

Providing Authorization with Apache Ranger

Date of Publish: 2019-08-26



Contents

| | |
|---|-----------|
| Using Ranger to Provide Authorization in Hadoop..... | 4 |
| Ranger Policies Overview..... | 4 |
| Ranger Tag-Based Policies..... | 4 |
| Tags and Policy Evaluation..... | 5 |
| Apache Ranger Access Conditions..... | 7 |
| Using the Ranger Console..... | 9 |
| Opening and Closing the Ranger Console..... | 10 |
| Ranger Console Navigation..... | 11 |
| Resource-Based Services and Policies..... | 13 |
| Configuring Resource-Based Services..... | 13 |
| Configure a Resource-based Service: HBase..... | 14 |
| Configure a Resource-based Service: HDFS..... | 16 |
| Configure a Resource-based Service: Hive..... | 18 |
| Configure a Resource-based Service: Kafka..... | 21 |
| Configure a Resource-based Service: Knox..... | 23 |
| Configure a Resource-based Service: Solr..... | 24 |
| Configure a Resource-based Service: Storm..... | 26 |
| Configure a Resource-based Service: YARN..... | 27 |
| Configure a Resource-based Service: Atlas..... | 29 |
| Configure a Resource-based Service: NiFi..... | 30 |
| Configure a Resource-based Service: NiFi Registry..... | 32 |
| Configuring Resource-Based Policies..... | 34 |
| Configure a Resource-based Policy: HBase..... | 34 |
| Configure a Resource-based Policy: HDFS..... | 37 |
| Configure a Resource-based Policy: Hive..... | 38 |
| Configure a Resource-based Policy: Kafka..... | 42 |
| Configure a Resource-based Policy: Knox..... | 43 |
| Configure a Resource-based Policy: Solr..... | 45 |
| Configure a Resource-based Policy: Storm..... | 47 |
| Configure a Resource-based Policy: YARN..... | 49 |
| Configure a Resource-based Policy: Atlas..... | 51 |
| Configure a Resource-based Policy: NiFi..... | 52 |
| Configure a Resource-based Policy: NiFi Registry..... | 54 |
| Wildcards and Variables in Resource-based Policies..... | 56 |
| Importing and Exporting Resource-Based Policies..... | 57 |
| Import Resource-Based Policies for a Specific Service..... | 59 |
| Import Resource-Based Policies for All Services..... | 61 |
| Export Resource-Based Policies for a Specific Service..... | 64 |
| Export All Resource-Based Policies for All Services..... | 65 |
| Row-level Filtering and Column Masking in Hive..... | 67 |
| Row-level Filtering in Hive with Ranger Policies..... | 67 |
| Dynamic Resource-Based Column Masking in Hive with Ranger Policies..... | 70 |
| Dynamic Tag-Based Column Masking in Hive with Ranger Policies..... | 74 |

| | |
|--|------------|
| Tag-Based Services and Policies..... | 77 |
| Adding a Tag-based Service..... | 77 |
| Adding Tag-Based Policies..... | 78 |
| Using Tag Attributes and Values in Ranger Tag-Based Policy Conditions..... | 81 |
| Adding a Tag-Based PII Policy..... | 83 |
| Default EXPIRES ON Tag Policy..... | 86 |
| Importing and Exporting Tag-Based Policies..... | 88 |
| Import Tag Based Policies..... | 90 |
| Export Tag-Based Policies..... | 92 |
| | |
| Create a Time-bound Policy..... | 94 |
| | |
| Ranger Security Zones..... | 96 |
| Overview..... | 96 |
| Adding a Ranger Security Zone..... | 97 |
| | |
| Administering Ranger Users, Groups, and Permissions..... | 101 |
| Add a User..... | 102 |
| Edit a User..... | 103 |
| Delete a User..... | 104 |
| Add a Group..... | 106 |
| Edit a Group..... | 106 |
| Delete a Group..... | 108 |
| Add/Edit Permissions..... | 109 |
| | |
| Administering Ranger Reports..... | 111 |
| View Ranger Reports..... | 111 |
| Search Ranger Reports..... | 112 |
| Export Reports..... | 113 |
| | |
| Adding a New Component to Apache Ranger..... | 114 |
| | |
| Configuring Advanced Authorization Settings..... | 116 |
| Developing a Custom Authorization Module..... | 117 |
| Special Requirements for High Availability Environments..... | 118 |
| Configure Advanced Usersync Settings..... | 118 |
| Configure User Sync LDAP SSL..... | 121 |
| Set Up Database Users Without Sharing DBA Credentials..... | 122 |
| Updating Ranger Admin Passwords..... | 122 |
| Ranger Password Requirements..... | 123 |

Using Ranger to Provide Authorization in Hadoop

Ranger manages access control through a user interface that ensures consistent policy administration across Hadoop data access components. Security administrators can define security policies at the database, table, column, and file levels, and can administer permissions for specific LDAP-based groups or individual users. Rules based on dynamic conditions such as time or geolocation can also be added to an existing policy rule. The Ranger authorization model is pluggable and can be easily extended to any data source using a service-based definition.

Once a user has been authenticated, their access rights must be determined. Authorization defines user access rights to resources. For example, a user may be allowed to create a policy and view reports, but not allowed to edit users and groups. You can use Ranger to set up and manage access to Hadoop services.

Ranger enables you to create services for specific Hadoop resources (HDFS, HBase, Hive, etc.) and add access policies to those services. You can also create tag-based services and add access policies to those services. Using tag-based policies enables you to control access to resources across multiple Hadoop components without creating separate services and policies in each component. You can also use Ranger TagSync to synchronize the Ranger tag store with an external metadata service such as Apache Atlas.

For more information on Ranger authorization, see the “HDP Security Features” Authorization overview.

Related Information

[HDP Security Features](#)

Ranger Policies Overview

Ranger has two types of policies: resource-based and tag-based.

Resource-based policies

Ranger enables you to configure resource-based services (HDFS, HBase, Hive, etc.) and add access policies to those services.

Tag-based policies

Ranger enables you to create tag-based services and add access policies to those services.

Ranger Tag-Based Policies

Ranger enables you to create tag-based services and add access policies to those services.

Tag-Based Policies Overview

- An important feature of Ranger tag-based authorization is the separation of resource-classification from access-authorization. For example, resources (HDFS file/directory, Hive database/table/column etc.) containing sensitive data such as social security numbers, credit card numbers, or sensitive health care data can be tagged with PII/PCI/PHI – either as the resource enters the Hadoop ecosystem or at a later time. Once a resource is tagged, the authorization for the tag would be automatically enforced, thus eliminating the need to create or update policies for the resource.
- Using tag-based policies also enables you to control access to resources across multiple Hadoop components without creating separate services and policies in each component.
- Tag details are stored in a tag store. Ranger TagSync can be used to synchronize the tag store with an external metadata service such as Apache Atlas.

Tag Store

Details of tags associated with resources are stored in a tag store. Apache Ranger plugins retrieve the tag details from the tag store for use during policy evaluation. To minimize the performance impact during policy evaluation (in finding tags for resources), Apache Ranger plugins cache the tags and periodically poll the tag store for any changes. When a change is detected, the plugins update the cache. In addition, the plugins store the tag details in a local cache file – just as the policies are stored in a local cache file. On component restart, the plugins will use the tag data from the local cache file if the tag store is not reachable.

Apache Ranger plugins download the tag details from the store managed by Ranger Admin. Ranger Admin persists the tag details in its policy store and provides a REST interface for the plugins to download the tag details.

Tags

Ranger Tags can have attributes. Tag attribute values can be used in Ranger tag-based policies to influence the authorization decision.

For example, to deny access to a resource after a specific date:

1. Add the EXPIRES_ON tag to the resource.
2. Add an expiry_date tag attribute and set its value to the expiry date.
3. Create a Ranger policy for the EXPIRES_ON tag.
4. Add a condition in this policy to deny access when the date specified in the expiry_date tag attribute is later than the current date.

Note that the EXPIRES_ON tag policy is created as the default policy in tag service instances.

TagSync

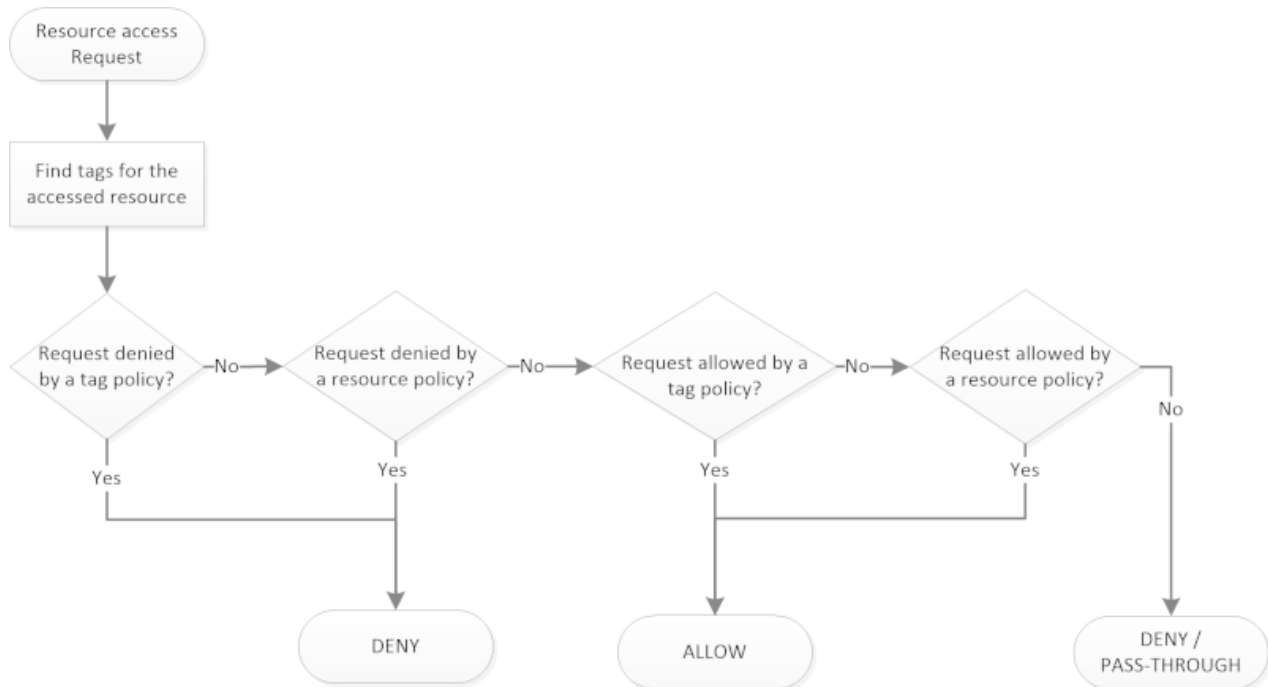
Ranger TagSync is used to synchronize the tag store with an external metadata service such as Apache Atlas. TagSync is a daemon process similar to the Ranger UserSync process.

Ranger TagSync receives tag details from Apache Atlas via change notifications. As tags are added to, updated, or deleted from resources in Apache Atlas, Ranger TagSync receives notifications and updates the tag store.

Tags and Policy Evaluation

When authorizing an access request, an Apache Ranger plugin evaluates applicable Ranger policies for the resource being accessed. The following diagram shows the details of the policy evaluation flow. More details on the steps in this workflow are provided in the subsequent sections.

Apache Ranger Policy Evaluation Flow with Tags



Apache Ranger Policy Evaluation Flow with Tags

Finding Tags

Apache Ranger supports a service to register context enrichers, which are used to update context data to the access request.

The Ranger Tag service, which is part of the tag-based policies feature, adds a context enricher named `RangerTagEnricher`. This context enricher is responsible for finding tags for the requested resource and adding the tag details to the request context. This context enricher keeps a cache of the available tags; while processing an access request, it finds the tags applicable for the requested resource and adds the tags to the request context. The context enricher keeps the cache updated by periodically polling Ranger Admin for changes.

Evaluating Tag-Based Policies

Once the list of tags for the requested resource is found, the Apache Ranger policy engine evaluates the tag-based policies applicable to the tags. If a policy for one of these tag results in a deny, access will be denied. If none of the tags are denied, and if a policy allows for one of the tags, access will be allowed. If there is no result for any tag, or if there are no tags for the resource, the policy engine will evaluate the resource-based policies to make the authorization decision.

Using Tags in Conditions

Apache Ranger allows the use of custom conditions while evaluating authorization policies. The Apache Ranger policy engine makes various request details – such as user, groups, resource, and context – available to the conditions. Tags in the request context, which are added by the enricher, are available to the conditions and can be used to influence the authorization decision.

The default policy in tag service instances, the `EXPIRES_ON` tag, uses such condition to check to see if the request date is later than the value specified in tag attribute `expiry_date`. This default policy does not work unless an `EXPIRES_ON` tag has been created in Atlas.

Related Information

[Apache Ranger Wiki> Context Enrichers](#)

Apache Ranger Access Conditions

The Apache Ranger access policy model consists of two major components: specification of the resources a policy is applied to, such as HDFS files and directories, Hive databases, tables, and columns, HBase tables, column-families, and columns, and so on; and the specification of access conditions for specific users and groups

Allow Deny and Exclude Conditions

Apache Ranger supports the following access conditions:

- Allow
- Exclude from Allow
- Deny
- Exclude from Deny

These access conditions enable you to set up fine-grained access control policies.

For example, you can allow access to a "finance" database to all users in the "finance" group, but deny access to all users in the "interns" group. Let's say that one of the members of the "interns" group, "scott", needs to work on an assignment that requires access to the "finance" database. In that case, you can add an Exclude from Deny condition that will allow user "scott" to access the "finance" database. The following image shows how this policy would be set up in Apache Ranger:

Policy Details :

Policy ID: 15

Policy Name: finance database enabled

Hive Database: finance Include

table: * Include **Resource**

Hive Column: * Include

Description: authorization for finance database

Audit Logging: YES

Allow Conditions

Allow Conditions:

| Select Group | Select User | Permissions | Delegate Admin |
|--------------|-------------|-------------|-------------------------------------|
| finance | Select User | All | <input checked="" type="checkbox"/> |

Exclude from Allow Conditions:

Deny Conditions

Deny Conditions:

| Select Group | Select User | Permissions | Delegate Admin |
|--------------|-------------|-------------|-------------------------------------|
| interns | Select User | All | <input checked="" type="checkbox"/> |

Exclude from Deny Conditions:

Deny Excludes

Exclude from Deny Conditions:

| Select Group | Select User | Permissions | Delegate Admin |
|--------------|-------------|-------------|--------------------------|
| Select Group | scot* | select | <input type="checkbox"/> |

