

Configuring Apache Ranger Authentication with UNIX, LDAP, or AD

Date of Publish: 2018-07-15

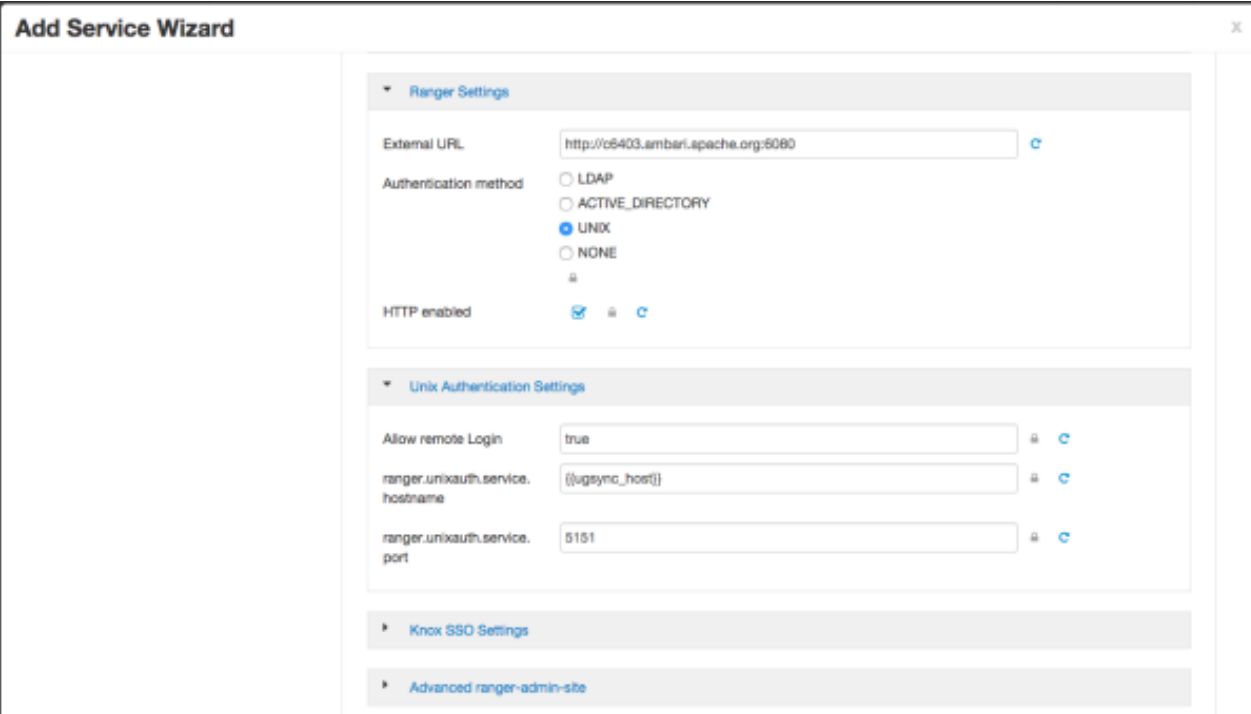


Contents

- Configuring Ranger Authentication with UNIX, LDAP, or AD.....3**
- Configure Ranger Authentication for UNIX..... 3**
- Configure Ranger Authentication for AD..... 4**
- Configure Ranger Authentication for LDAP..... 7**
- Ranger AD Integration..... 9**
 - Ranger UI Authentication.....13
 - Ranger UI Authorization..... 16
 - Ranger Usersync..... 17
 - Ranger User Management..... 23
 - Known Issue: Ranger Group Mapping..... 24

Configuring Ranger Authentication with UNIX, LDAP, or AD

This section describes how to configure the authentication method that determines who is allowed to login to the Ranger web interface. The options are local Unix, AD, or LDAP.



The screenshot shows the 'Add Service Wizard' window for configuring Ranger. It is divided into two main sections: 'Ranger Settings' and 'Unix Authentication Settings'. In the 'Ranger Settings' section, the 'External URL' is set to 'http://c6403.ambari.apache.org:6080'. The 'Authentication method' is set to 'UNIX' (selected with a radio button). The 'HTTP enabled' checkbox is checked. In the 'Unix Authentication Settings' section, 'Allow remote Login' is set to 'true'. The 'ranger.unikauth.service.hostname' is set to '({jgsync_host})'. The 'ranger.unikauth.service.port' is set to '5151'. There are also sections for 'Knox SSO Settings' and 'Advanced ranger-admin-site' which are currently collapsed.

Configure Ranger Authentication for UNIX

How to configure Ranger to use Unix for user authentication.

About this task

You can configure Ranger authentication in two ways:

- During installation: **Ranger Customize Services > Advanced tab > Ranger Settings**
- After installation: **Ambari > Ranger > Configs > Advanced > Ranger Settings**

The screenshot shows the 'Add Service Wizard' window. The 'Ranger Settings' tab is active, displaying the following configuration options:

- External URL:** A text field containing 'http://o5403.ambari.apache.org:6080'.
- Authentication method:** Radio buttons for LDAP, ACTIVE_DIRECTORY, UNIX (selected), and NONE.
- HTTP enabled:** A checkbox that is checked.

The 'Unix Authentication Settings' tab is also visible, showing the following configuration options:

- Allow remote Login:** A text field containing 'true'.
- ranger.unixauth.service.hostname:** A text field containing '([logsync_host])'.
- ranger.unixauth.service.port:** A text field containing '5151'.

Below these tabs are links for 'Knox SSO Settings' and 'Advanced ranger-admin-site'.

Procedure

- From the **Ranger Settings** tab:
 - Enter the external URL, e.g. `http://my-vm.hortonworks.com:6080`.
 - Under **Authentication method**, select **UNIX**.
 - Under **HTTP enabled**, make a selection. This option enables you to select HTTP/HTTPS communication for Ranger admin console. If you disable HTTP, only HTTPS is allowed. HTTP is enabled by default.
- From the **UNIX Authentication Settings** tab, enter the following values:

Table 1: UNIX Authentication Settings

Configuration Property	Description	Default Value	Example Value	Required
Allow remote Login	Flag to enable/disable remote login via UNIX Authentication Mode.	TRUE	TRUE	No.
ranger.unixauth.service.hostname	The FQDN where the ranger-usersync module is running (along with the UNIX Authentication Service).	localhost	myunixhost.domain.com	Yes, if select
ranger.unixauth.service.port	The port number where the ranger-usersync module is running the UNIX Authentication Service.	5151	5151	Yes, if select

Configure Ranger Authentication for AD

How to configure Ranger to use AD for user authentication.

About this task

You can configure Ranger authentication in two ways:

- During installation: **Ranger Customize Services > Advanced tab > Ranger Settings**
- After installation: **Ambari > Ranger > Configs > Advanced > Ranger Settings**

The screenshot shows the 'Add Service Wizard' interface. It has two main sections: 'Ranger Settings' and 'AD Settings'.

Ranger Settings:

- External URL:** A text field containing 'http://06403.ambari.apache.org:6080'.
- Authentication method:** Radio buttons for LDAP, **ACTIVE_DIRECTORY** (selected), UNIX, and NONE.
- HTTP enabled:** A checkbox that is checked.

