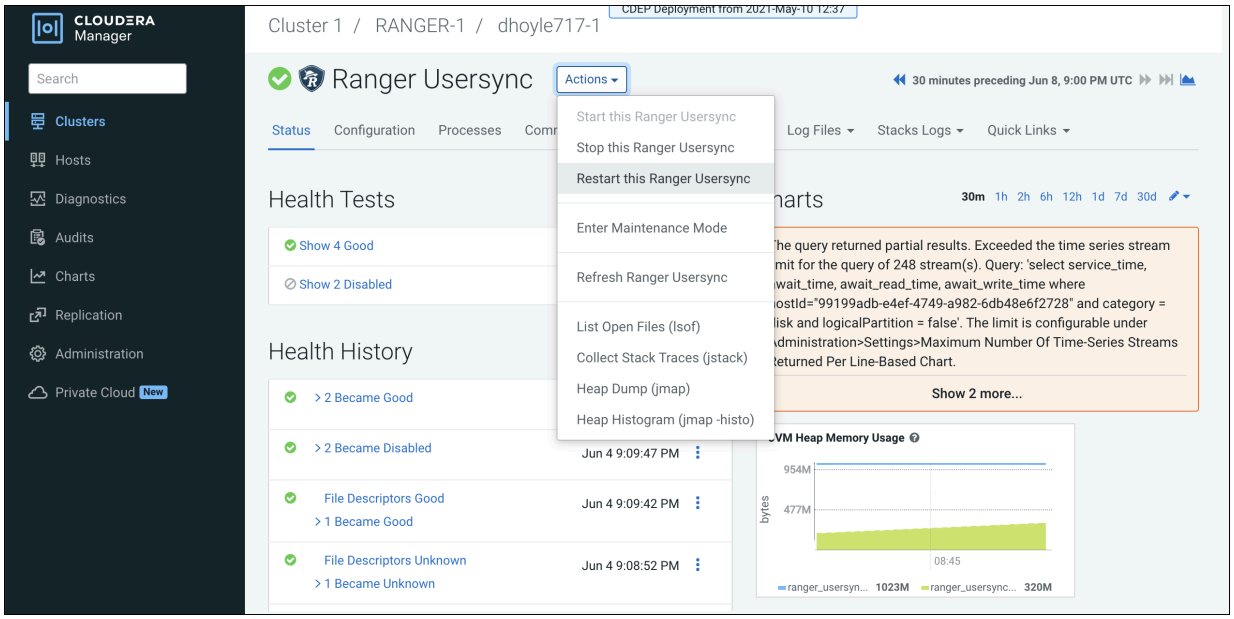

Note:

- Sync source is tracked when processing Ranger users and groups for deletion. If the same user name for a separate sync source already exists in Ranger DB, that user will not be updated or marked as hidden.
- For AD/LDAP sync:
  - After marking a user or group as deleted/hidden in Ranger, the user or group status does not change automatically. The user or group must be manually deleted (or deleted using the cli "delete" script). Usersync must be restarted to reflect any changes to the same user name in the source.
  - For example, a user (Bob) from one OU (say Engineering) is deleted from the source and is marked as deleted in Ranger admin. If the same user name (Bob) is subsequently added back to the same OU, the user status will not be automatically enabled. The user must be manually deleted and Usersync must be restarted to implement the changes.
  - If an identical user name (say Bob) is deleted from one OU (say Engineering) and added to a different OU (say Finance) between the sync cycles, user Bob is marked as hidden/deleted only when the delete cycle is triggered. Until then there is a security risk that user Bob from Finance will be granted the permissions for Bob from Engineering.