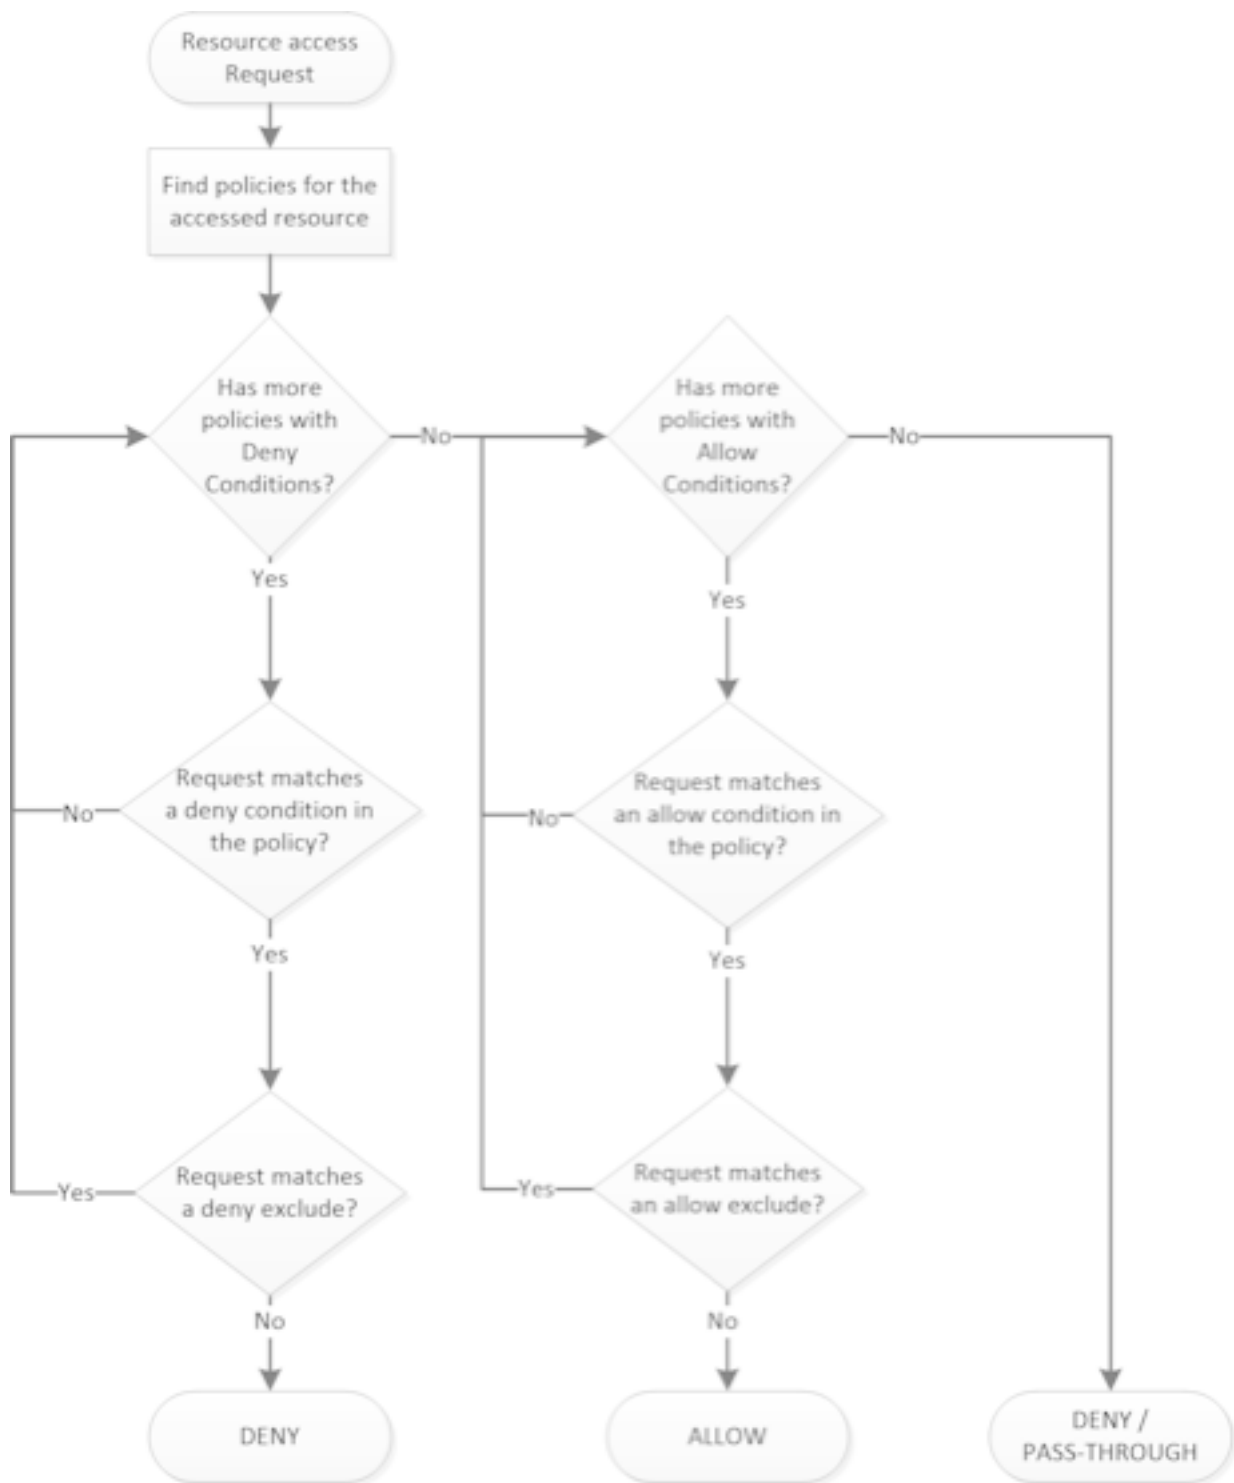
Enable Deny Conditions for Policies

The deny condition in policies is disabled by default and must be enabled for use.

1. From Ambari>Ranger>Configs>Advanced>Custom ranger-admin-site, add ranger.servicedef.enableDenyAndExceptionsInPolicies=true .
2. Restart Ranger.

Policy Evaluation of Access Conditions

Apache Ranger policies are evaluated in a specific order to ensure predictable results (if there is no access policy that allows access, the authorization request will typically be denied). The following diagram shows the policy evaluation work-flow:



Apache Ranger Policy Evaluation Flow

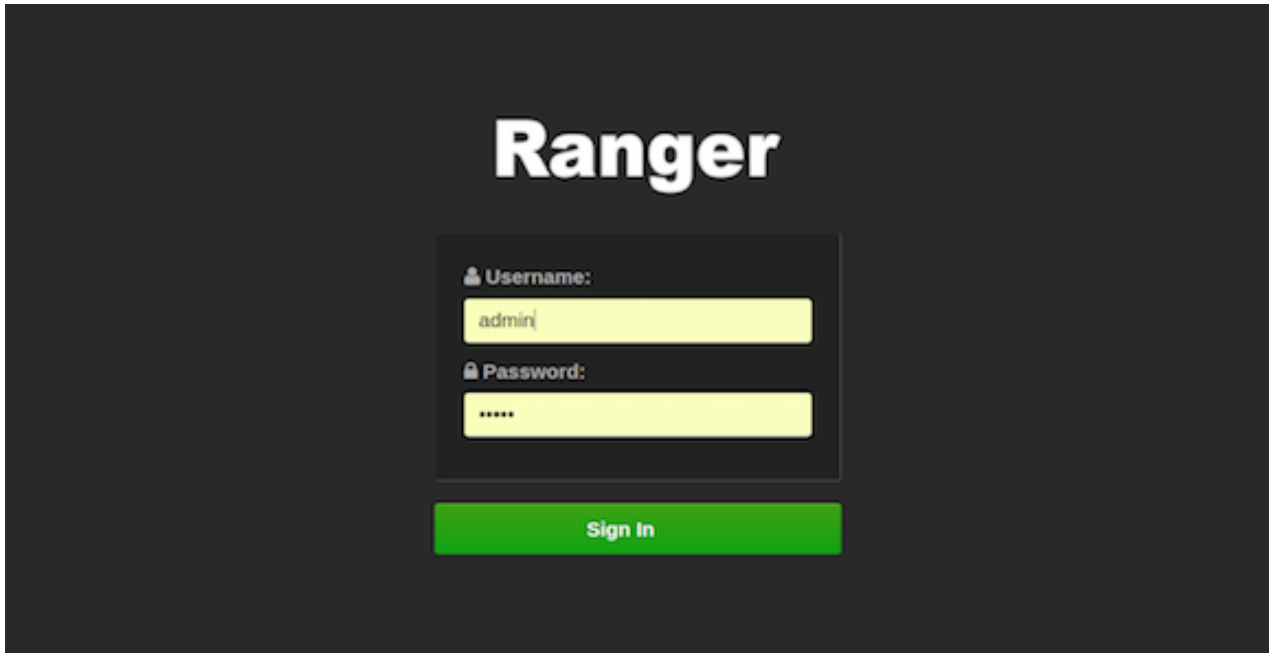
Using the Ranger Console

This chapter contains an overview of the Ranger console.

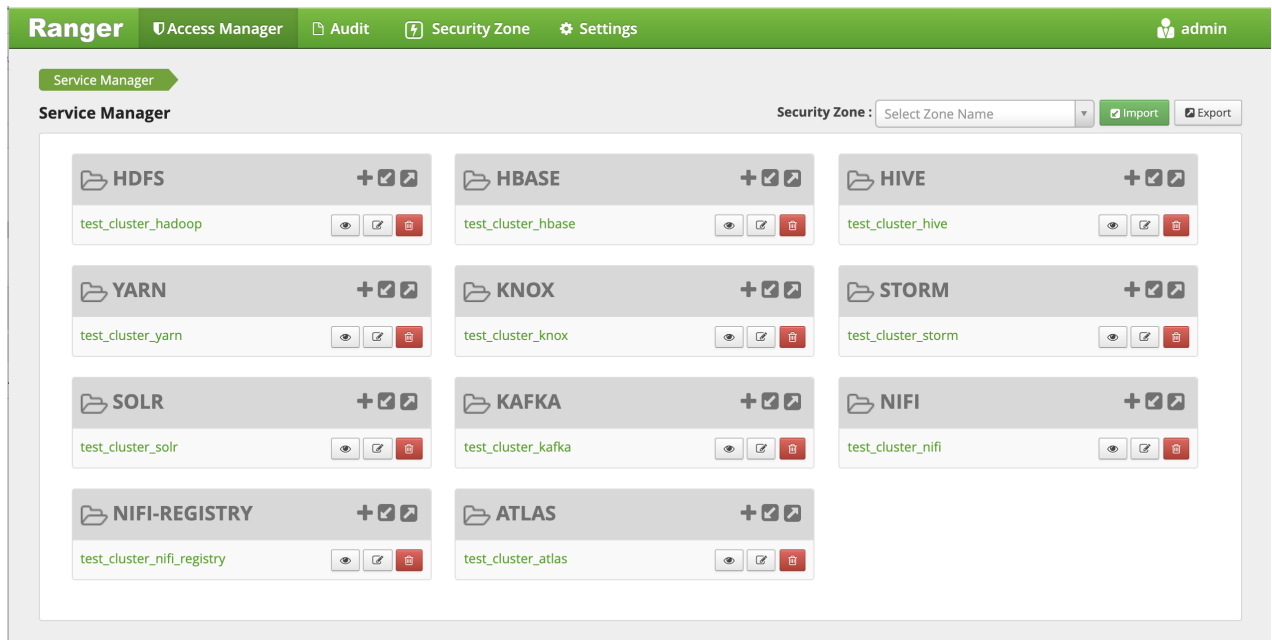
Opening and Closing the Ranger Console

Overview of how to open and close the Ranger console.

To open the Ranger Console, log in to the Ranger portal at `http://<your_ranger_server_address>:6080`. To log in, enter your user name and password, then click Sign In.



Ranger Console Home Page



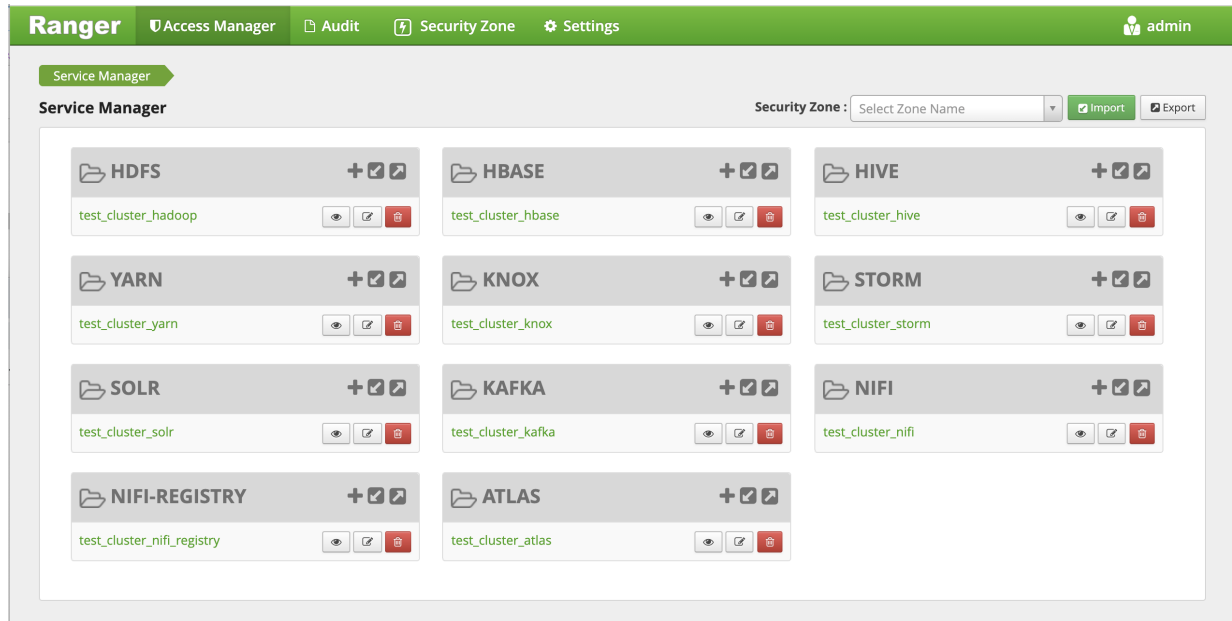
Ranger Login Console

After you log in, your user name is displayed at the top right of the Ranger Console.

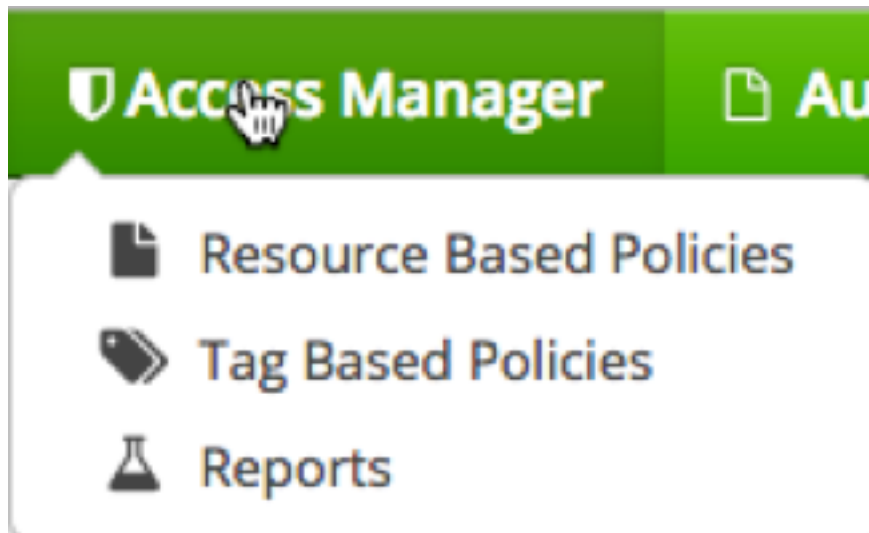
Ranger Console Navigation

Explains the basic Ranger console/GUI.

- The Service Manager for Resource Based Policies page is displayed when you log in to the Ranger Console. You can use this page to create services for Hadoop resources (HDFS, HBase, Hive, etc.) and add access policies to those resources.



Clicking Access Manager in the top menu opens the Service Manager for Resource Based Policies page, and also displays a submenu with links to Resource Based Policies, Tag Based Policies, and Reports (this submenu is also displayed when you pass the mouse over the Access Manager link).



- Access Manager > Resource Based Policies -- Opens the Service Manager for Resource Based Policies page. You can use this page to create services for resources (HDFS, HBase, Hive, etc.) and add access policies to those services.
- Access Manager > Tag Based Policies -- Opens the Service Manager for Tag Based Policies page. You can use this page to create tag-based services and add access policies to those services. Using tag-based policies enables you to control access to resources across multiple components without creating separate services and policies in each component.

- Access Manager > Reports -- Opens the Reports page. You can use this page to generate user access reports for resource and tag-based policies based on search criteria such as policy name, resource, group, and user name.
- Audit -- You can use the Audit page to monitor user activity at the resource level, and also to set up conditional auditing based on users, groups, or time. The Audit page includes the Access, Admin, Login Sessions, Plugins, Plugin Status, and User Sync tabs.