AD Settings:

- ranger.ldap.ad.base.dn:** A text field containing 'dc=example,dc=com'.
- ranger.ldap.ad.bind.dn:** A text field containing '{{ranger_ug_ldap_bind_dn}}'.
- ranger.ldap.ad.bind.password:** A password field with a masked input.
- Domain Name (Only for AD):** A text field containing 'dc=hwqa,dc=hortonworks,dc=com'.
- ranger.ldap.ad.referral:** A text field containing 'ignore'.
- ranger.ldap.ad.uri:** A text field containing '{{ranger_ug_ldap_uri}}'.
- ranger.ldap.ad.user.searchfilter:** A text field containing '{{ranger_ug_ldap_user_searchfilter}}'.

Below the AD Settings section, there are two expandable sections: 'Knox SSO Settings' and 'Advanced ranger-admin-site'.

Procedure

1. From the **Ranger Settings** tab:

- Enter the external URL, e.g. `http://my-vm.hortonworks.com:6080`.
- Under **Authentication method**, select **ACTIVE_DIRECTORY**.
- Under **HTTP enabled**, make a selection. This option enables you to select HTTP/HTTPS communication for Ranger admin console. If you disable HTTP, only HTTPS is allowed. HTTP is enabled by default.

2. From the **AD Settings** tab, enter the following values:

Property	Description	Default value	Sample values
ranger.ldap.ad.base.dn	The Distinguished Name (DN) of the starting point for directory server searches.	dc=example,dc=com	dc=example,dc=com
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. This is a macro variable value that is derived from the Bind User value from Ranger User Info > Common Configs.	{{ranger_ug_ldap_bind_dn}}	{{ranger_ug_ldap_bind_dn}}
ranger.ldap.ad.bind.password	Password for the bind.dn. This is a macro variable value that is derived from the Bind User Password value from Ranger User Info > Common Configs.		
Domain Name (Only for AD)	The domain name of the AD Authentication service.		dc=example,dc=com

Property	Description	Default value	Sample values
ranger.ldap.ad.referral*	See below.	ignore	follow ignore throw
ranger.ldap.ad.url	The AD server URL. This is a macro variable value that is derived from the LDAP/AD URL value from Ranger User Info > Common Configs.	{{ranger_ug_ldap_url }}	{{ranger_ug_ldap_url }}
ranger.ldap.ad.user.searchfilter	The search filter used for Bind Authentication. This is a macro variable value that is derived from the User Search Filter value from Ranger User Info > User Configs.	{{ranger_ug_ldap_user_searchfilter}}	{{ranger_ug_ldap_user_searchfilter}}

3. Optional: Custom ranger-admin-site Settings for Active Directory:

a) Select Custom ranger-admin-site, then click Add Property.

The screenshot shows the Ranger configuration interface. The 'AD Settings' section is expanded, showing 'ranger.ldap.ad.domain' set to 'localhost' and 'ranger.ldap.ad.url' set to 'ldap://ad.xasecure.net:389'. Below this, the 'LDAP Settings' section is collapsed. The 'Advanced ranger-admin-site' section is collapsed. The 'Advanced ranger-env' section is collapsed. The 'Advanced ranger-ugsync-site' section is collapsed. The 'Custom admin-properties' section is collapsed. The 'Custom ranger-admin-site' section is expanded, and the 'Add Property ...' button is highlighted with a red box. Below this, the 'Custom ranger-site' section is collapsed. The 'Custom ranger-ugsync-site' section is collapsed. The 'Custom usersync-properties' section is collapsed.

b) The following table shows the Custom ranger-admin-site settings required for Active Directory (AD) authentication:

Key	Value
ranger.ldap.ad.base.dn	dc=example,dc=com
ranger.ldap.ad.bind.dn	cn=adadmin,cn=Users,dc=example,dc=com
ranger.ldap.ad.bind.password	Secret123!

Key	Value
ranger ldap.ad.referral*	follow ignore throw

Custom ranger-site

ranger ldap.ad.base.dn

dc=example,dc=com

+

-

ranger ldap.ad.bind.dn

cn=adadmin,cn=Users,dc=example,dc=com

+

-

ranger ldap.ad.bind.password

secret123!

+

-

ranger ldap.ad.referral

follow

+

-

Add Property ...

*

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 theReferralException 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.

Configure Ranger Authentication for LDAP

How to configure Ranger to use LDAP for user authentication.

About this task

You can configure Ranger authentication in two ways:

- During installation: **Ranger Customize Services > Advanced tab > Ranger Settings**
- After installation: **Ambari > Ranger > Configs > Advanced > Ranger Settings**

The screenshot shows the 'Add Service Wizard' interface. It has two main sections: 'Ranger Settings' and 'LDAP Settings'. Below these are several expandable sections: 'Knox SSO Settings', 'Advanced ranger-admin-site', 'Advanced ranger-env', 'Advanced ranger-ugsync-site', and 'Custom admin-properties'.

Ranger Settings

- External URL:
- Authentication method: ☒ LDAP, ☐ ACTIVE_DIRECTORY, ☐ UNIX, ☐ NONE
- HTTP enabled: ☒

LDAP Settings

- ranger.ldap.base.dn:
- Bind User:
- Bind User Password:
- ranger.ldap.group.roleattribute:
- ranger.ldap.referral:
- LDAP URL:
- ranger.ldap.user.dnpattern:
- User Search Filter:

Procedure

- From the **Ranger Settings** tab:
 - Enter the external URL, e.g. `http://my-vm.hortonworks.com:6080`.
 - Under **Authentication method**, select **LDAP**.
 - Under HTTP enabled, make a selection. This option enables you to select HTTP/HTTPS communication for Ranger admin console. If you disable HTTP, only HTTPS is allowed. HTTP is enabled by default.
- From the **LDAP Settings** tab, enter the following values:

Property	Description	Default value	Sample values
Group Search Base		<code>{{ranger_ug_ldap_group_searchbase}}</code>	
Group Search Filter		<code>{{ranger_ug_ldap_group_searchfilter}}</code>	
LDAP URL		<code>{{ranger_ug_ldap_url}}</code>	
Bind User		<code>{{ranger_ug_ldap_bind_dn}}</code>	
Bind User Password		N/A	
User Search Filter		<code>(uid={0})</code>	
ranger.ldap.base.dn		<code>dc=example,dc=com</code>	

Property	Description	Default value	Sample values
ranger.ldap.group.roleattribute		cn	
ranger.ldap.referral	See below.	ignore	follow throw ignore
ranger.ldap.user.dnpattern		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.

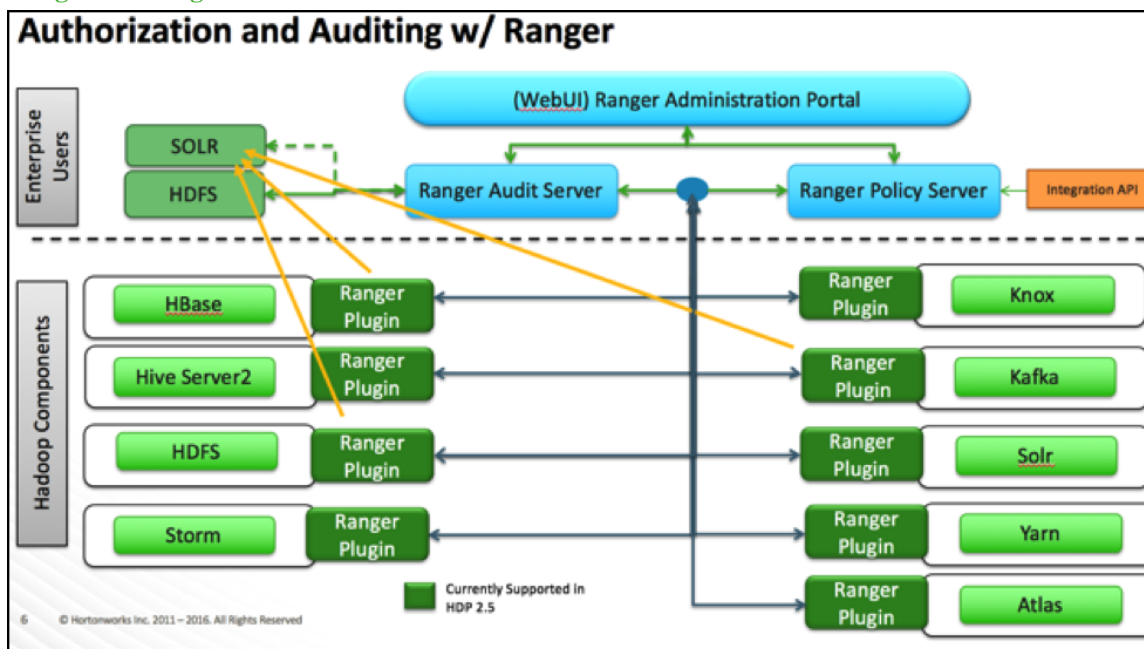
When searching a directory, the server might return several search results, along with a few continuation references that show where to obtain further results. These results and references might be interleaved at the protocol level.

- 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.

Ranger AD Integration

A conceptual overview of Ranger-AD integration architecture.

Ranger AD Integration: Architecture Overview



When a Ranger plugin for a component (like HBase or HDFS) is activated, Ranger will be in full control of any access. There is a two-way communication between the Ranger plugin and 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 allow for 30 sec. 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 dropdown list when defining policies.)

The first communication channel (Plugins 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 if the second does not work, but you will not have autocomplete during policy definition.

Configuration details on both communication channels are configured on both Ambari configuration for the component and on the RAP.

Example for HDFS plugin:

Advanced ranger-hdfs-plugin-properties

Enable Ranger for HDFS

Ranger repository config password

Ranger repository config user

common.name.for.certificate

hadoop.rpc.protection

Policy user for HDFS

Advanced ranger-hdfs-policymgr-ssl

xasecure.policymgr.clientssl.keystore

xasecure.policymgr.clientssl.keystore.credential.file

xasecure.policymgr.clientssl.keystore.password

xasecure.policymgr.clientssl.truststore

xasecure.policymgr.clientssl.truststore.credential.file

xasecure.policymgr.clientssl.truststore.password

The ‘Ranger repository config user’ is the one that involved the second communication channel (RPS to components) for getting metadata from HDFS (like HDFS folders) across. The settings on the HDFS configuration have to match those set at the Ranger end (Access Manager > Resource Based Policies > HDFS >



:

Ranger Access Manager Audit Settings

Service Manager Edit Service

Edit Service

Service Details :

Service Name * HDP_hadoop

Description hdfs repo

Active Status ☒ Enabled ☐ Disabled

Select Tag Service Select Tag Service

Config Properties :

Username * hadoop

Password * *****

Namenode URL * hdfs://hdp25-m-01:8020

Authorization Enabled Yes

Authentication Type * Kerberos

To verify if the paramount first communication channel (Plugins to RPS) works can be done by having a look in the RAP at Audit > Plugins:

Ranger Access Manager Audit Settings admin

Access Admin Login Sessions Plugins

Search for your plugins...

Last Updated Time : 12/14/2016 10:23:58 AM

Export Date (CET) *	Service Name	Plugin Id	Plugin IP	Http Response Code	Status
12/13/2016 01:13:30 PM	HDP_hive	hiveServer2@hdp25-m-02-HDP_hive	172.26.	200	Policies synced to plugin
12/13/2016 01:12:00 PM	HDP_hive	hiveServer2@hdp25-m-02-HDP_hive	172.26.	200	Policies synced to plugin
12/13/2016 11:09:15 AM	HDP_atlas	atlas@hdp25-m-01-HDP_atlas	172.26.	200	Policies synced to plugin
12/13/2016 11:00:14 AM	HDP_atlas	atlas@hdp25-m-01-HDP_atlas	172.26.	200	Policies synced to plugin
12/13/2016 10:43:12 AM	HDP_atlas	atlas@hdp25-m-01-HDP_atlas	172.26.	200	Policies synced to plugin
12/12/2016 10:58:18 PM	HDP_kafka	kafka@hdp25-s-03-HDP_kafka	172.26.	200	Policies synced to plugin

To verify the second communication channel (RPS to components) press the 'Test Connection' button (Access Manager > Resource Based Policies > HDFS >



:

Authorization Enabled	Yes												
Authentication Type *	Kerberos												
hadoop.security.auth_to_local	RULE:[1:\$1@\$0](ambari-qa-hdp_rj)												
dfs.datanode.kerberos.principal	dn/-hdp25-m-01@FIELD.HORTO												
dfs.namenode.kerberos.principal	nn/-hdp25-m-01@FIELD.HORTO												
dfs.secondary.namenode.kerberos.principal	nn/-hdp25-m-01@FIELD.HORTO												
RPC Protection Type	Authentication												
Common Name for Certificate													
Add New Configurations	<table border="1"> <thead> <tr> <th>Name</th> <th>Value</th> <th></th> </tr> </thead> <tbody> <tr> <td>ambari.service.check.user</td> <td>ambari-qa</td> <td>×</td> </tr> <tr> <td>tag.download.auth.users</td> <td>hdfs</td> <td>×</td> </tr> <tr> <td>policy.download.auth.users</td> <td>hdfs</td> <td>×</td> </tr> </tbody> </table>	Name	Value		ambari.service.check.user	ambari-qa	×	tag.download.auth.users	hdfs	×	policy.download.auth.users	hdfs	×
Name	Value												
ambari.service.check.user	ambari-qa	×											
tag.download.auth.users	hdfs	×											
policy.download.auth.users	hdfs	×											
	+												
<div>Test Connection</div>													
<div>Save Cancel Delete</div>													

If the settings are right you'll get:



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 sink.