The screenshot shows the Ranger console interface with the 'Audit' tab selected. The top navigation bar includes 'Ranger', 'Access Manager', 'Audit', 'Security Zone', and 'Settings'. The user 'admin' is logged in. Below the navigation, there are tabs for 'Access', 'Admin', 'Login Sessions', 'Plugins', 'Plugin Status', and 'User Sync'. A search bar contains 'START DATE: 07/09/2019'. The main content area displays a table of audit entries with the following columns: Policy ID, Policy Version, Event Time, Application, User, Service Name / Type, Resource Name / Type, Access Type, Result, and Access Enforcer. The table shows several entries for 'hdfs' and 'hbaseRegional' applications, with results marked as 'Allowed'.

| Policy ID | Policy Version | Event Time | Application | User | Service Name / Type | Resource Name / Type | Access Type | Result | Access Enforcer |
|-----------|----------------|------------------------|---------------|-------|--------------------------|---------------------------------|--------------|---------|-----------------|
| -- | | 07/09/2019 11:41:28 AM | hdfs | oozie | test_cluster_hadoop hdfs | /user/oozie/share/lib path | READ_EXECUTE | Allowed | hadoop-acl a27 |
| -- | | 07/09/2019 11:41:27 AM | hdfs | spark | test_cluster_hadoop hdfs | /spark2-history path | READ_EXECUTE | Allowed | hadoop-acl a27 |
| -- | | 07/09/2019 11:41:27 AM | hdfs | spark | test_cluster_hadoop hdfs | /spark2-history/.27070b... path | WRITE | Allowed | hadoop-acl a27 |
| -- | | 07/09/2019 11:41:27 AM | hdfs | spark | test_cluster_hadoop hdfs | /spark2-history path | WRITE | Allowed | hadoop-acl a27 |
| 24 | 2 | 07/09/2019 11:41:24 AM | hbaseRegional | atlas | test_cluster_hbase hbase | atlas_janus/m column-family | get | Allowed | ranger-acl a27 |
| 24 | 2 | 07/09/2019 11:41:24 AM | hbaseRegional | atlas | test_cluster_hbase hbase | atlas_janus/s column-family | get | Allowed | ranger-acl a27 |

- Security Zone -- Lets you organize resource and tag-based services and policies into separate security zones. You can assign one or more administrators for each security zone. Security zone administrators can then create and update policies for their security zone.

The screenshot shows the Ranger console interface with the 'Security Zone' tab selected. The top navigation bar includes 'Ranger', 'Access Manager', 'Audit', 'Security Zone', and 'Settings'. The user 'admin' is logged in. The main content area displays the configuration for 'security-zone1'. It includes a search bar, a list of security zones, and a detailed view for 'security-zone1'. The 'Zone Administration' section shows 'Admin Users' as 'admin' and 'Auditor Users' as 'auditor1'. The 'Zone Tag Services' section shows 'tag_service1'. The 'Services' section shows a table with columns for Service Name, Service Type, and Resource.

| Service Name | Service Type | Resource |
|-------------------|--------------|-----------------|
| test_cluster_hive | HIVE | database : hive |

- Settings -- Enables you to manage and assign policy permissions to users and groups. Clicking or passing the mouse over Settings displays a submenu with links to the Users/Groups and Permissions pages.

| <input type="checkbox"/> | User Name | Email Address | Role | User Source | Groups | Visibility |
|--------------------------|----------------|---------------|-------|-------------|--------|------------|
| <input type="checkbox"/> | admin | | Admin | Internal | -- | Visible |
| <input type="checkbox"/> | rangerusersync | | Admin | Internal | -- | Visible |
| <input type="checkbox"/> | rangertagsync | | Admin | Internal | -- | Visible |
| <input type="checkbox"/> | yarn-ats | | User | External | hadoop | Visible |
| <input type="checkbox"/> | hive | | User | External | hadoop | Visible |
| <input type="checkbox"/> | infra-solr | | User | External | hadoop | Visible |
| <input type="checkbox"/> | elasticsearch | | User | External | hadoop | Visible |

Resource-Based Services and Policies

Ranger enables you to configure resource-based services for Hadoop components (e.g. HBase, Kafka, Storm, etc.) and add access policies to those services.

Configuring Resource-Based Services

The Service Manager for Resource Based Policies page is displayed when you log in to the Ranger Console. You can also access this page by selecting Access Manager > Resource Based Policies. You can use this page to create services for Hadoop resources (HDFS, HBase, Hive, etc.) and add access policies to those resources.

- To add a new resource-based service, click the Add icon



(in the applicable box on the Service Manager page. Enter the required configuration settings, then click **Add**.)

- To edit a resource-based service, click the Edit icon



(at the right of the service. Edit the service settings, then click Save to save your changes.)

- To delete a resource-based service, click the Delete icon



(at the right of the service. Deleting a service also deletes all of the policies for that service.)

The screenshot shows the Ranger Service Manager interface. At the top, there is a navigation bar with 'Ranger', 'Access Manager', 'Audit', 'Security Zone', and 'Settings'. The user is logged in as 'admin'. Below the navigation bar, there is a 'Service Manager' section with a 'Security Zone' dropdown menu and 'Import' and 'Export' buttons. The main area displays a grid of service cards for HDFS, HBASE, HIVE, YARN, KNOX, STORM, SOLR, KAFKA, NIFI, NIFI-REGISTRY, and ATLAS. Each card shows the service name, a cluster name (e.g., test_cluster_hadoop), and three icons: a plus sign for adding, a pencil for editing, and a trash can for deleting. Blue arrows point from the text 'Add Service', 'Edit Service', and 'Delete Service' to the respective icons on the ATLAS service card.

Configure a Resource-based Service: HBase

How to add an HBase service.

Procedure

1. On the Service Manager page, click the Add icon



(next to HBase.)

The Create Service page appears.

2. Enter the following information on the Create Service page:

Table 1: Service Details

| Field name | Description |
|--------------------|--|
| Service Name | The name of the service; required when configuring agents. |
| Description | A description of the service. |
| Active Status | Enabled or Disabled. |
| Select Tag Service | Select a tag-based service to apply the service and its tag-based policies to HBase. |

Table 2: Configuration Properties

| Field name | Description |
|------------|--|
| Username | The end system username that can be used for connection. |

| Field name | Description |
|-------------------------------------|--|
| Password | The password for the username entered above. |
| hadoop.security.authorization | The complete connection URL, including port and database name. (Default port: 10000.) For example, on the sandbox, jdbc:hive2://sandbox:10000/. |
| hbase.master.kerberos.principal | The Kerberos principal for the HBase Master. (Required only if Kerberos authentication is enabled.) |
| hbase.security.authentication | As noted in the hadoop configuration file hbase-site.xml. |
| hbase.zookeeper.property.clientPort | As noted in the hadoop configuration file hbase-site.xml. |
| hbase.zookeeper.quorum | As noted in the hadoop configuration file hbase-site.xml. |
| zookeeper.znode.parent | As noted in the hadoop configuration file hbase-site.xml. |
| Common Name for Certificate | The name of the certificate. This field is interchangeably named Common Name For Certificate and Ranger Plugin SSL CName in Create Service pages. |
| Add New Configurations | Add any other new configuration(s). |

3. Click **Test Connection**.

4. Click **Add**.

Configure a Resource-based Service: HDFS

How to add an HDFS service.

Procedure

1. On the Service Manager page, click the Add icon



(next to HDFS.)

The Create Service page appears.

2. Enter the following information on the Create Service page:

Table 3: Service Details

| Field name | Description |
|--------------------|---|
| Service Name | The name of the service; required when configuring agents. |
| Description | A description of the service. |
| Active Status | Enabled or Disabled. |
| Select Tag Service | Select a tag-based service to apply the service and its tag-based policies to HDFS. |

Table 4: Configuration Properties

| Field name | Description |
|------------|--|
| Username | The end system username that can be used for connection. |

| Field name | Description |
|---|--|
| Password | The password for the username entered above. |
| NameNode URL | hdfs://NAMENODE_FQDN:8020 The location of the Hadoop HDFS service, as noted in the hadoop configuration file core-site.xml OR (if this is a HA environment) the path for the primary NameNode. This field was formerly named fs.defaultFS. |
| Authorization Enabled | Authorization involves restricting access to resources. If enabled, user need authorization credentials. |
| Authentication Type | The type of authorization in use, as noted in the hadoop configuration file core-site.xml; either simple or Kerberos. (Required only if authorization is enabled). This field was formerly named hadoop.security.authorization. |
| hadoop.security.auth_to_local | Maps the login credential to a username with Hadoop; use the value noted in the hadoop configuration file, core-site.xml. |
| dfs.datanode.kerberos.principal | The principal associated with the datanode where the service resides, as noted in the hadoop configuration file hdfs-site.xml. (Required only if Kerberos authentication is enabled). |
| dfs.namenode.kerberos.principal | The principal associated with the NameNode where the service resides, as noted in the hadoop configuration file hdfs-site.xml. (Required only if Kerberos authentication is enabled). |
| dfs.secondary.namenode.kerberos.principal | The principal associated with the secondary NameNode where the service resides, as noted in the hadoop configuration file hdfs-site.xml. (Required only if Kerberos authentication is enabled). |
| RPC Protection Type | Only authorised user can view, use, and contribute to a dataset. A list of protection values for secured SASL connections. Values: Authentication, Integrity, Privacy |
| Common Name For Certificate | The name of the certificate. This field is interchangeably named Common Name For Certificate and Ranger Plugin SSL CName in Create Service pages. |
| Add New Configurations | Add any other new configuration(s). |

3. Click **Test Connection**.

4. Click **Add**.

Configure a Resource-based Service: Hive

How to add a Hive service.

Procedure

1. On the Service Manager page, click the Add icon



(next to Hive.)

The Create Service page appears.

2. Enter the following information on the Create Service page:

Table 5: Service Details

| Field name | Description |
|--------------------|---|
| Service Name | The name of the service; required when configuring agents. |
| Description | A description of the service. |
| Active Status | Enabled or Disabled. |
| Select Tag Service | Select a tag-based service to apply the service and its tag-based policies to Hive. |

Table 6: Configuration Properties

| Field name | Description |
|-----------------------|---|
| Username | The end system username that can be used for connection. |
| Password | The password for the username entered above. |
| jdbc.driver ClassName | The full classname of the driver used for Hive connections. Default: org.apache.hive.jdbc.HiveDriver |
| jdbc.url | The complete connection URL, including port and database name. (Default port: 10000.) For example, on the sandbox, jdbc:hive2://sandbox:10000/. |

| Field name | Description |
|-----------------------------|--|
| Common Name For Certificate | The name of the certificate. This field is interchangeably named Common Name For Certificate and Ranger Plugin SSL CName in Create Service pages. |
| Add New Configurations | Add any other new configuration(s). |

3. Click **Test Connection**.

4. Click **Add**.

What to do next

Usually, the Ranger Hive service definition uses the HiveServer2 (HS2) JDBC driver to fetch Hive database/table info for resource lookup and testing the connection. Alternatively, you can configure the service definition to use Hive metastore libraries connecting to the Hive metastore database directly. This is recommended when it is difficult to set up HiveServer2 on your cluster, such as when using HDCloud for AWS.

1. Under Ambari>Hive>Configs>Advanced, edit Hive properties:
2. Add the below properties to custom ranger-hive-plugin-properties:

```
ranger.service.config.param.enable.hive.metastore.lookup = true
```

```
ranger.service.config.param.hive.site.file.path = /etc/hive/conf/hive-site.xml
```



3. Save and restart required components.

4. To test the configuration is successful, create a new Hive service and specify the jdbc.url as "none", then run **Test**

Config Properties :

Username *

Password *

jdbc.driverClassName *

jdbc.url *

Common Name for Certificate

Add New Configurations

| Name | Value | |
|------------------------------|---|----------------------------------|
| enable.hive.metastore.lookup | <input type="text" value="true"/> | <input type="button" value="x"/> |
| hive.site.file.path | <input type="text" value="/etc/hive/conf/hive-site.xml"/> | <input type="button" value="x"/> |
| ambari.service.check.user | <input type="text" value="ambari-qa"/> | <input type="button" value="x"/> |

Connection.

Configure a Resource-based Service: Kafka

How to add a Kafka service.

Procedure

1. On the Service Manager page, click the Add icon



next to Kafka.

The Create Service page appears.

2. Enter the following information on the Create Service page:

Table 7: Service Details

| Field name | Description |
|--------------------|--|
| Service Name | The name of the service; required when configuring agents. |
| Description | A description of the service. |
| Active Status | Enabled or Disabled. |
| Select Tag Service | Select a tag-based service to apply the service and its tag-based policies to Kafka. |

Table 8: Configuration Properties

| Field name | Description |
|--------------------------|---|
| Username | The end system username that can be used for connection. |
| Password | The password for the username entered above. |
| ZooKeeper Connect String | Defaults to localhost:2181 (Provide FQDN of zookeeper host : 2181). |
| Ranger Plugin SSL CName | Provide common.name.for.certificate which is registered with Ranger (in Wire Encryption environment). This field is interchangeably named Common Name For Certificate and Ranger Plugin SSL CName in Create Service pages. |

| Field name | Description |
|------------------------|-------------------------------------|
| Add New Configurations | Add any other new configuration(s). |

3. Click **Test Connection**.
4. Click **Add**.

Configure a Resource-based Service: Knox

How to add a Knox service.

Procedure

1. On the Service Manager page, click the Add icon



The Create Service page appears.

2. Enter the following information on the Create Service page:

Table 9: Service Details

| Field name | Description |
|---------------|--|
| Service Name | The name of the service; required when configuring agents. |
| Description | A description of the service. |
| Active Status | Enabled or Disabled. |

| Field name | Description |
|--------------------|---|
| Select Tag Service | Select a tag-based service to apply the service and its tag-based policies to Knox. |

Table 10: Configuration Properties

| Field name | Description |
|-----------------------------|--|
| Username | The end system username that can be used for connection. |
| Password | The password for the username entered above. |
| knox.url | The Gateway URL for Knox. |
| Common Name For Certificate | The name of the certificate. This field is interchangeably named Common Name For Certificate and Ranger Plugin SSL CName in Create Service pages. |
| Add New Configurations | Add any other new configuration(s). |

3. Click **Test Connection**.

4. Click **Add**.

Configure a Resource-based Service: Solr

How to add a Solr service.

Procedure

1. On the Service Manager page, click the Add icon



(next to Solr.)

The Create Service page appears.

2. Enter the following information on the Create Service page:

Table 11: Service Details

| Field name | Description |
|--------------------|---|
| Service Name | The name of the service; required when configuring agents. |
| Description | A description of the service. |
| Active Status | Enabled or Disabled. |
| Select Tag Service | Select a tag-based service to apply the service and its tag-based policies to Solr. |

Table 12: Configuration Properties

| Field name | Description |
|-------------------------|---|
| Username | The end system username that can be used for connection. |
| Password | The password for the username entered above. |
| Solr URL | For HDP Search's Solr Instance: http://Solr_host:8983 For Ambari Infra's Solr Instance: http://Solr_host:8886 |
| Ranger Plugin SSL CName | Provide common.name.for.certificate which is registered with Ranger (in Wire Encryption environment). This field is interchangeably named Common Name For Certificate and Ranger Plugin SSL CName in Create Service pages. |

| Field name | Description |
|------------------------|-------------------------------------|
| Add New Configurations | Add any other new configuration(s). |

3. Click **Test Connection**.
4. Click **Add**.

Configure a Resource-based Service: Storm

How to add a Storm service.

Procedure

1. On the Service Manager page, click the Add icon



The Create Service page appears.

2. Enter the following information on the Create Service page:

Table 13: Service Details

| Field name | Description |
|---------------|--|
| Service Name | The name of the service; required when configuring agents. |
| Description | A description of the service. |
| Active Status | Enabled or Disabled. |

| Field name | Description |
|--------------------|--|
| Select Tag Service | Select a tag-based service to apply the service and its tag-based policies to Storm. |

Table 14: Configuration Properties

| Field name | Description |
|-----------------------------|--|
| Username | The end system username that can be used for connection. |
| Password | The password for the username entered above. |
| Nimbus URL | Host name of nimbus format, in the form: http://ipaddress:8080. This field was formerly named nimbus.url. |
| Common Name For Certificate | The name of the certificate. This field is interchangeably named Common Name For Certificate and Ranger Plugin SSL CName in Create Service pages. |
| Add New Configurations | Add any other new configuration(s). |

3. Click **Test Connection**.
4. Click **Add**.

Configure a Resource-based Service: YARN

How to add a YARN service.

Procedure

1. On the Service Manager page, click the Add icon



(next to YARN.)

The Create Service page appears.