Policy ID	Event Time *	User	Service Name / Type	Resource Name / Type	Access Type	Result	Access Enforcer	Client IP	Event Count	Tags
-	12/14/2016 11:08:31 AM	spark	HDP_RK_hadoop hdfs	/spark2-history/ path	READ_EXECUTE	Allowed	hadoop-acl	172.26.1.1	1	
-	12/14/2016 11:08:31 AM	spark	HDP_RK_hadoop hdfs	/spark2-history/ path	WRITE	Allowed	hadoop-acl	172.26.1.1	1	
-	12/14/2016 11:08:31 AM	spark	HDP_RK_hadoop hdfs	/spark2-history/ path	READ_EXECUTE	Allowed	hadoop-acl	172.26.1.1	1	
-	12/14/2016 11:08:31 AM	spark	HDP_RK_hadoop hdfs	/spark2-history/ path	WRITE	Allowed	hadoop-acl	172.26.1.1	1	
-	12/14/2016 11:08:31 AM	spark	HDP_RK_hadoop hdfs	/spark2-history/ path	WRITE	Allowed	hadoop-acl	172.26.1.1	1	
-	12/14/2016 11:08:31 AM	spark	HDP_RK_hadoop hdfs	/spark2-history/ path	WRITE	Allowed	hadoop-acl	172.26.1.1	1	
-	12/14/2016 11:08:26 AM	rangeragync	HDP_RK_kafka kafka	ATLAS_ENTITIES topic	describe	Denied	ranger-acl	172.26.1.1	11	
-	12/14/2016 11:08:26 AM	rangeragync	HDP_RK_kafka kafka	ATLAS_ENTITIES topic	describe	Denied	ranger-acl	172.26.1.1	7	
14	12/14/2016 11:08:25 AM	atlas	HDP_RK_hbase hbase	atlas_star/m column-family	get	Allowed	ranger-acl	172.26.1.1	1	
8	12/14/2016 11:08:25 AM	atlas	HDP_RK_kafka kafka	ATLAS_HDOK topic	consume	Allowed	ranger-acl	172.26.1.1	49	
-	12/14/2016 11:08:25 AM	rangeragync	HDP_RK_kafka kafka	ATLAS_ENTITIES topic	describe	Denied	ranger-acl	172.26.1.1	7	

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 (ACL's). For HDFS that would be the user : rwx / group : rwx / other : rwx ACL's on folders and files.

Once this defaulting to component ACL's happens the audit events show a '-' in the 'Policy ID' column instead of a policy number. If a Ranger policy was in control of allowing/denying the policy number is shown.

Ranger AD Integration: Overview

Rangers AD Integration has 2 levels:

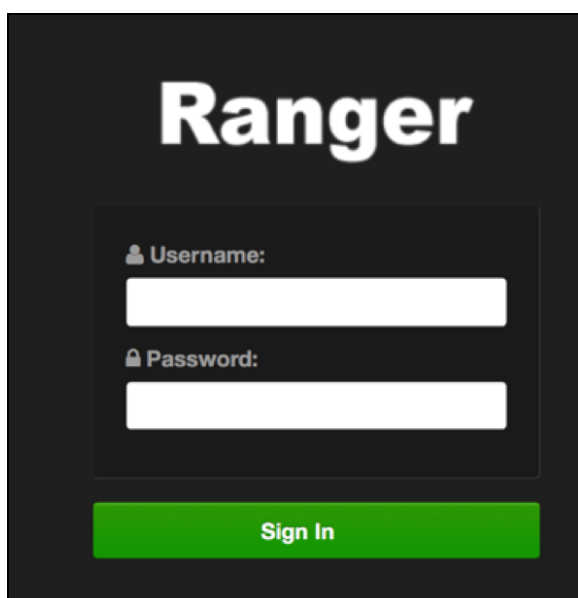
1. Ranger UI authentication (which users may log on to Ranger itself?)
2. Ranger User / group sync (which users / groups to define policies for?)

The configuration of both is done entirely on Ambari.

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 working with AD there are 2 options for defining who can access the Ranger UI; LDAP or ACTIVE_DIRECTORY. There is not much difference between them, just another set of properties.

Some of the configuration is in fact shared with the configuration of Ranger usersync as can be seen by the property with formats like `ranger_ug_ldap_bind_dn`. These properties are provided at runtime only with the value of another property by that name.

ACTIVE_DIRECTORY

The configuration for it is on Ambari > Ranger > Configs > Advanced:

 The image shows the Ambari Ranger configuration page for Active Directory settings. At the top, under "Authentication method", the "ACTIVE_DIRECTORY" option is selected with a blue dot. Below this, there are fields for "ranger.ldap.ad.base.dn" (set to "OU=...,OU=User Accounts,OU=CorpUsers,DC=field,DC=hortonworks,DC=com"), "ranger.ldap.ad.bind.dn" (set to "{{ranger_ug_ldap_bind_dn}}"), "ranger.ldap.ad.bind.password" (masked with asterisks), "Domain Name (Only for AD)" (set to "FIELD.HORTONWORKS.COM"), "ranger.ldap.ad.referral" (set to "follow"), "ranger.ldap.ad.url" (set to "{{ranger_ug_ldap_url}}"), and "ranger.ldap.ad.user.searchfilter" (set to "{sAMAccountName={0}}"). Each field has a lock icon and a refresh icon.

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

The `ranger.ldap.ad.user.searchfilter` 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.

With `ACTIVE_DIRECTORY` it is not possible to limit the scope of users that can access Ranger UI any further by refining the `ranger.ldap.ad.user.searchfilter` even further to :

```
(&(memberOf=CN=Hdp_admins,OU=Company,OU=User
Accounts,OU=CorpUsers,DC=field,DC=hortonworks,DC=com)(sAMAccountName={0}))
```

This does NOT work with the `ACTIVE_DIRECTORY` option.

LDAP

The other LDAP related properties do allow for more fine tuning:

The screenshot shows the Ranger configuration interface. The top section is 'Ranger Settings' and the bottom section is 'LDAP Settings'.

Ranger Settings:

- External URL: `http://[redacted]-hdp25-m-02:6080`
- Authentication method: ☒ LDAP, ☐ ACTIVE_DIRECTORY, ☐ UNIX, ☐ NONE
- HTTP enabled: ☒

LDAP Settings:

- `ranger.ldap.base.dn`: `OU=[redacted],OU=User Accounts,OU=CorpUsers,DC=field,DC=hortonworks,DC=com`
- `Bind User`: `{{(ranger_ug_ldap_bind_dn)}}`
- `Bind User Password`: Two masked password fields.
- `ranger.ldap.group.roleattribute`: `cn`
- `ranger.ldap.referral`: `follow`
- `LDAP URL`: `{{(ranger_ug_ldap_url)}}`
- `ranger.ldap.user.dnpattern`: `DC=intentionally,DC=wrong`
- `User Search Filter`: `(&(objectclass=user)(memberOf=CN=Hdp_admins,OU=[redacted],OU=User Accounts,OU=CorpUsers,DC=field,DC=hortonworks,DC=com)(sAMAccountName={0}))`

There is 1 catch though; the `ranger.ldap.user.dnpattern` is evaluated first, so usually putting a value like:

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

Would work, but has 2 by-effects; first 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 in stead of the `k.reshi` short name variant.

Second 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.

That adverse behavior can be worked around by intentionally putting a DN pattern (`DC=intentionally,DC=wrong`) in the `ranger.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=field,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, then the **User Search Filter** is not applied.