2. Enter the following information on the Create Service page:

Table 15: Service Details

| Field name | Description |
|--------------------|---|
| Service Name | The name of the service; required when configuring agents. |
| Description | A description of the service. |
| Active Status | Enabled or Disabled. |
| Select Tag Service | Select a tag-based service to apply the service and its tag-based policies to YARN. |

Table 16: Configuration Properties

| Field name | Description |
|---------------------|--|
| Username | The end system username that can be used for connection. |
| Password | The password for the username entered above. |
| YARN REST URL | Http or https://RESOURCEMANAGER_FQDN:8088. |
| Authentication Type | The type of authorization in use, as noted in the hadoop configuration file core-site.xml; either simple or Kerberos. (Required only if authorization is enabled). This field was formerly named hadoop.security.authorization. |

| Field name | Description |
|-----------------------------|--|
| Common Name For Certificate | The name of the certificate. This field is interchangeably named Common Name For Certificate and Ranger Plugin SSL CName in Create Service pages. |
| Add New Configurations | Add any other new configuration(s). |

3. Click **Test Connection**.
4. Click **Add**.

Configure a Resource-based Service: Atlas

How to add an Atlas service.

Procedure

1. On the Service Manager page, click the Add icon



The Create Service page appears.

2. Enter the following information on the Create Service page:

Table 17: Service Details

| Field name | Description |
|--------------------|--|
| Service Name | The name of the service; required when configuring agents. |
| Description | A description of the service. |
| Active Status | Enabled or Disabled. |
| Select Tag Service | Select a tag-based service to apply the service and its tag-based policies to Atlas. |

Table 18: Configuration Properties

| Field name | Description |
|-----------------------------|--|
| Username | The end system username that can be used for connection. |
| Password | The password for the username entered above. |
| atlas.rest.address | Atlas host and port: : http://atlas_host_FQDN:21000. |
| Common Name For Certificate | The name of the certificate. This field is interchangeably named Common Name For Certificate and Ranger Plugin SSL CName in Create Service pages. |
| Add New Configurations | Add any other new configuration(s). |

3. Click **Test Connection**.

4. Click **Add**.

Configure a Resource-based Service: NiFi

How to add a NiFi service.

Procedure

1. On the Service Manager page, click the Add icon



(next to NiFi.

The Create Service page appears.

)

Ranger Access Manager Audit Security Zone Settings admin

Service Manager > Create Service

Create Service

Service Details :

Service Name *

Description

Active Status Enabled Disabled

Select Tag Service

Config Properties :

NiFi URL *

Authentication Type *

Keystore

Keystore Type

Keystore Password

Truststore

Truststore Type

Truststore Password

Add New Configurations

| Name | Value |
|----------------------|----------------------|
| <input type="text"/> | <input type="text"/> |

- Enter the following information on the Create Service page:

Table 19: Service Details

| Field name | Description |
|--------------------|---|
| Service Name | The name of the service; required when configuring agents. |
| Description | A description of the service. |
| Active Status | Enabled or Disabled. |
| Select Tag Service | Select a tag-based service to apply the service and its tag-based policies to NiFi. |

Table 20: Configuration Properties

| Field name | Description |
|---------------------|------------------------------------|
| NiFi URL | The complete NiFi host URL |
| Authentication Type | None or SSL |
| Keystore Type | The keystore type (JKS or PKCS12). |
| Keystore Password | The keystore password. |

| Field name | Description |
|------------------------|---|
| Truststore | The truststore to use when Ranger makes an https connection to NiFi. This truststore contains the public key of the certificate authority that signed the NiFi server certificates. |
| Truststore Type | The truststore type (JKS or PKCS12). |
| Truststore Password | The truststore password. |
| Keystore | The keystore to use when Ranger makes an https connection to NiFi. This keystore contains the certificate that represents the Ranger server. |
| Add New Configurations | Add any other new configuration(s). |

3. Click **Test Connection**.

4. Click **Add**.

Configure a Resource-based Service: NiFi Registry

How to add a NiFi Registry service.

Procedure

1. On the Service Manager page, click the Add icon



(next to NiFi Registry.

The Create Service page appears.

)

2. Enter the following information on the Create Service page:

Table 21: Service Details

| Field name | Description |
|--------------------|---|
| Service Name | The name of the service; required when configuring agents. |
| Description | A description of the service. |
| Active Status | Enabled or Disabled. |
| Select Tag Service | Select a tag-based service to apply the service and its tag-based policies to NiFi. |

Table 22: Configuration Properties

| Field name | Description |
|---------------------|------------------------------------|
| NiFi Registry URL | The complete NiFi Registry URL |
| Authentication Type | None or SSL |
| Keystore Type | The keystore type (JKS or PKCS12). |
| Keystore Password | The keystore password. |

| Field name | Description |
|------------------------|--|
| Truststore | The truststore to use when Ranger makes an https connection to the NiFi Registry. This truststore contains the public key of the certificate authority that signed the NiFi server certificates. |
| Truststore Type | The truststore type (JKS or PKCS12). |
| Truststore Password | The truststore password. |
| Keystore | The keystore to use when Ranger makes an https connection to the NiFi Registry. This keystore contains the certificate that represents the Ranger server. |
| Add New Configurations | Add any other new configuration(s). |

3. Click **Test Connection**.


4. Click **Add**.

Configuring Resource-Based Policies

To view the policies associated with a service, click the service name on the Resource Based Policies Service Manager page. The policies for that service will be displayed in a list, along with a search box.


- To add a new resource-based policy to the service, click **Add New Policy**.
- To edit a resource-based policy, click the Edit icon















()
at the right of the entry for that service. Edit the policy settings, then click Save to save your changes.

- To delete a resource-based policy, click the Delete icon



()
at the right of the entry for that service.

The screenshot shows the Ranger web interface. The top navigation bar includes 'Ranger', 'Access Manager', 'Audit', 'Security Zone', and 'Settings'. The user 'admin' is logged in. The breadcrumb trail is 'Service Manager > test_cluster_hbase Policies'. Below this, the title is 'List of Policies : test_cluster_hbase'. There is a search box for policies and an 'Add New Policy' button. The main content is a table with the following data:

| Policy ID | Policy Name | Policy Labels | Status | Audit Logging | Groups | Users | Action |
|-----------|-------------------------------------|---------------|---------|---------------|--------|-----------|---|
| 3 | all - table, column-family, column | -- | Enabled | Enabled | -- | hbase |    |
| 4 | Service Check User Policy for Hbase | -- | Enabled | Enabled | -- | ambari-qa |    |
| 24 | grant-1561403615836 | -- | Enabled | Enabled | -- | atlas |    |
| 25 | grant-1561403615982 | -- | Enabled | Enabled | -- | atlas |    |

Related Information

[Importing and Exporting Resource-Based Policies](#)

Configure a Resource-based Policy: HBase

How to add a new policy to an existing HBase service.

Procedure

1. On the Service Manager page, select an existing HBase service.

The List of Policies page appears.

2. Click **Add New Policy**.

The Create Policy page appears.

The screenshot shows the 'Create Policy' page in the Ranger interface. The page has a green header with 'Ranger' and navigation links for 'Access Manager', 'Audit', 'Security Zone', and 'Settings'. The user is logged in as 'admin'. The breadcrumb trail is 'Service Manager > test_cluster_hbase Policies > Create Policy'. The main form is titled 'Create Policy' and is divided into 'Policy Details' and 'Allow Conditions' sections. In the 'Policy Details' section, there are input fields for 'Policy Name *', 'Policy Label', 'HBase Table *', 'HBase Column-family *', 'HBase Column *', and 'Description'. There are also toggle switches for 'enabled/normal', 'include', and 'Audit Logging'. A 'Add Validity Period' button is visible in the top right of the form. A dropdown menu for 'add/edit permissions' is open, showing a list of permissions: Read (checked), Write, Create, Admin, and Select/Deselect All. The 'Allow Conditions' section includes 'Select Group', 'Select User', and 'Delegate Admin' fields, along with an 'Add Permissions' button.

3. Complete the Create Policy page as follows:

Table 23: Policy Details

| Label | Description |
|---------------------|---|
| Policy Name | Enter an appropriate policy name. This name cannot be duplicated across the system. This field is mandatory. |
| normal/override | Enables you to specify an override policy. When override is selected, the access permissions in the policy override the access permissions in existing policies. This feature can be used with Add Validity Period to create temporary access policies that override existing policies. |
| HBase Table | Select the appropriate database. Multiple databases can be selected for a particular policy. This field is mandatory. |
| HBase Column-family | For the selected table, specify the column families to which the policy applies. |
| HBase Column | For the selected table and column families, specify the columns to which the policy applies. |
| Description | (Optional) Describe the purpose of the policy. |
| Audit Logging | Specify whether this policy is audited. (De-select to disable auditing). |
| Policy Label | Specify a label for this policy. You can search reports and filter policies based on these labels. |

| Label | Description |
|---------------------|--|
| Add Validity Period | Specify a start and end time for the policy. |

Table 24: Allow Conditions

| Label | Description |
|----------------|--|
| Select Group | Specify the groups to which this policy applies. To designate a group as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. The public group contains all users, so granting access to the public group grants access to all users. |
| Select User | Specify the users to which this policy applies. To designate a user as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |
| Permissions | Add or edit permissions: Read, Write, Create, Admin, Select/Deselect All. |
| Delegate Admin | You can use Delegate Admin to assign administrator privileges to the users or groups specified in the policy. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |

- You can use the Plus (+) symbol to add additional conditions. Conditions are evaluated in the order listed in the policy. The condition at the top of the list is applied first, then the second, then the third, and so on.
- Click **Add**.

What to do next

Provide User Access to HBase Database Tables from the Command Line

HBase provides the means to manage user access to HBase database tables directly from the command line. The most commonly-used commands are:

- GRANT

Syntax:

```
grant '<user-or-group>', '<permissions>', '<table>'
```

For example, to create a policy that grants user1 read/write permission on the table usertable, the command would be:

```
grant 'user1', 'RW', 'usertable'
```

The syntax is the same for granting CREATE and ADMIN rights.

- REVOKE

Syntax:

```
revoke '<user-or-group>', '<usertable>'
```

For example, to revoke the read/write access of user1 to the table usertable, the command would be:

```
revoke 'user1', 'usertable'
```

**Note:**

Unlike Hive, HBase has no specific revoke commands for each user privilege.

Related Information

[Wildcards and Variables in Resource-based Policies](#)

Configure a Resource-based Policy: HDFS

How to add a new policy to an existing HDFS service.

About this task

Through configuration, Apache Ranger enables both Ranger policies and HDFS permissions to be checked for a user request. When the NameNode receives a user request, the Ranger plugin checks for policies set through the Ranger Service Manager. If there are no policies, the Ranger plugin checks for permissions set in HDFS.

We recommend that permissions be created at the Ranger Service Manager, and to have restrictive permissions at the HDFS level.

Procedure

1. On the Service Manager page, select an existing HDFS service.

The List of Policies page appears.

2. Click **Add New Policy**.

The Create Policy page appears.

3. Complete the Create Policy page as follows:

Table 25: Policy Details

| Field | Description |
|-------------|--|
| Policy Name | Enter a unique name for this policy. The name cannot be duplicated anywhere in the system. |

| Field | Description |
|---------------------|---|
| normal/override | Enables you to specify an override policy. When override is selected, the access permissions in the policy override the access permissions in existing policies. This feature can be used with Add Validity Period to create temporary access policies that override existing policies. |
| Resource Path | Define the resource path for the policy folder/file. The default recursive setting specifies that the resource path is recursive; you can also specify a non-recursive path. |
| Description | (Optional) Describe the purpose of the policy. |
| Audit Logging | Specify whether this policy is audited. (De-select to disable auditing). |
| Policy Label | Specify a label for this policy. You can search reports and filter policies based on these labels. |
| Add Validity Period | Specify a start and end time for the policy. |

Table 26: Allow Conditions

| Label | Description |
|----------------|--|
| Select Group | Specify the groups to which this policy applies. To designate a group as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. The public group contains all users, so granting access to the public group grants access to all users. |
| Select User | Specify the users to which this policy applies. To designate a user as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |
| Permissions | Add or edit permissions: Read, Write, Execute, Select/Deselect All. |
| Delegate Admin | You can use Delegate Admin to assign administrator privileges to the users or groups specified in the policy. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |

- You can use the Plus (+) symbol to add additional conditions. Conditions are evaluated in the order listed in the policy. The condition at the top of the list is applied first, then the second, then the third, and so on.
- Click **Add**.

Related Information

[Wildcards and Variables in Resource-based Policies](#)

Configure a Resource-based Policy: Hive

How to add a new policy to an existing Hive service.

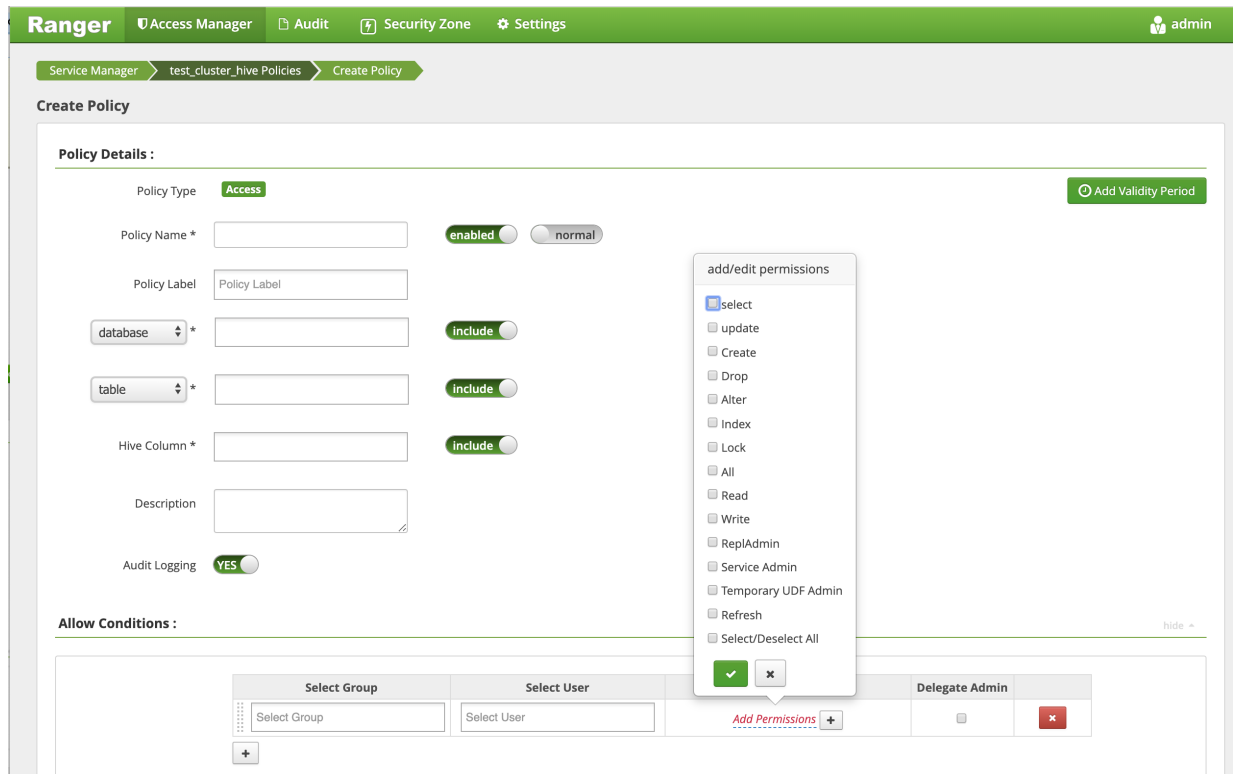
Procedure

- On the Service Manager page, select an existing Hive service.

The List of Policies page appears.

- Click **Add New Policy**.

The Create Policy page appears.



3. Complete the Create Policy page as follows:

Table 27: Policy Details

| Field | Description |
|-----------------|---|
| Policy Name | Enter an appropriate policy name. This name cannot be duplicated across the system. This field is mandatory. The policy is enabled by default. |
| normal/override | Enables you to specify an override policy. When override is selected, the access permissions in the policy override the access permissions in existing policies. This feature can be used with Add Validity Period to create temporary access policies that override existing policies. |
| Database | Type in the applicable database name. The autocomplete feature displays available databases based on the entered text. Include is selected by default to allow access. Select Exclude to deny access.. |
| table/udf | Specifies a table-based or UDF-based policy. Select table or udf, then type in the applicable table or UDF name. The autocomplete feature displays available tables based on the entered text. Include is selected by default to allow access. Select Exclude to deny access. |

| Field | Description |
|---------------------|--|
| Hive Column | <p>Type in the applicable Hive column name. The autocomplete feature displays available columns based on the entered text.</p> <p>Include is selected by default to allow access. Select Exclude to deny access.</p> <p>If using the Ranger Hive plugin with HiveServer2 or HiveServer2-LLAP, where column or description permissions include all, you must set a parameter for Hive columns to display as expected: in Ambari>Hive, under ranger-hive-security.xml, enter: <code>xasecure.hive.describetable.showcolumns.authorization.option=show-all</code>. Failure to set this parameter will result in the error message <code>HiveAccessControlException</code>.</p> |
| URL | <p>Specify the cloud storage path (for example <code>s3a://dev-admin/demo/campaigns.txt</code>) where the end-user permission is needed to read/write the Hive data from/to a cloud storage path.</p> <p>Permissions: READ operation on the URL permits the user to perform HiveServer2 operations which use S3 as data source for Hive tables. WRITE operation on the URL permits the user to perform HiveServer2 operations which write data to the specified S3 location.</p> <p>This feature is a Technical Preview: it is not ready for production deployment.</p> |
| URI | <p>Hive INSERT OVERWRITE queries require a Ranger URI policy to allow write operations, even if the user has write privilege granted through HDFS policy.</p> <p>Failure to specify this field will result in the following error: Error while compiling statement: FAILED: HiveAccessControlException Permission denied: user [jdoe] does not have [WRITE] privilege on [/tmp/*] (state=42000,code=40000)</p> <p>Example value: <code>/tmp/*</code></p> |
| Description | <p>(Optional) Describe the purpose of the policy.</p> <p>If using the Ranger Hive plugin with HiveServer2 or HiveServer2-LLAP, where column or description permissions include all, you must set a parameter for Hive columns to display as expected: in Ambari>Hive, under ranger-hive-security.xml, enter: <code>xasecure.hive.describetable.showcolumns.authorization.option=show-all</code>. Failure to set this parameter will result in the error message <code>HiveAccessControlException</code>.</p> |
| Hive Service Name | <p>hiveservice is used only in conjunction with Permissions=Service Admin. Enables a user who has Service Admin permission in Ranger to run the kill query API: <code>kill query <queryID></code>. Supported value: <code>*</code>. (Required)</p> |
| Audit Logging | <p>Specify whether this policy is audited. (De-select to disable auditing).</p> |
| Policy Label | <p>Specify a label for this policy. You can search reports and filter policies based on these labels.</p> |
| Add Validity Period | <p>Specify a start and end time for the policy.</p> |

Table 28: Allow Conditions

| Label | Description |
|----------------|---|
| Select Group | Specify the groups to which this policy applies. To designate a group as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. The public group contains all users, so granting access to the public group grants access to all users. |
| Select User | Specify the users to which this policy applies. To designate a user as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |
| Permissions | Add or edit permissions: Select, Update, Create, Drop, Alter, Index, Lock, All, ReplAdmin, Service Admin, Select/Deselect All. If using the Ranger Hive plugin with HiveServer2 or HiveServer2-LLAP, where column or description permissions include all, you must set a parameter for Hive columns to display as expected: in Ambari>Hive, under ranger-hive-security.xml, enter: xasecure.hive.describetable.showcolumns.authorization.option=show-all. Failure to set this parameter will result in the error message HiveAccessControlException. In order to execute repl dump, repl load, or repl status commands, you must set a parameter: in Ambari>Hive, under hive-site.xml, enter: hive.distcp.privileged.doAs=hive. Service Admin is used in conjunction with Hive Service Name and the kill query API: kill query <queryID> . |
| Delegate Admin | You can use Delegate Admin to assign administrator privileges to the users or groups specified in the policy. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |

4. You can use the Plus (+) symbol to add additional conditions. Conditions are evaluated in the order listed in the policy. The condition at the top of the list is applied first, then the second, then the third, and so on.

5. Click **Add**.

What to do next

Provide User Access to Hive Database Tables from the Command Line

Hive provides the means to manage user access to Hive database tables directly from the command line. The most commonly-used commands are:

- GRANT

Syntax:

```
grant <permissions> on table <table> to user <user or group>;
```

For example, to create a policy that grants user1 SELECT permission on the table default-hivesmoke22074, the command would be:

```
grant select on table default.hivesmoke22074 to user user1;
```

The syntax is the same for granting UPDATE, CREATE, DROP, ALTER, INDEX, LOCK, ALL, and ADMIN rights.

- REVOKE

Syntax:

```
revoke <permissions> on table <table> from user <user or group>;
```

For example, to revoke the SELECT rights of user1 to the table default.hivesmoke22074, the command would be:

```
revoke select on table default.hivesmoke22074 from user user1;
```

The syntax is the same for revoking UPDATE, CREATE, DROP, ALTER, INDEX, LOCK, ALL, and ADMIN rights.

Related Information

[Wildcards and Variables in Resource-based Policies](#)

Configure a Resource-based Policy: Kafka

How to add a new policy to an existing Kafka service.

Procedure

1. On the Service Manager page, select an existing Kafka service.
The List of Policies page appears.
2. Click **Add New Policy**.
The Create Policy page appears.

3. Complete the Create Policy page as follows:

Table 29: Policy Details

| Field | Description |
|-----------------|---|
| Policy Name | Enter an appropriate policy name. This name cannot be duplicated across the system. This field is mandatory. |
| normal/override | Enables you to specify an override policy. When override is selected, the access permissions in the policy override the access permissions in existing policies. This feature can be used with Add Validity Period to create temporary access policies that override existing policies. |

| Field | Description |
|---|--|
| Policy Label | Specify a label for this policy. You can search reports and filter policies based on these labels. |
| Topic | Kafka resource type. A topic is a category or feed name to which messages are published. |
| Transactional ID | Kafka resource type, uniquely identifies producers in a persistent way. |
| Cluster | Kafka resource type. |
| Delegation Token | Kafka resource type for authentication. |
| Description | (Optional) Describe the purpose of the policy. |
| Audit Logging | Specify whether this policy is audited. (De-select to disable auditing). |
| Add Validity Period | Specify a start and end time for the policy. |
| Policy Conditions (applied at the policy level) | Click the + icon, then specify an IP address range. |

Table 30: Allow Conditions

| Label | Description |
|---|--|
| Select Group | Specify the groups to which this policy applies. To designate a group as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. The public group contains all users, so granting access to the public group grants access to all users. |
| Select User | Specify the users to which this policy applies. To designate a user as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |
| Policy Conditions (applied at the item level) | Specify an IP address range. |
| Permissions | Add or edit permissions: Publish, Consume, Configure, Describe, Create, Delete, Describe Configs, Alter Configs, Select/Deselect All. |
| Delegate Admin | You can use Delegate Admin to assign administrator privileges to the users or groups specified in the policy. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |

- You can use the Plus (+) symbol to add additional conditions. Conditions are evaluated in the order listed in the policy. The condition at the top of the list is applied first, then the second, then the third, and so on.
- Click **Add**.

Related Information

[Wildcards and Variables in Resource-based Policies](#)

Configure a Resource-based Policy: Knox

How to add a new policy to an existing Knox service.

Procedure

- On the Service Manager page, select an existing Knox service.
The List of Policies page appears.
- Click **Add New Policy**.
The Create Policy page appears.

Ranger Access Manager Audit Security Zone Settings admin

Service Manager > test_cluster_knox Policies > Create Policy

Create Policy

Policy Details :

Policy Type: **Access** + Add Validity Period

Policy Name * enabled normal

Policy Label

Knox Topology * include

Knox Service * include

Description

Audit Logging **YES**

Policy Conditions +

No Conditions

Allow Conditions : hide >

| Select Group | Select User | Policy Conditions | Permissions | Delegate Admin |
|----------------------|----------------------|--|---|---|
| <input type="text"/> | <input type="text"/> | Add Conditions + | Add Permissions + | <input type="checkbox"/> × |

3. Complete the Create Policy page as follows:

Table 31: Policy Details

| Field | Description |
|---|---|
| Policy Name | Enter an appropriate policy name. This name cannot be duplicated across the system. This field is mandatory. |
| normal/override | Enables you to specify an override policy. When override is selected, the access permissions in the policy override the access permissions in existing policies. This feature can be used with Add Validity Period to create temporary access policies that override existing policies. |
| Knox Topology | Enter an appropriate Topology Name. |
| Knox Service | Enter an appropriate Service Name. |
| Description | (Optional) Describe the purpose of the policy. |
| Audit Logging | Specify whether this policy is audited. (De-select to disable auditing). |
| Policy Label | Specify a label for this policy. You can search reports and filter policies based on these labels. |
| Add Validity Period | Specify a start and end time for the policy. |
| Policy Conditions (applied at the policy level) | Click the + icon, then specify an IP address range. |

Table 32: Allow Conditions

| Label | Description |
|---|--|
| Select Group | Specify the groups to which this policy applies. To designate a group as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. The public group contains all users, so granting access to the public group grants access to all users. |
| Select User | Specify the users to which this policy applies. To designate a user as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |
| Policy Conditions (applied at the item level) | Specify an IP address range. |
| Permissions | Add or edit permissions: Allow |
| Delegate Admin | You can use Delegate Admin to assign administrator privileges to the users or groups specified in the policy. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |

Since Knox does not provide a command line methodology for assigning privileges or roles to users, the User and Group Permissions portion of the Knox Create Policy form is especially important.

4. You can use the Plus (+) symbol to add additional conditions. Conditions are evaluated in the order listed in the policy. The condition at the top of the list is applied first, then the second, then the third, and so on.
5. Click **Add**.

Related Information

[Wildcards and Variables in Resource-based Policies](#)

Configure a Resource-based Policy: Solr

How to add a new policy to an existing Solr service.

Procedure

1. On the Service Manager page, select an existing Solr service.
The List of Policies page appears.
2. Click **Add New Policy**.
The Create Policy page appears.

3. Complete the Create Policy page as follows:

Table 33: Policy Details

| Field | Description |
|---|---|
| Policy Name | Enter an appropriate policy name. This name cannot be duplicated across the system. This field is mandatory. |
| normal/override | Enables you to specify an override policy. When override is selected, the access permissions in the policy override the access permissions in existing policies. This feature can be used with Add Validity Period to create temporary access policies that override existing policies. |
| Solr Collection | For HDP Search's Solr Instance: http:host_ip:8983/solr For Ambari Infra's Solr Instance: http:host_ip:8886/solr |
| Description | (Optional) Describe the purpose of the policy. |
| Audit Logging | Specify whether this policy is audited. (De-select to disable auditing). |
| Policy Label | Specify a label for this policy. You can search reports and filter policies based on these labels. |
| Add Validity Period | Specify a start and end time for the policy. |
| Policy Conditions (applied at the policy level) | Click the + icon, then specify an IP address range. |

Table 34: Allow Conditions

| Label | Description |
|---|--|
| Select Group | Specify the groups to which this policy applies. To designate a group as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. The public group contains all users, so granting access to the public group grants access to all users. |
| Select User | Specify the users to which this policy applies. To designate a user as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |
| Policy Conditions (applied at the item level) | Specify an IP address range. |
| Permissions | Add or edit permissions: Query, Update, Others, Solr Admin, Select/Deselect All. |
| Delegate Admin | You can use Delegate Admin to assign administrator privileges to the users or groups specified in the policy. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |

- You can use the Plus (+) symbol to add additional conditions. Conditions are evaluated in the order listed in the policy. The condition at the top of the list is applied first, then the second, then the third, and so on.
- Click **Add**.

Related Information

[Wildcards and Variables in Resource-based Policies](#)

Configure a Resource-based Policy: Storm

How to add a new policy to an existing Storm service.

Procedure

- On the Service Manager page, select an existing Storm service.
The List of Policies page appears.
- Click **Add New Policy**.
The Create Policy page appears.

3. Complete the Create Policy page as follows:

Table 35: Policy Details

| Label | Description |
|---------------------|---|
| Policy Name | Enter an appropriate policy name. This name is cannot be duplicated across the system. This field is mandatory. |
| normal/override | Enables you to specify an override policy. When override is selected, the access permissions in the policy override the access permissions in existing policies. This feature can be used with Add Validity Period to create temporary access policies that override existing policies. |
| Storm Topology | Enter an appropriate Topology Name. |
| Description | (Optional) Describe the purpose of the policy. |
| Audit Logging | Specify whether this policy is audited. (De-select to disable auditing). |
| Policy Label | Specify a label for this policy. You can search reports and filter policies based on these labels. |
| Add Validity Period | Specify a start and end time for the policy. |

Table 36: Allow Conditions

| Label | Description |
|--------------|--|
| Select Group | Specify the groups to which this policy applies. To designate a group as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. The public group contains all users, so granting access to the public group grants access to all users. |

| Label | Description |
|-----------------------------------|---|
| Select User | Specify the users to which this policy applies. To designate a user as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |
| Storm User and Group Permissions* | Add or edit permissions. See the table below. |
| Delegate Admin | You can use Delegate Admin to assign administrator privileges to the users or groups specified in the policy. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |

Because Storm does not provide a command line methodology for assigning privileges or roles to users, the User and Group Permissions portion of the Storm Create Policy form is especially important.

Table 37: * Storm User and Group Permissions

| Actions | Description |
|-----------------------|---|
| Submit Topology | Allows a user to submit a topology. |
| File upload | Allows a user to upload files. |
| File Download | Allows a user to download files. |
| Kill Topology | Allows a user to kill the topology. |
| Rebalance | Allows a user to rebalance topologies. |
| Activate | Allows a user to activate a topology. |
| Deactivate | Allows a user to deactivate a topology. |
| Get Topology Conf | Allows a user to access a topology configuration. |
| Get Topology | Allows a user to access a topology. |
| Get User Topology | Allows a user to access a user topology. |
| Get Topology Info | Allows a user to access topology information. |
| Upload New Credential | Allows a user to upload a new credential. |
| Select/Deselect All | Select or deselect all permissions. |

4. You can use the Plus (+) symbol to add additional conditions. Conditions are evaluated in the order listed in the policy. The condition at the top of the list is applied first, then the second, then the third, and so on.
5. Click **Add**.

Related Information

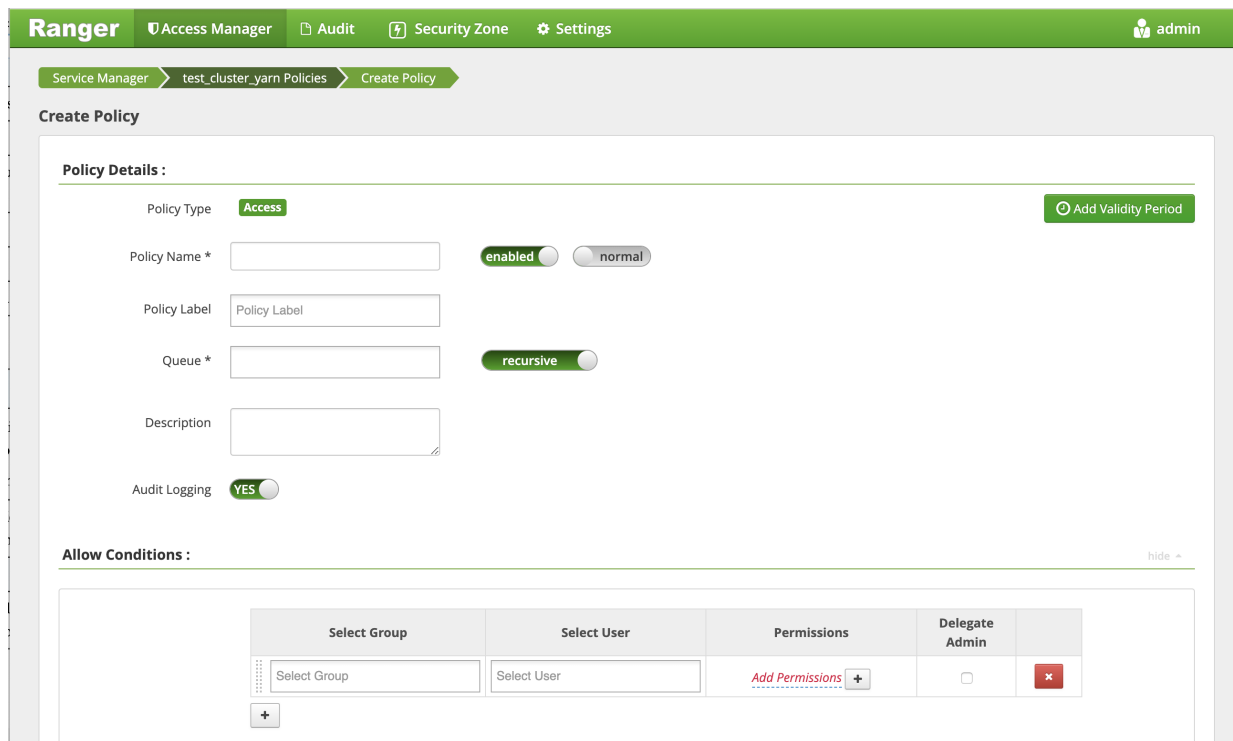
[Wildcards and Variables in Resource-based Policies](#)

Configure a Resource-based Policy: YARN

How to add a new policy to an existing YARN service.