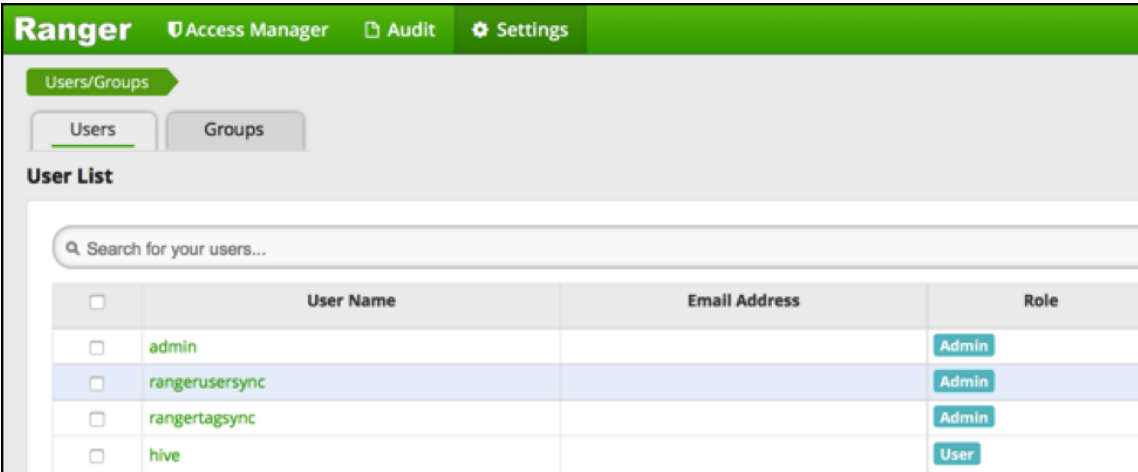
Ranger has a very simple approach to the internal user list that is kept in a relational schema. That list contains all users that were synced with AD ever, and all those users can potentially log on to Ranger UI. But only admin users can really do anything policy related things on the Ranger UI (see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

Beware that all this is still only about authentication to Ranger. Someone from the 'Hdp_admins' group would still not have a Ranger admin role.

Ranger UI Authorization

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

The Ranger authorization model is quite simple. It is maintained on the Ranger UI at Settings>Users/Groups :



The screenshot shows the Ranger UI interface. At the top is a green navigation bar with 'Ranger' and tabs for 'Access Manager', 'Audit', and 'Settings'. Below this is a 'Users/Groups' section with two tabs: 'Users' (selected) and 'Groups'. Under the 'Users' tab, there is a 'User List' section with a search bar labeled 'Search for your users...'. Below the search bar is a table with the following data:

<input type="checkbox"/>	User Name	Email Address	Role
<input type="checkbox"/>	admin		Admin
<input type="checkbox"/>	rangerusersync		Admin
<input type="checkbox"/>	rangertagsync		Admin
<input type="checkbox"/>	hive		User

A user can be either a normal user or an admin:

The screenshot shows the 'User Detail' form in the Ranger interface. At the top, there are two tabs: 'Users/Groups' and 'User Edit'. Below these is the 'User Detail' section. The 'Basic Info' tab is selected, showing fields for 'User Name *' (Kvothe), 'First Name *' (Kvothe), 'Last Name' (Reshi), 'Email Address' (empty), 'Select Role *' (a dropdown menu with 'Admin' selected and 'User' visible), and 'Group' (Rothfuss). At the bottom, there are 'Save' and 'Cancel' buttons.

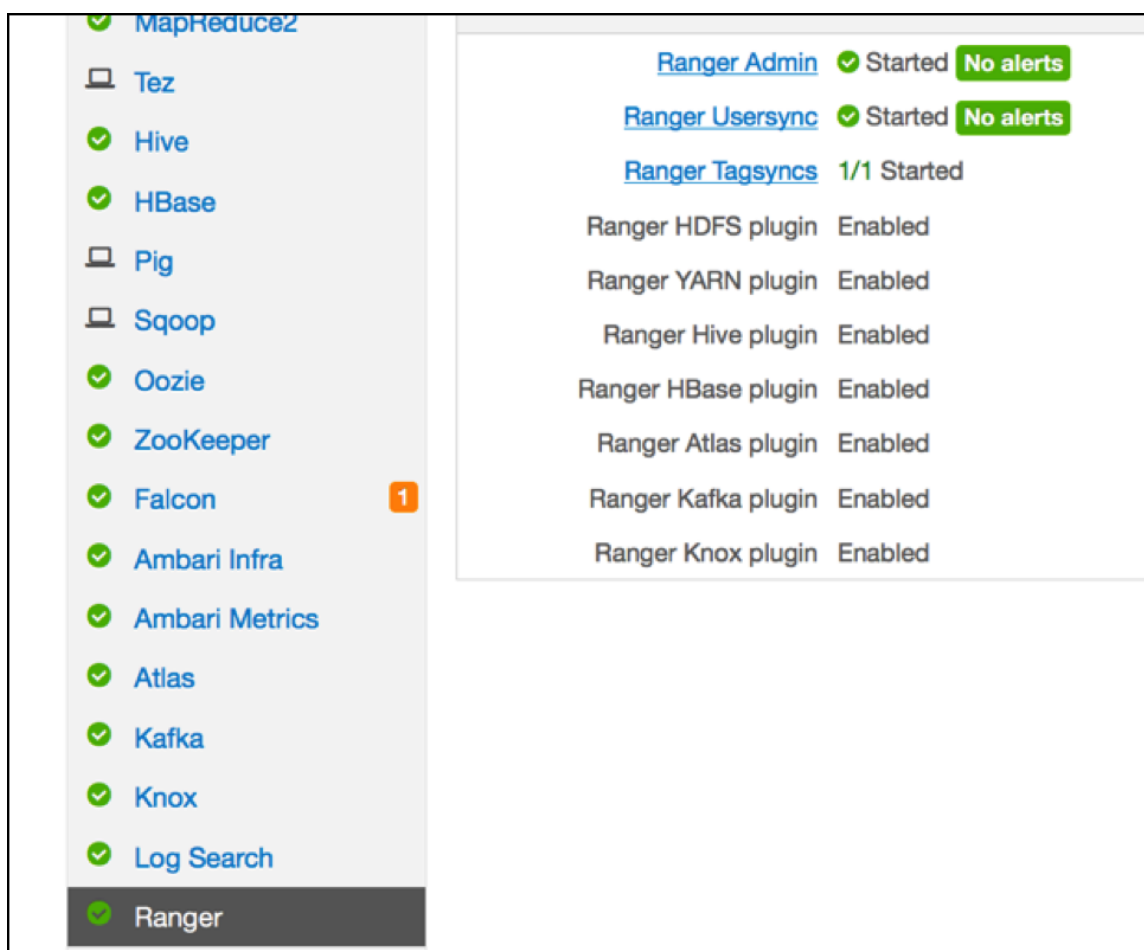
Only user with an Admin role can view or alter policies in Ranger.

Ranger Usersync

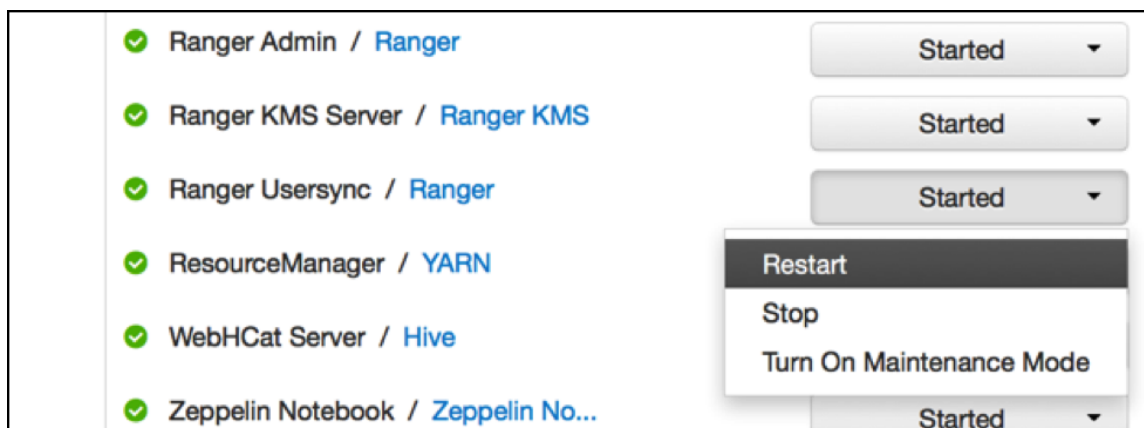
Reference information on Ranger usersync, when configuring Ranger AD integration.

A vital part of the Ranger architecture is the ability to get users and groups from the corporate AD to use in policy definitions.

Ranger usersync runs as separate daemon:



It can also be (re)started separately from Ambari:



Ranger Usersync Configuration

Usersync has a lot of moving parts and can have very different outcomes. Two main sets of properties govern the way users and groups are synchronized.

Without **Enable Group Search First** (a setting on the tab **Group Configs**) the primary access pattern is user based and groups will only be searched/added based on the users it finds first. In contrast, with **Enable Group Search First** enabled, the primary access pattern is group based (in turn based on the group search filter) and users will only be searched/added based on the group memberships it finds first

Sync Source

LDAP/AD

Common Configs User Configs Group Configs

Username Attribute

sAMAccountName

User Object Class

User

User Search Base

OU=CorpUsers,DC=field,DC=hortonworks,DC=com

User Search Filter

((memberOf=CN=Hdp_admins,OU= ,OU=User Accounts,OU=CorpUsers,DC=field,DC=horton

User Search Scope

sub

User Group Name Attribute

sAMAccountName

Group User Map Sync

Yes

Enable User Search

Yes

Value of 'User Search Base':
 OU=CorpUsers,DC=field,DC=hortonworks,DC=com

Value of 'User Search Filter':
 (| (memberOf=CN=Hdp_admins,OU=Company,OU=User
 Accounts,OU=CorpUsers,DC=field,DC=hortonworks,DC=com)
 (memberOf=CN=Hdp_users,OU=Company,OU=User
 Accounts,OU=CorpUsers,DC=field,DC=hortonworks,DC=com))

Value of 'User Group Name Attribute':
 sAMAccountName

Ranger User Info

Enable User Sync

Yes

Sync Source

LDAP/AD

Common Configs

User Configs

Group Configs

Enable Group Sync

Yes

Group Member Attribute

member

Group Name Attribute

name

Group Object Class

group

Group Search Base

OU= ,OU=User Accounts,OU=CorpUsers,DC=field,DC=hortonworks,DC=com

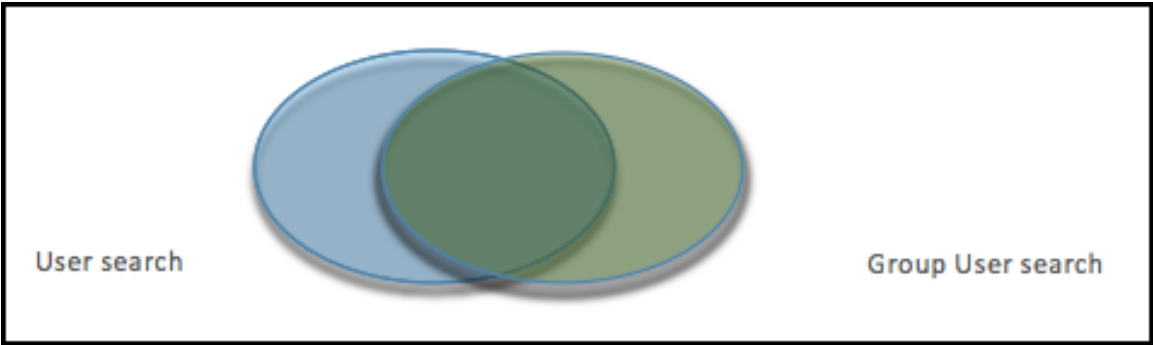
Group Search Filter

(|(CN=Hdp_users)(CN=Hdp_admins))

Enable Group Search First

Yes

Value of 'Group Search Base':
(|(CN=Hdp_users)(CN=Hdp_admins))

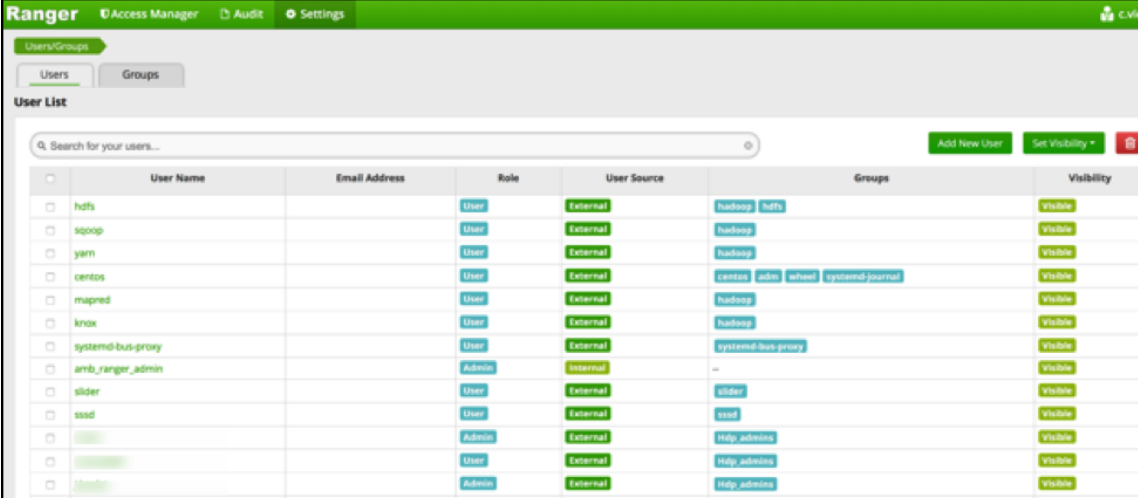


Beware that the filters on group level limit the returns on the user search and vice versa. In the graph above if the left oval would be the results of all users queried by the settings on the **User configs** and the right oval all users queried by **Group configs** settings, the eventual set of users that make it to the Ranger usersync is the overlap between the two.

Hortonworks therefore recommends to have the filters on both ends set exactly the same to potentially have a 100% overlap in the ovals.

In the example configuration given the scope of the usersync would be all members of both the groups 'Hdp_admins' and 'Hdp_users'.

The result in Ranger User list:



	User Name	Email Address	Role	User Source	Groups	Visibility
<input type="checkbox"/>	hdfs		User	External	hadoop hdfs	Visible
<input type="checkbox"/>	sqoop		User	External	hadoop	Visible
<input type="checkbox"/>	yarn		User	External	hadoop	Visible
<input type="checkbox"/>	centos		User	External	centos audit selinux systemd-journal	Visible
<input type="checkbox"/>	mapred		User	External	hadoop	Visible
<input type="checkbox"/>	knox		User	External	hadoop	Visible
<input type="checkbox"/>	systemd-bus-proxy		User	External	systemd-bus-proxy	Visible
<input type="checkbox"/>	amb_ranger_admin		Admin	Internal	-	Visible
<input type="checkbox"/>	slider		User	External	slider	Visible
<input type="checkbox"/>	sssd		User	External	sssd	Visible
<input type="checkbox"/>	hdp		Admin	External	hdp-admins	Visible
<input type="checkbox"/>	hdpadmin		User	External	hdp-admins	Visible
<input type="checkbox"/>	hdpadmin		Admin	External	hdp-admins	Visible

Regarding the other switches on the user and group sync pages, best of both worlds is to have **Enable Group Search First** and **Enable User Search** enabled at the same time.