Procedure

1. On the Service Manager page, select an existing YARN service.
The List of Policies page appears.
2. Click **Add New Policy**.
The Create Policy page appears.



3. Complete the Create Policy page as follows:

Table 38: Policy Details

| Field | Description |
|---------------------|---|
| Policy Name | Enter an appropriate policy name. This name cannot be duplicated across the system. This field is mandatory. |
| normal/override | Enables you to specify an override policy. When override is selected, the access permissions in the policy override the access permissions in existing policies. This feature can be used with Add Validity Period to create temporary access policies that override existing policies. |
| Queue | The YARN queue to which the policy applies. |
| Recursive | The default recursive setting specifies that the policy will also be applied to all sub-queues; you can also specify a non-recursive path. |
| Description | (Optional) Describe the purpose of the policy. |
| Audit Logging | Specify whether this policy is audited. (Deselect to disable auditing). |
| Policy Label | Specify a label for this policy. You can search reports and filter policies based on these labels. |
| Add Validity Period | Specify a start and end time for the policy. |

Table 39: Allow Conditions

| Label | Description |
|--------------|--|
| Select Group | Specify the groups to which this policy applies. To designate a group as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. The public group contains all users, so granting access to the public group grants access to all users. |

| Label | Description |
|----------------|---|
| Select User | Specify the users to which this policy applies. To designate a user as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |
| Permissions | Add or edit permissions: submit-app, admin-queue, Select/Deselect All. |
| Delegate Admin | You can use Delegate Admin to assign administrator privileges to the users or groups specified in the policy. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |

4. Click **Add**.

Related Information

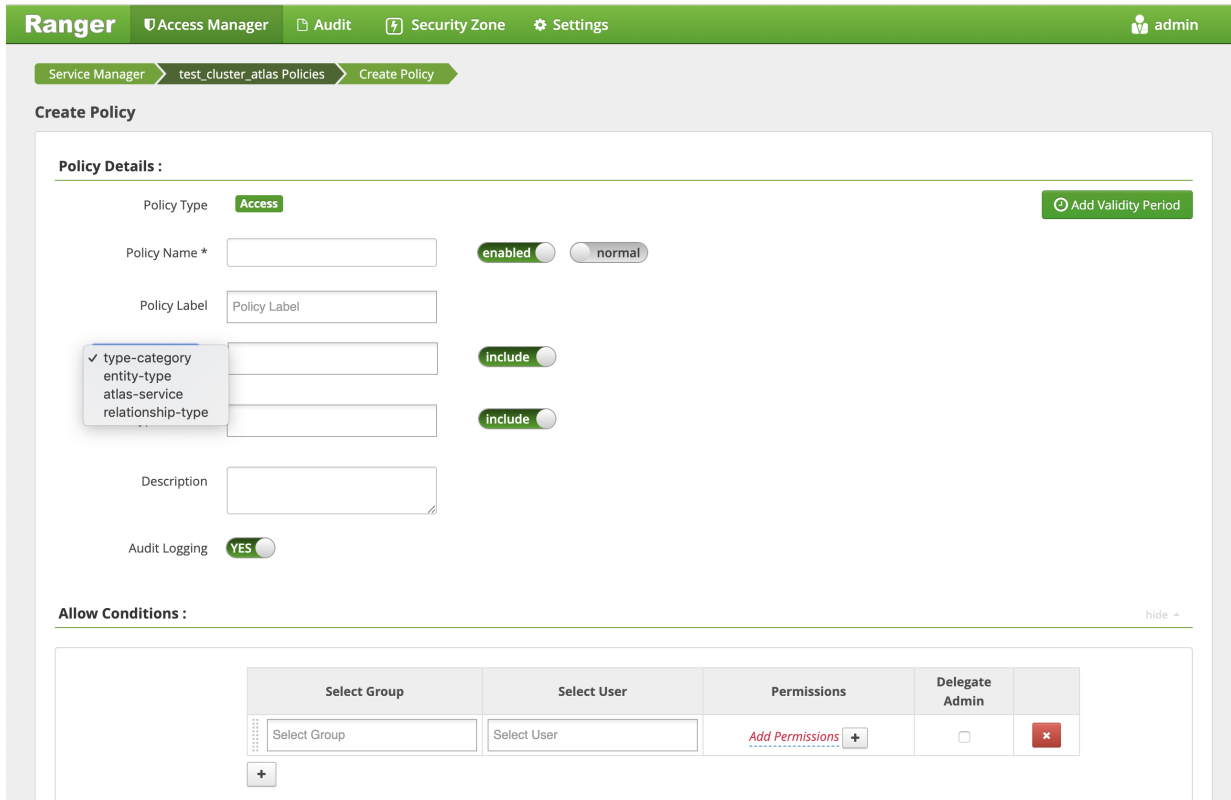
[Wildcards and Variables in Resource-based Policies](#)

Configure a Resource-based Policy: Atlas

How to add a new policy to an existing Atlas service.

Procedure

1. On the Service Manager page, select an existing Atlas service.
The List of Policies page appears.
2. Click **Add New Policy**.
The Create Policy page appears.



3. Complete the Create Policy page as follows:

Table 40: Policy Details

| Field | Description |
|---------------------|---|
| Policy Name | Enter an appropriate policy name. This name cannot be duplicated across the system. This field is mandatory. |
| normal/override | Enables you to specify an override policy. When override is selected, the access permissions in the policy override the access permissions in existing policies. This feature can be used with Add Validity Period to create temporary access policies that override existing policies. |
| type-category | Select type-category, entity-type, atlas-service, or relationship-type. |
| Description | (Optional) Describe the purpose of the policy. |
| Audit Logging | Specify whether this policy is audited. (De-select to disable auditing). |
| Policy Label | Specify a label for this policy. You can search reports and filter policies based on these labels. |
| Add Validity Period | Specify a start and end time for the policy. |

Table 41: Allow Conditions

| Label | Description |
|----------------|--|
| Select Group | Specify the groups to which this policy applies. To designate a group as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. The public group contains all users, so granting access to the public group grants access to all users. |
| Select User | Specify the users to which this policy applies. To designate a user as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |
| Permissions | Add or edit permissions: Create Type, Update Type, Delete Type, Select/Deselect All. |
| Delegate Admin | You can use Delegate Admin to assign administrator privileges to the users or groups specified in the policy. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |

- You can use the Plus (+) symbol to add additional conditions. Conditions are evaluated in the order listed in the policy. The condition at the top of the list is applied first, then the second, then the third, and so on.
- Click **Add**.

Related Information

[Wildcards and Variables in Resource-based Policies](#)

Configure a Resource-based Policy: NiFi

How to add a new policy to an existing Atlas service.

Procedure

- On the Service Manager page, select an existing NiFi service.
The List of Policies page appears.
- Click **Add New Policy**.

The Create Policy page appears.

The screenshot shows the 'Create Policy' page in the Ranger interface. The page is titled 'Create Policy' and is part of the 'test_cluster_nifi Policies' section. The form includes the following fields and controls:

- Policy Details:**
 - Policy Type: **Access** (with a green 'Add Validity Period' button)
 - Policy Name *: (with 'enabled' and 'normal' toggle buttons)
 - Policy Label:
 - NIFI Resource Identifier *:
 - Description:
 - Audit Logging: **YES** (with a toggle button)
- Allow Conditions:** (with a 'hide -' link)

| Select Group | Select User | Permissions | Delegate Admin | |
|---|--|--------------------------|--------------------------|----------|
| <input type="text" value="Select Group"/> | <input type="text" value="Select User"/> | Add Permissions + | <input type="checkbox"/> | x |

3. Complete the Create Policy page as follows:

Table 42: Policy Details

| Field | Description |
|--------------------------|---|
| Policy Name | Enter an appropriate policy name. This name cannot be duplicated across the system. This field is mandatory. |
| normal/override | Enables you to specify an override policy. When override is selected, the access permissions in the policy override the access permissions in existing policies. This feature can be used with Add Validity Period to create temporary access policies that override existing policies. |
| NiFi Resource Identifier | In a NiFi cluster, all nodes must be granted the ability to view and modify component data in order for user to list or empty queues in processor component outbound connections. With Ranger this can be accomplished by using a wildcard to grant all of the NiFi nodes read and write access to the /data/* NiFi resource. |
| Description | (Optional) Describe the purpose of the policy. |
| Audit Logging | Specify whether this policy is audited. (De-select to disable auditing). |
| Policy Label | Specify a label for this policy. You can search reports and filter policies based on these labels. |
| Add Validity Period | Specify a start and end time for the policy. |

Table 43: Allow Conditions

| Label | Description |
|----------------|--|
| Select Group | Specify the groups to which this policy applies. To designate a group as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. The public group contains all users, so granting access to the public group grants access to all users. |
| Select User | Specify the users to which this policy applies. To designate a user as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |
| Permissions | Add or edit permissions: Read, Write, Select/Deselect All. |
| Delegate Admin | You can use Delegate Admin to assign administrator privileges to the users or groups specified in the policy. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |

- You can use the Plus (+) symbol to add additional conditions. Conditions are evaluated in the order listed in the policy. The condition at the top of the list is applied first, then the second, then the third, and so on.
- Click **Add**.

Related Information

[NiFi Ranger based policy descriptions](#)

[What policy in Ranger should be used for connections in NiFi?](#)

Configure a Resource-based Policy: NiFi Registry

How to add a new policy to an existing Atlas service.

Procedure

- On the Service Manager page, select an existing NiFi Registry service.
The List of Policies page appears.
- Click **Add New Policy**.
The Create Policy page appears.

3. Complete the Create Policy page as follows:

Table 44: Policy Details

| Field | Description |
|-----------------------------------|---|
| Policy Name | Enter an appropriate policy name. This name cannot be duplicated across the system. This field is mandatory. |
| normal/override | Enables you to specify an override policy. When override is selected, the access permissions in the policy override the access permissions in existing policies. This feature can be used with Add Validity Period to create temporary access policies that override existing policies. |
| NiFi Registry Resource Identifier | In a NiFi cluster, all nodes must be granted the ability to view and modify component data in order for user to list or empty queues in processor component outbound connections. With Ranger this can be accomplished by using a wildcard to grant all of the NiFi nodes read and write access to the /data/* NiFi resource. |
| Description | (Optional) Describe the purpose of the policy. |
| Audit Logging | Specify whether this policy is audited. (De-select to disable auditing). |
| Policy Label | Specify a label for this policy. You can search reports and filter policies based on these labels. |
| Add Validity Period | Specify a start and end time for the policy. |

Table 45: Allow Conditions

| Label | Description |
|----------------|--|
| Select Group | Specify the groups to which this policy applies. To designate a group as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. The public group contains all users, so granting access to the public group grants access to all users. |
| Select User | Specify the users to which this policy applies. To designate a user as an Administrator, select the Delegate Admin check box. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |
| Permissions | Add or edit permissions: Read, Write, Delete, Select/Deselect All. |
| Delegate Admin | You can use Delegate Admin to assign administrator privileges to the users or groups specified in the policy. Administrators can edit or delete the policy, and can also create child policies based on the original policy. |

- You can use the Plus (+) symbol to add additional conditions. Conditions are evaluated in the order listed in the policy. The condition at the top of the list is applied first, then the second, then the third, and so on.
- Click **Add**.

Related Information

[NiFi Ranger based policy descriptions](#)

[What policy in Ranger should be used for connections in NiFi?](#)

Wildcards and Variables in Resource-based Policies

Reference for wildcards and variables in resource-based policies.

Ranger Authorization Resource Policy Wildcard Characters

Wildcard characters can be included in the resource path, the database name, the table name, or the column name:

- * indicates zero or more occurrences of characters
- ? indicates a single character

Ranger Authorization Resource Policy {USER} Variable

The variable {USER} can be used to autofill the accessing user, for example:

In **Select User**, choose {USER}.

In **Resource Path**, enter data_{USER}.

Ranger Authorization Resource Policy {USER} Variable Recommended Practices and Customizability

Ranger requires that string '{USER}' is used to represent accessing user as the user in the policy-item in a Ranger policy. However, Ranger provides flexible way of customizing the string that is used as shorthand to represent the accessing user's name in the policy resource specification. By default, Ranger policy resource specification expects characters '{' and '}' as delimiters for string 'USER', however, ranger supports customizable way of specifying delimiter characters, escaping those delimiters, and the string 'USER' itself by prefixing it with another, user-specified string on a per resource-level basis in the service definition of each component supported by Ranger.

For example, if for a certain HDFS installation, if the path names may contain '{' or '}' as valid characters, but not '%' character, then the service-definition for HDFS can be specified as:

```
"resources" : [
{
```



```

    "itemId": 1,
    "name": "path",
    "type": "path",
    "level": 10,
    "parent": "",
    "mandatory": true,
    "lookupSupported": true,
    "recursiveSupported": true,
    "excludesSupported": false,
    "matcher":
"org.apache.ranger.plugin.resourcematcher.RangerPathResourceMatcher",
    "matcherOptions": {"wildcard": true, "ignoreCase": false},
    "replaceTokens": true, "tokenDelimiterStart": "%", "tokenDelimiterEnd": "%",
    "tokenDelimiterPrefix": "rangerToken:" }
    "validationRegex": "",
    "validationMessage": "",
    "uiHint": "",
    "label": "Resource Path",
    "description": "HDFS file or directory
path"
}
]

```

Corresponding ranger policy for the use case for HDFS will be written as follow:

```

resource: path=/home/%rangerToken:USER%
user: {USER}
permissions: all, delegateAdmin=true

```

The following customizable matcherOptions are available for this feature:

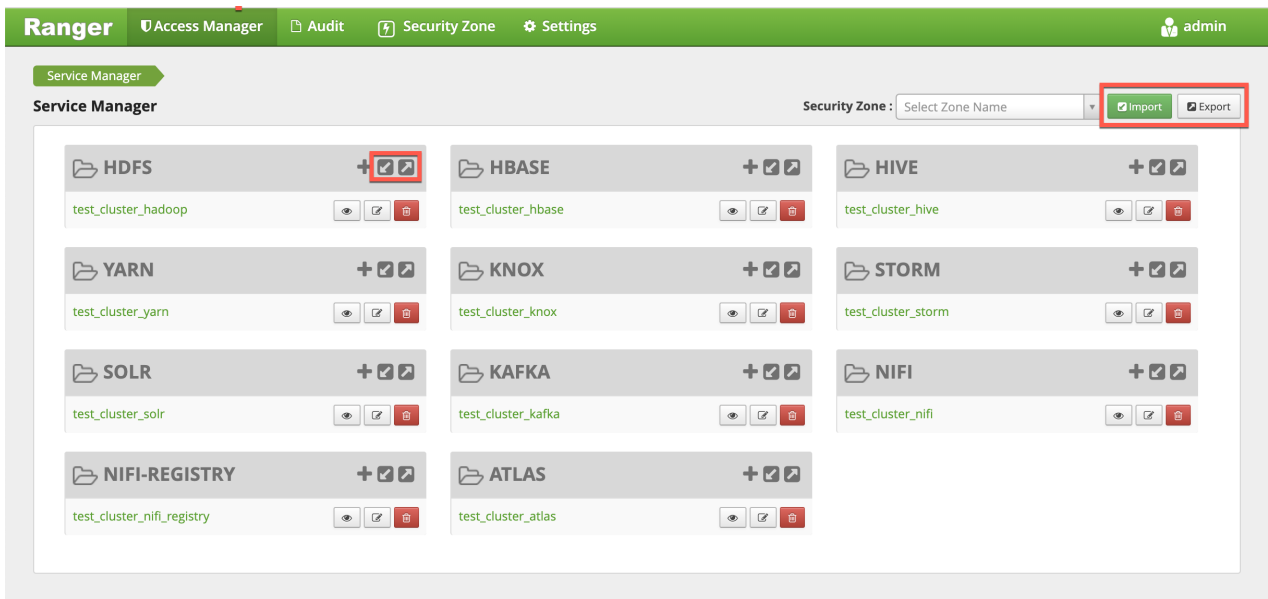
- `replaceTokens`: true if short-hand for user in resource-spec needs to be replaced at run-time with current-user's name; false if the resource-spec needs to be interpreted as it is. Default value: true.
- `tokenDelimiterStart`: Identifies start character of short-hand for current-user in resource specification. Default value: {.
- `tokenDelimiterEnd`: Identifies end character of short-hand for current-user in resource specification. Default value: }.
- `tokenDelimiterEscape`: Identifies escape character for escaping `tokenDelimiterStart` or `tokenDelimiterEnd` values in resource specification. Default value: \.
- `tokenDelimiterPrefix`: Identifies special prefix which together with string 'USER' makes up short-hand for current-user's name in the resource specification. Default value: .