The logging of a run of the usersync daemon can be retrieved from /var/log/ranger/usersync/usersync.log on the server hosting Ranger Admin. A successful run might output logging like below:

```

[restarted: true, user-search-enabled: true, ldap-errors: ignore]
08 Dec 2016 19:40:05 INFO UserGroupSync [UnixUserSyncThread] - Begin: Initial load of user/group from source=sink
08 Dec 2016 19:40:05 INFO LdapUserGroupBuilder [UnixUserSyncThread] - LdapUserGroupBuilder updateSink started
08 Dec 2016 19:40:05 INFO LdapUserGroupBuilder [UnixUserSyncThread] - Performing Group search first
08 Dec 2016 19:40:05 INFO LdapUserGroupBuilder [UnixUserSyncThread] - Adding Hdp_users to user
08 Dec 2016 19:40:05 INFO LdapUserGroupBuilder [UnixUserSyncThread] - Adding Hdp_users to user
08 Dec 2016 19:40:05 INFO LdapUserGroupBuilder [UnixUserSyncThread] - No. of members in the group Hdp_users = 2
08 Dec 2016 19:40:05 INFO LdapUserGroupBuilder [UnixUserSyncThread] - Adding Hdp_admins to user
08 Dec 2016 19:40:05 INFO LdapUserGroupBuilder [UnixUserSyncThread] - Adding Hdp_admins to user
08 Dec 2016 19:40:05 INFO LdapUserGroupBuilder [UnixUserSyncThread] - Adding Hdp_admins to user
08 Dec 2016 19:40:05 INFO LdapUserGroupBuilder [UnixUserSyncThread] - No. of members in the group Hdp_admins = 3
08 Dec 2016 19:40:05 INFO LdapUserGroupBuilder [UnixUserSyncThread] - LDAPUserGroupBuilder.getGroups() completed with group count: 2
08 Dec 2016 19:40:05 INFO LdapUserGroupBuilder [UnixUserSyncThread] - User search is enabled and hence computing user membership.
08 Dec 2016 19:40:05 INFO LdapUserGroupBuilder [UnixUserSyncThread] - Updating username for
08 Dec 2016 19:40:05 INFO LdapUserGroupBuilder [UnixUserSyncThread] - Updating username for
08 Dec 2016 19:40:05 INFO LdapUserGroupBuilder [UnixUserSyncThread] - Updating username for
08 Dec 2016 19:40:05 INFO LdapUserGroupBuilder [UnixUserSyncThread] - Updating username for
08 Dec 2016 19:40:05 INFO LdapUserGroupBuilder [UnixUserSyncThread] - Updating username for
08 Dec 2016 19:40:05 INFO LdapUserGroupBuilder [UnixUserSyncThread] - LDAPUserGroupBuilder.getUsers() completed with user count: 5
08 Dec 2016 19:40:05 INFO UserGroupSync [UnixUserSyncThread] - End: Initial load of user/group from source=sink
08 Dec 2016 19:40:05 INFO UserGroupSync [UnixUserSyncThread] - Done initializing user/group source and sink

```

From that log it clearly shows that the groups are synced first and that all users belonging to those groups are then retrieved according to its own settings, after which the user parts are enriched/overwritten by the returns from the user queries.

Beware:

If you don't enable **Enable User Search** that enrichment does NOT happen. Logging for such a run looks like this:

```
08 Dec 2016 19:24:28 INFO LDAPUserGroupBuilder [UnixSyncThread] - LDAPUserGroupBuilder Initialization completed with -- ldapuri: ldap://adml.fidd.hortonworks.com:389, ldapbind: binduserid@fidd.hortonworks.com, ldapbindpw: test, ldapsearchscope: SUBTREE_SCOPE, ldapchangelog: userSearchAttribute: {O=}, ldapfilterAttributes: {(memberofOfGroupAdmin,O=)}, ldapObjectClasses: {(memberofOfGroupAdmin,O=)}, ldapObjectClassNames: {(memberofOfGroupAdmin,O=)}, ldapObjectClassAliases: {}, ldapObjectClassAliasesMap: {}, ldapObjectClassAliasesSet: {}, ldapObjectClassAliasesList: {}, ldapObjectClassAliasesUnset: {}, ldapObjectClassAliasesUnsetList: {}, ldapObjectClassAliasesUnsetSet: {}, ldapObjectClassAliasesUnsetLabel: {}  
08 Dec 2016 19:24:28 INFO LDAPUserGroupBuilder [UnixSyncThread] - LDAPUserGroupBuilder postInit started  
08 Dec 2016 19:24:28 INFO LDAPUserGroupBuilder [UnixSyncThread] - Performing Group search First  
08 Dec 2016 19:24:28 INFO PolicyManagerGroupBuilder [UnixSyncThread] - Using principal = nangenusersync/rjk-hd25-e-6BDFTELO.HORTONWORKS.COM and keytab = /etc/security/keytabs/nangenusersync.service.keytab  
08 Dec 2016 19:24:28 INFO LDAPUserGroupBuilder [UnixSyncThread] - Adding Hdp_users to user  
08 Dec 2016 19:24:28 INFO LDAPUserGroupBuilder [UnixSyncThread] - Adding Hdp_users to user  
08 Dec 2016 19:24:28 INFO LDAPUserGroupBuilder [UnixSyncThread] - No. of members in the group Hdp_admins = 2  
08 Dec 2016 19:24:28 INFO PolicyManagerGroupBuilder [UnixSyncThread] - Using principal = nangenusersync/rjk-hd25-e-6BDFTELO.HORTONWORKS.COM and keytab = /etc/security/keytabs/nangenusersync.service.keytab  
08 Dec 2016 19:24:28 INFO LDAPUserGroupBuilder [UnixSyncThread] - Adding Hdp_admins to user  
08 Dec 2016 19:24:28 INFO LDAPUserGroupBuilder [UnixSyncThread] - Adding Hdp_admins to user  
08 Dec 2016 19:24:28 INFO LDAPUserGroupBuilder [UnixSyncThread] - User Search is disabled hence using the group member attribute for username.  
08 Dec 2016 19:24:28 INFO LDAPUserGroupBuilder [UnixSyncThread] - LongName:  
08 Dec 2016 19:24:28 INFO LDAPUserGroupBuilder [UnixSyncThread] - LongName:  
08 Dec 2016 19:24:28 INFO LDAPUserGroupBuilder [UnixSyncThread] - LongName:  
08 Dec 2016 19:24:28 INFO LDAPUserGroupBuilder [UnixSyncThread] - LongName:  
08 Dec 2016 19:24:28 INFO LDAPUserGroupBuilder [UnixSyncThread] - End Initial load of user/group from source=  
08 Dec 2016 19:24:28 INFO UserGroupSync [UnixSyncThread] - Done Initializing user/group source and sink
```

The result in Ranger UI are other user names (LongUserName) derived from ‘member’ group attributes full DN. You get the long name ‘James Kirk’ in the Ranger userlist in stead of j.kirk.