Importing and Exporting Resource-Based Policies

You can export and import policies from the Ranger Admin UI for cluster resiliency (backups), during recovery operations, or when moving policies from test clusters to production clusters. You can export/import a specific subset of policies (such as those that pertain to specific resources or user/groups) or clone the entire repository (or multiple repositories) via Ranger Admin UI.

Interfaces

You can import and export policies from the Service Manager page:



You can also export policies from the Reports page:

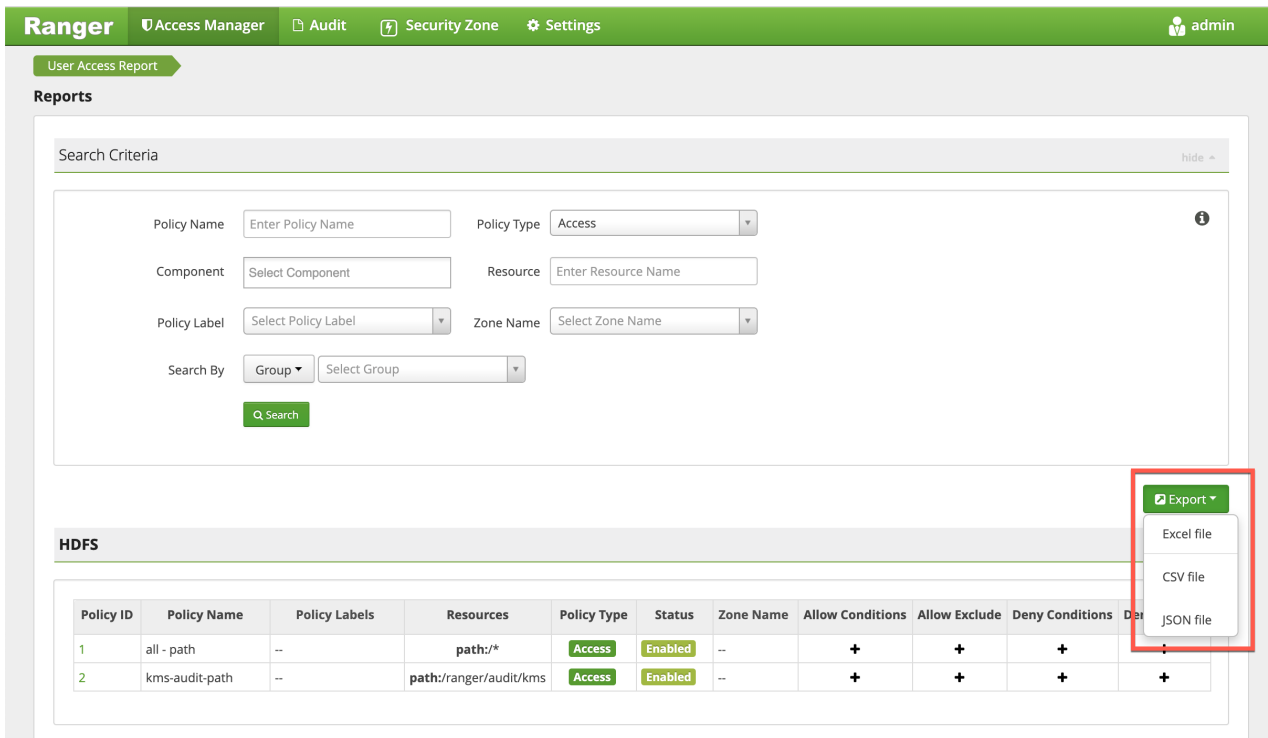


Table 46: Export Policy Options

| | Service Manager Page | Reports Page |
|-------------------------|----------------------|----------------------|
| Formats | JSON | JSON Excel CSV |
| Filtering Supported | No | Yes |
| Specific Service Export | Yes | Via filtering |

Filtering

When exporting from the Reports page, you can apply filters before saving the file.

Export Formats

You can export policies in the following formats:

- Excel
- JSON
- CSV

Note: CSV format is not supported for importing policies.

When you export policies from the Service Manager page, the policies are automatically downloaded in JSON format. If you wish to export in Excel or CSV format, export the policies from the Reports page dropdown menu.

Required User Roles

The Ranger admin user can import and export only Resource & Tag based policies. The credentials for this user are set in Ranger Configs > Advanced ranger-env in the fields labeled admin_username (default: admin/admin).

The Ranger KMS keyadmin user can import and export only KMS policies. The default credentials for this user are keyadmin/keyadmin.

Limitations

To successfully import policies, use the following database versions:

- MariaDB: 10.1.16+
- MySQL: 5.6.x+
- Oracle: 11gR2+
- PostgreSQL: 8.4+
- MS SQL: 2008 R2+

Partial import is not supported.

Related Information

[Configuring Resource-Based Policies](#)

[Importing and Exporting Tag-Based Policies](#)

Import Resource-Based Policies for a Specific Service

How to import the policies for a specific service (HBase, YARN, etc).

Procedure

1. On the Service Manager page, click the Import icon for the service:



The Import Policy page appears.

2. Select the file to import.

You can only import policies in JSON format.

Import Policy ✕

Select File :

Select file

Override Policy :

Ranger_Policies_20190717_190622.json ✕

i All services gets listed on service destination when Zone destination is blank. When zone is selected at destination, then only services associated with that zone will be listed.

Specify Zone Mapping :

| Source | To | Destination |
|--------|----|---|
| | To | No zone selected ▼ |

Specify Service Mapping :

| Source | To | Destination |
|--|----|--|
| cm_hdfs ✕ ▼ | To | Select service name ▼ ✕ |

Cancel

Import

3. (Optional) Configure the import operation:

- a) The Override Policy option deletes all policies of the destination repositories.
- b) Zone Mapping – when no destination is selected, all services are imported. When a destination is selected, only the services associated with that security zone are imported.
- c) Service Mapping maps the downloaded file repository, i.e. source repository to destination repository. You can use the red x symbols to remove services from the import. Scroll down to view all service mappings.

Import Policy ✕

Specify Zone Mapping :

| Source | | Destination |
|--------|----|--------------------|
| | To | No zone selected ▼ |

Specify Service Mapping :

| Source | | Destination |
|---|----|--|
| cm_hdfs ✕ ▼ | To | Select service name ▼ ✕ |
| cm_hbase ✕ ▼ | To | Select service name ▼ ✕ |
| cm_yarn ✕ ▼ | To | Select service name ▼ ✕ |
| cm_hive ✕ ▼ | To | Select service name ▼ ✕ |
| cm_knox ✕ ▼ | To | Select service name ▼ ✕ |
| cm_storm ✕ ▼ | To | Select service name ▼ ✕ |

Cancel
Import

4. Click **Import.**

A confirmation message appears after the file is imported.

Related Information

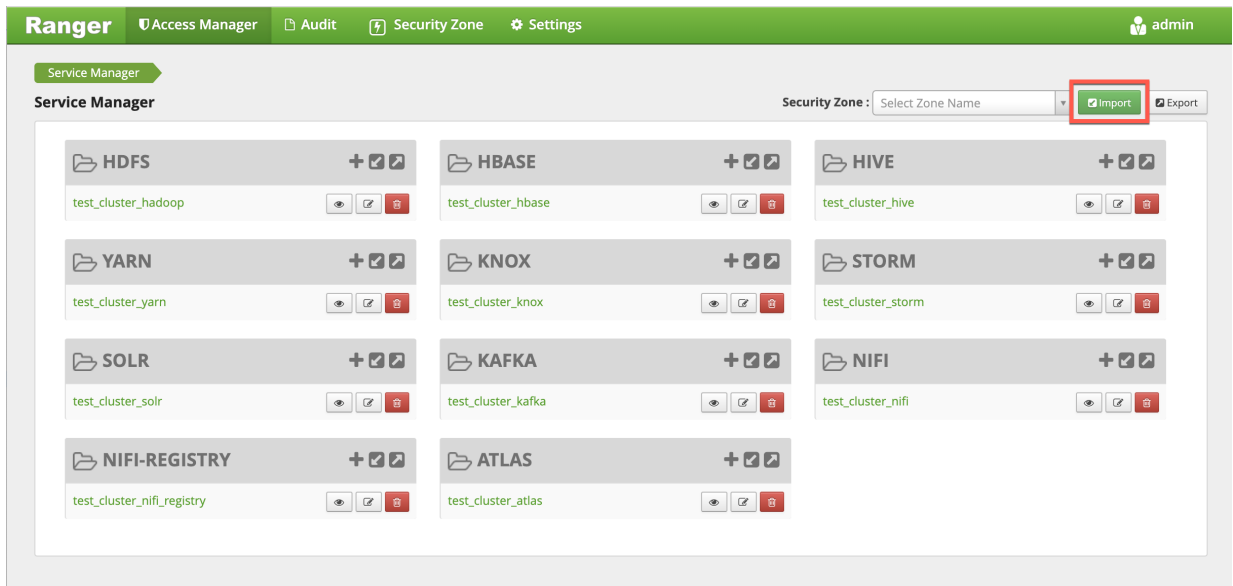
[Import Resource-Based Policies for All Services](#)

Import Resource-Based Policies for All Services

How to import policies for all services.

Procedure

1. On the Service Manager page, click **Import**.



The Import Policy page appears.

Import Policy ✕

Select File :

Select file

Override Policy :

Ranger_Policies_20190717_190622.json ✕

i All services gets listed on service destination when Zone destination is blank. When zone is selected at destination, then only services associated with that zone will be listed.

Specify Zone Mapping :

| Source | To | Destination | |
|--------|----|---|--|
| | To | No zone selected ▼ | |

Specify Service Mapping :

| Source | To | Destination | |
|--|----|--|---|
| cm_hdfs ✕ ▼ | To | Select service name ▼ | ✕ |

Cancel
Import

2. Select the file to import.
You can only import policies in JSON format.
3. (Optional) Configure the import operation:
 - a) The Override Policy option deletes all policies of the destination repositories.
 - b) Zone Mapping – when no destination is selected, all services are imported. When a destination is selected, only the services associated with that security zone are imported.
 - c) Service Mapping maps the downloaded file repository, i.e. source repository to destination repository. You can use the red x symbols to remove services from the import. Scroll down to view all service mappings.

Import Policy ✕

Specify Zone Mapping :

| Source | To | Destination |
|--------|----|---|
| | To | No zone selected ▼ |

Specify Service Mapping :

| Source | To | Destination |
|---|----|---|
| cm_hdfs ✕ ▼ | To | Select service name ▼ ✕ |
| cm_hbase ✕ ▼ | To | Select service name ▼ ✕ |
| cm_yarn ✕ ▼ | To | Select service name ▼ ✕ |
| cm_hive ✕ ▼ | To | Select service name ▼ ✕ |
| cm_knox ✕ ▼ | To | Select service name ▼ ✕ |
| cm_storm ✕ ▼ | To | Select service name ▼ ✕ |

Cancel
Import

4. Click **Import**.

A confirmation message appears after the file is imported.

Related Information

[Import Resource-Based Policies for a Specific Service](#)

Export Resource-Based Policies for a Specific Service

How to export the policies for a specific service (HBase, YARN, etc).

About this task

If you wish to export in Excel or CSV format, export the policies from the Reports page dropdown menu.

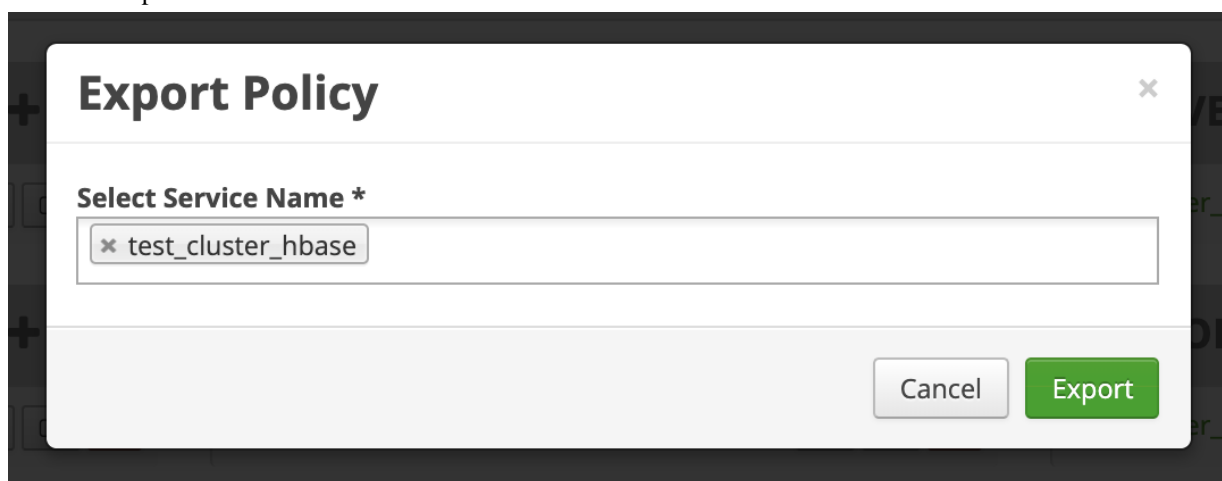
Procedure

1. On the Service Manager page, click the Export icon for the service:



The Export Policy page appears.

2. Click the Export button.



The file downloads in your browser as a JSON file.

Related Information

[Export All Resource-Based Policies for All Services](#)

Export All Resource-Based Policies for All Services

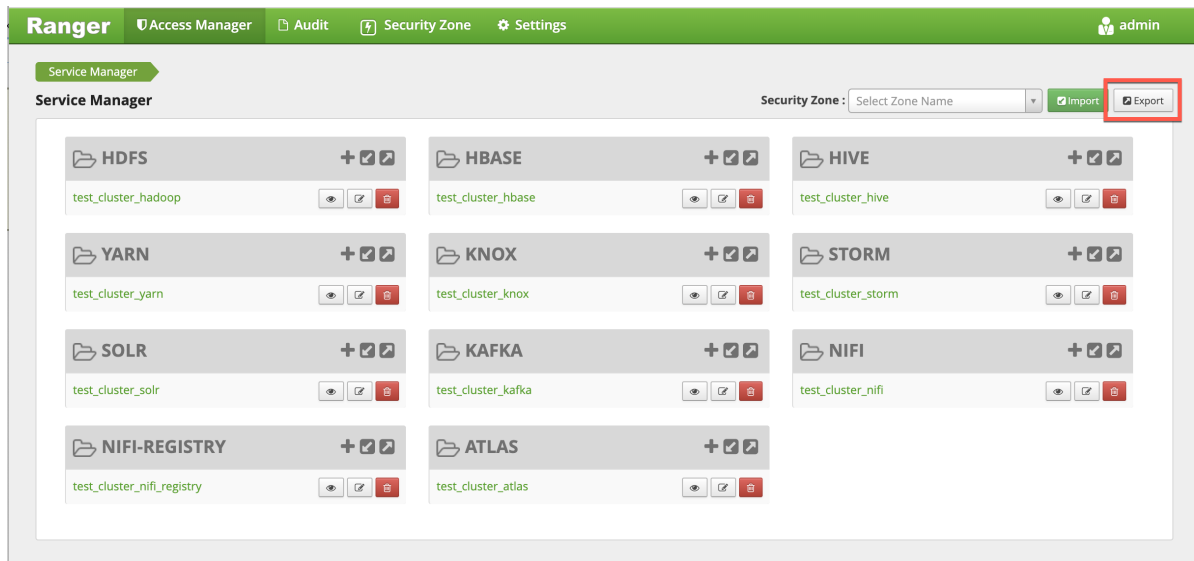
How to export the policies for all service.