Ranger does not treat those as one and the same user:

Ranger					
Access Manager		Audit	Settings		
<input type="checkbox"/> mapred					
<input type="checkbox"/> knox			User	External	hadoop
<input type="checkbox"/> systemd-bus-proxy			User	External	systemd-bus-proxy
<input type="checkbox"/> amb_ranger_admin			Admin	Internal	--
<input type="checkbox"/> slider			User	External	slider
<input type="checkbox"/> sssd			User	External	sssd
<input type="checkbox"/> [redacted]			Admin	External	Hdp_admins
<input type="checkbox"/> [redacted]			User	External	Hdp_admins
<input type="checkbox"/> [redacted]			Admin	External	Hdp_admins
<input type="checkbox"/> [redacted]			User	External	Hdp_users
<input type="checkbox"/> [redacted]			User	External	Hdp_users
<input type="checkbox"/> hadoop			User	External	--
<input type="checkbox"/> rangerlookup			User	External	--
<input type="checkbox"/> [redacted]			User	External	Hdp_users
<input type="checkbox"/> [redacted]			User	External	Hdp_users
<input type="checkbox"/> [redacted]			User	External	Hdp_admins
<input type="checkbox"/> [redacted]			User	External	Hdp_admins
<input type="checkbox"/> [redacted]			User	External	Hdp_admins

Policies that were defined for user 'k.reshi' will not map to the user 'Kvothe Reshi' and vice versa. To prevent any confusion it is probably best to delete the long username versions from Rangers userlist.

Beware:

On the first page of Rangers user list there are lots of HDP system users. Most of them were put there by the Ranger installer and during the plugins installs:

Ranger

Access Manager

Audit

Settings

Search for your users...

<input type="checkbox"/>	User Name	Email Address	Role	User Source	
<input type="checkbox"/>	admin		Admin	Internal	--
<input type="checkbox"/>	rangerusersync		Admin	Internal	--
<input type="checkbox"/>	rangertagsync		Admin	Internal	--
<input type="checkbox"/>	hive		User	External	hadoop
<input type="checkbox"/>	infra-solr		User	External	hadoop
<input type="checkbox"/>	atlas		User	External	hadoop
<input type="checkbox"/>	ams		User	External	hadoop
<input type="checkbox"/>	falcon		User	External	hadoop users
<input type="checkbox"/>	systemd-network		User	External	systemd-network
<input type="checkbox"/>	ranger		User	External	hadoop ranger
<input type="checkbox"/>	kms		User	External	hadoop
<input type="checkbox"/>	polkitd		User	External	polkitd
<input type="checkbox"/>	nfsnobody		User	External	nfsnobody
<input type="checkbox"/>	spark		User	External	hadoop
<input type="checkbox"/>	hbase		User	External	hadoop
<input type="checkbox"/>	hcat		User	External	hadoop
<input type="checkbox"/>	zookeeper		User	External	hadoop
<input type="checkbox"/>	oozie		User	External	hadoop users
<input type="checkbox"/>	tez		User	External	hadoop users
<input type="checkbox"/>	zeppelin		User	External	hadoop
<input type="checkbox"/>	logsearch		User	External	hadoop
<input type="checkbox"/>	livy		User	External	hadoop

Do NOT remove those system users!

There are basic access policies based on those system users designed to keep a Ranger governed HDP component working after Ranger is given all control over that components authorizations. Without those policies/users many HDP components will be in serious trouble.

Ranger User Management

Reference information on Ranger user management, when configuring Ranger AD integration.

☐

Bast

☒

Auri

☐

Felurian

☐

Cinder

ew User

Set Visibility

Visibility

Visible

User can be easily remove from Ranger by checking the username in the list and hit the red **Delete** button. Ranger takes care of referential integrity so that user will also be removed from any policy.

Known Issue: Ranger Group Mapping

For Ranger AD integration, there is an issue with Ranger not being able to map a user on a group 'Hdp_admins' to a policy that allows/denies access to the group 'Hdp_admins'. The issue is on the capital characters that might be on a AD group name definition.

Most HDP components get the group information for a user via the SSSD daemon. When asked for the groups the user 'd.threpe' belongs to we get:

```
[centos@rjk-hdp25-m-01 ~]$ groups d.threpe
d.threpe : domain_users hdp_admins hadoop
```

So 'hdp_admins' all in lower case. Ranger does not treat this as the same value as 'Hdp_admins' which came via the group sync and was applied to some policies.

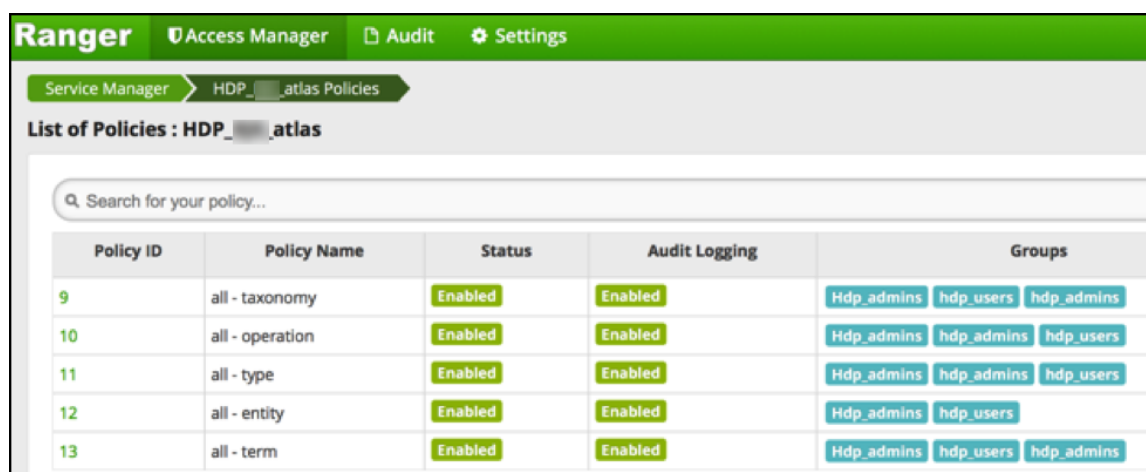
There is no way to make the group sync write or retrieve the group names all in lower case since there is no AD attribute that rewrites it in lowercase.

This issue can be worked around fortunately (till it gets solved). The solution is to define a local group in Ranger as a shadow group of a real group from AD, but then all in lower case:



<input type="checkbox"/>	system-auth-proxy	External
<input type="checkbox"/>	slider	External
<input type="checkbox"/>	sssd	External
<input type="checkbox"/>	Hdp_users	External
<input type="checkbox"/>	Hdp_admins	External
<input type="checkbox"/>	hdp_admins	Internal
<input type="checkbox"/>	hdp_users	Internal

If we now create policies and use that lower case 'shadow' group literal the result is that policies are correctly mapped to the AD groups again:



Policy ID	Policy Name	Status	Audit Logging	Groups
9	all - taxonomy	Enabled	Enabled	Hdp_admins hdp_users hdp_admins
10	all - operation	Enabled	Enabled	Hdp_admins hdp_admins hdp_users
11	all - type	Enabled	Enabled	Hdp_admins hdp_admins hdp_users
12	all - entity	Enabled	Enabled	Hdp_admins hdp_users
13	all - term	Enabled	Enabled	Hdp_admins hdp_users hdp_admins

*The 'Hdp_admins' entry does not have to be there, it is shown for clarification only. 'hdp_admins' is necessary to make it work.