About this task

If you would like to export in Excel or CSV format, export the policies from the Reports page dropdown menu.

Procedure

- From the Service Manager page:
 - a) Click Export:



The Export Policy page appears.

- b) Remove components or specific services, then click Export.



The file downloads in your browser as a JSON file.

- From the Reports page:
 - a) Apply filters before exporting the file.
 - b) Open the Export drop-down menu:

The screenshot shows the Ranger Admin interface for generating a User Access Report. The search criteria form is filled with the following values:

- Policy Name: Enter Policy Name
- Policy Type: Access
- Component: Select Component
- Resource: Enter Resource Name
- Policy Label: Select Policy Label
- Zone Name: Select Zone Name
- Search By: Group

The HDFS table contains the following data:

| Policy ID | Policy Name | Policy Labels | Resources | Policy Type | Status | Zone Name | Allow Conditions | Allow Exclude | Deny Conditions | Deny |
|-----------|----------------|---------------|-----------------------|-------------|---------|-----------|------------------|---------------|-----------------|------|
| 1 | all - path | -- | path:/* | Access | Enabled | -- | + | + | + | + |
| 2 | kms-audit-path | -- | path:ranger/audit/kms | Access | Enabled | -- | + | + | + | + |

- c) Select the file format.
The file downloads in your browser.

Related Information

[Export Resource-Based Policies for a Specific Service](#)

Row-level Filtering and Column Masking in Hive

You can use Apache Ranger row-level filters to set access policies for rows in Hive tables. You can also use Ranger column masking to set policies that mask data in Hive columns, for example to show only the first or last four characters of column data.

Row-level Filtering in Hive with Ranger Policies

Row-level filtering helps simplify Hive queries. By moving the access restriction logic down into the Hive layer, Hive applies the access restrictions every time data access is attempted. This helps simplify authoring of the Hive query, and provides seamless behind-the-scenes enforcement of row-level segmentation without having to add this logic to the predicate of the query.

About this task

Row-level filtering also improves the reliability and robustness of Hadoop. By providing row-level security to Hive tables and reducing the security surface area, Hive data access can be restricted to specific rows based on user characteristics (such as group membership) and the runtime context in which this request is issued.

Typical use cases where row-level filtering can be beneficial include:

- A hospital can create a security policy that allows doctors to view data rows only for their own patients, and that allows insurance claims administrators to view only specific rows for their specific site.
- A bank can create a policy to restrict access to rows of financial data based on the employee's business division, locale, or based on the employee's role (for example: only employees in the finance department are allowed to see customer invoices, payments, and accrual data; only European HR employees can see European employee data).

- A multi-tenant application can create logical separation of each tenant's data so that each tenant can see only their own data rows.

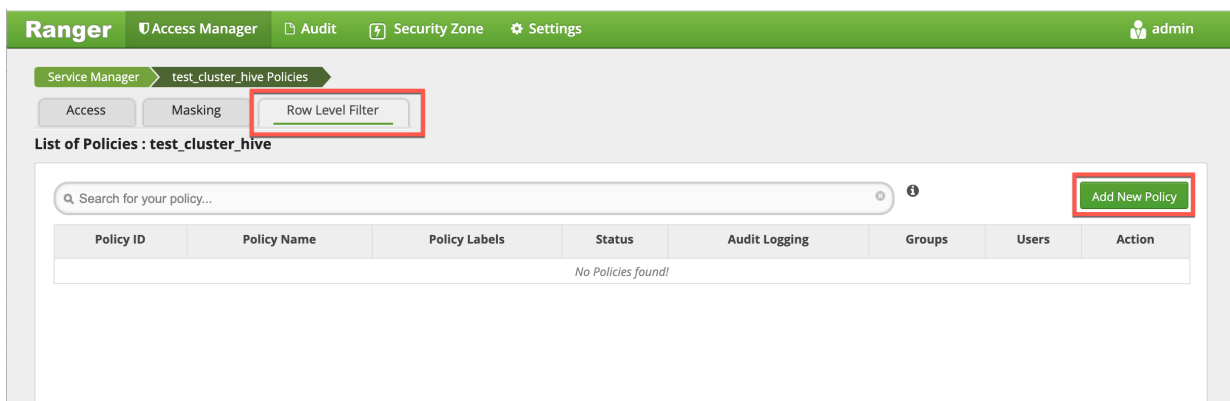
You can use Apache Ranger row-level filters to set access policies for rows in Hive tables. Row-level filter policies are similar to other Ranger access policies. You can set filters for specific users, groups, and conditions.

The following conditions apply when using row-level filters:

- The filter expression must be a valid WHERE clause for the table or view.
- Each table or view should have its own row-level filter policy.
- Wildcard matching is not supported on database or table names.
- Filters are evaluated in the order listed in the policy.
- An audit log entry is generated each time a row-level filter is applied to a table or view.

Procedure

1. On the Service Manager page, select an existing Hive Service.
2. Select the Row Level Filter tab, then click Add New Policy.



3. On the Create Policy page, add the following information for the row-level filter:

Table 47: Policy Details

| Field | Description |
|-----------------------------|---|
| Policy Name (required) | Enter an appropriate policy name. This name cannot be duplicated across the system. The policy is enabled by default. |
| normal/override | Enables you to specify an override policy. When override is selected, the access permissions in the policy override the access permissions in existing policies. This feature can be used with Add Validity Period to create temporary access policies that override existing policies. |
| Hive Database (required) | Type in the applicable database name. The auto-complete feature displays available databases based on the entered text. |
| Hive Table (required) | Type in the applicable table name. The auto-complete feature displays available tables based on the entered text. |
| Audit Logging | Audit Logging is set to Yes by default. Select No to turn off audit logging. |
| Description | Enter an optional description for the policy. |
| Add Validity Period | Specify a start and end time for the policy. |

Table 48: Row Filter Conditions

| Label | Description |
|----------------|---|
| Select Group | Specify the groups to which this policy applies. The public group contains all users, so granting access to the public group grants access to all users. |
| Select User | Specify one or more users to which this policy applies. |
| Access Types | Currently select is the only available access type. This will be used in conjunction with the WHERE clause specified in the Row Level Filter field. |
| Add Row Filter | <ul style="list-style-type: none"> To create a row filter for the specified users and groups, Click Add Row Filter, then type a valid WHERE clause in the Enter filter expression box. To allow Select access for the specified users and groups without row-level restrictions, do not add a row filter (leave the setting as "Add Row Filter"). Filters are evaluated in the order listed in the policy. The filter at the top of the Row Filter Conditions list is applied first, then the second, then the third, and so on. |

The screenshot displays the Ranger web interface for creating a policy. The 'Policy Details' section includes the following fields:

- Policy Type: Row Level Filter
- Policy Name: row-filter:hr.employee (with 'enabled' radio button selected)
- Policy Label: Policy Label
- Hive Database: hr
- Hive Table: employee
- Description: Row-level filter for hr.employee table
- Audit Logging: YES

The 'Row Filter Conditions' section shows a table with the following columns: Select Group, Select User, Access Types, and a filter expression. A modal dialog 'Enter filter expression' is open, showing a text input field and confirmation/cancel buttons. The table contains three rows of conditions, each with a dotted icon on the left for reordering.

- To move a condition in the Row Filter Conditions list (and therefore change the order in which it is evaluated), click the dotted rows icon at the left of the condition row, then drag the condition to a new position in the list.

Ranger Access Manager Audit Security Zone Settings admin

Service Manager test_cluster_hive Policies Create Policy

Create Policy

Policy Details :

Policy Type **Row Level Filter** Add Validity Period

Policy Name * row-filter:hr.employee **enabled** normal

Policy Label Policy Label

Hive Database * hr

Hive Table * employee

Description Row-level filter for hr.employee table

Audit Logging **YES**

Row Filter Conditions : hide

| Select Group | Select User | Access Types | Row Level Filter | |
|--------------|-------------|--------------|------------------|---|
| Select Group | admin | select | loc_state = 'CA' | ✖ |
| Select Group | ambari-qa | select | loc_state = 'CA' | ✖ |
| public | Select User | select | loc_state = 'CA' | ✖ |

Add **Cancel**

5. Click **Add** to add the new row-level filter policy.

Dynamic Resource-Based Column Masking in Hive with Ranger Policies

You can use Apache Ranger dynamic resource-based column masking capabilities to protect sensitive data in Hive in near real-time. You can set policies that mask or anonymize sensitive data columns (such as PII, PCI, and PHI) dynamically from Hive query output. For example, you can mask sensitive data within a column to show only the first or last four characters.

About this task

Dynamic column masking policies are similar to other Ranger access policies for Hive. You can set filters for specific users, groups, and conditions. With dynamic column-level masking, sensitive information never leaves Hive, and no changes are required at the consuming application or the Hive layer. There is also no need to produce additional protected duplicate versions of datasets.

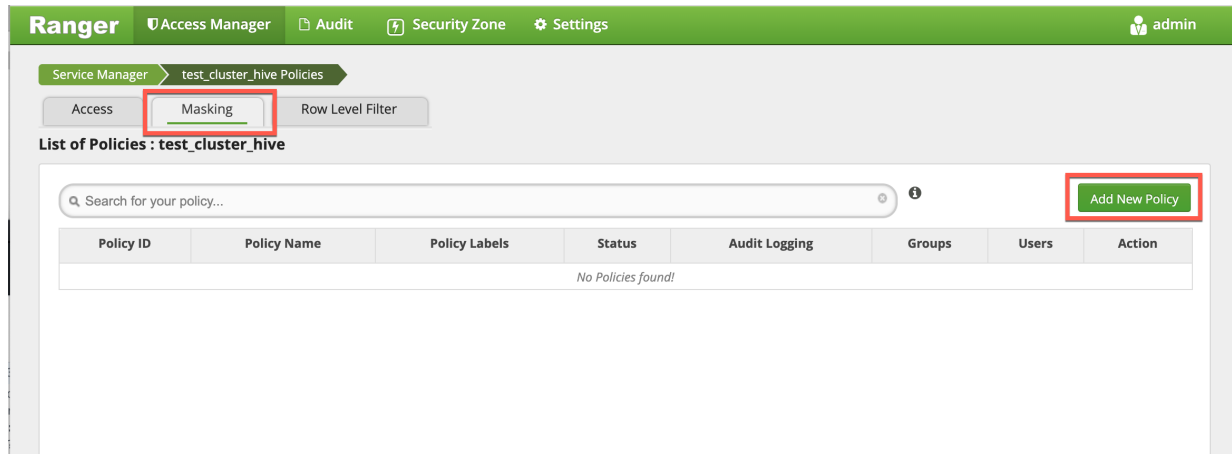
The following conditions apply when using Ranger column masking policies to mask data returned in Hive query results:

- A variety of masking types are available, such as show last 4 characters, show first 4 characters, Hash, Nullify, and date masks (show only year).
- You can specify a masking type for specific users, groups, and conditions.
- Wildcard matching is not supported.
- Each column should have its own masking policy.

- Masks are evaluated in the order listed in the policy.
- An audit log entry is generated each time a masking policy is applied to a column.

Procedure

1. On the Service Manager page, select an existing Hive Service.
2. Select the Masking tab, then click Add New Policy.



3. On the Create Policy page, add the following information for the column-masking filter:

Table 49: Policy Details

| Field | Description |
|-----------------------------|---|
| Policy Name (required) | Enter an appropriate policy name. This name cannot be duplicated across the system. The policy is enabled by default. |
| normal/override | Enables you to specify an override policy. When override is selected, the access permissions in the policy override the access permissions in existing policies. This feature can be used with Add Validity Period to create temporary access policies that override existing policies. |
| Hive Database (required) | Type in the applicable database name. The auto-complete feature displays available databases based on the entered text. |
| Hive Table (required) | Type in the applicable table name. The auto-complete feature displays available tables based on the entered text. |
| Hive Column (required) | Type in the applicable column name. The auto-complete feature displays available columns based on the entered text. |
| Audit Logging | Audit Logging is set to Yes by default. Select No to turn off audit logging. |
| Description | Enter an optional description for the policy. |
| Add Validity Period | Specify a start and end time for the policy. |

Table 50: Mask Conditions

| Label | Description |
|--------------|---|
| Select Group | Specify the groups to which this policy applies. The public group contains all users, so granting access to the public group grants access to all users. |

| Label | Description |
|---------------------|--|
| Select User | Specify one or more users to which this policy applies. |
| Access Types | Currently select is the only available access type. |
| Select Masking Type | <p>To create a row filter for the specified users and groups, click Select Masking Option, then select a masking type:</p> <ul style="list-style-type: none"> • Redact – mask all alphabetic characters with "x" and all numeric characters with "n". • Partial mask: show last 4 – Show only the last four characters. • Partial mask: show first 4 – Show only the first four characters. • Hash – Replace all characters with a hash of entire cell value. • Nullify – Replace all characters with a NULL value. • Unmasked (retain original value) – No masking is applied. • Date: show only year – Show only the year portion of a date string and default the month and day to 01/01 • Custom – Specify a custom masked value or expression. Custom masking can use any valid Hive UDF (Hive that returns the same data type as the data type in the column being masked). <p>Masking conditions are evaluated in the order listed in the policy. The condition at the top of the Masking Conditions list is applied first, then the second, then the third, and so on.</p> |

Ranger Access Manager Audit Security Zone Settings admin

Service Manager test_cluster_hive Policies Create Policy

Create Policy

Policy Details :

Policy Type: **Masking** Add Validity Period

Policy Name *: mask.hr.employee.ssn enabled normal

Policy Label: Policy Label

Hive Database *: hr

Hive Table *: employee

Hive Column *: ssn

Description: Masking for ssn column in hr.employee table

Audit Logging: **YES**

Mask Conditions :

| Select Group | Select User | Access Types | Masking Option | |
|--------------|-------------|--------------|----------------------------------|---|
| Select Group | admin | select | Unmasked (retain original value) | ✖ |
| Select Group | ambari-qa | select | Partial mask: show last 4 | ✖ |
| public | Select User | select | Nullify | ✖ |

hide

Add Cancel

- To move a condition in the Mask Conditions list (and therefore change the order in which it is evaluated), click the dotted rows icon at the left of the condition row, then drag the condition to a new position in the list.

The screenshot shows the Ranger web interface for creating a policy. The 'Policy Details' section is filled out with the following information:

- Policy Type: Masking
- Policy Name: mask.hr.employee.ssn
- Policy Label: Policy Label
- Hive Database: hr
- Hive Table: employee
- Hive Column: ssn
- Description: Masking for ssn column in hr.employee table
- Audit Logging: YES

The 'Mask Conditions' section contains a table with the following data:

| Select Group | Select User | Access Types | Masking Options |
|--------------|-------------|--------------|----------------------------------|
| Select Group | admin | select | Unmasked (retain original value) |
| Select Group | ambari-qa | select | Partial mask: show last 4 |
| + public | Select User | select | Nullify |

A modal window titled 'Select Masking Option' is open, showing the following options:

- Redact
- Partial mask: show last 4
- Partial mask: show first 4
- Hash
- Nullify
- Unmasked (retain original value)
- Date: show only year
- Custom

The 'Add' button is highlighted in red.

5. Click **Add** to add the new column masking filter policy.

Dynamic Tag-Based Column Masking in Hive with Ranger Policies

Where Ranger resource-based masking policy for Hive anonymizes data from a Hive column identified by the database, table, and column, tag-based masking policy anonymizes Hive column data based on tags and tag attribute values associated with Hive column (usually specified as metadata classification in Atlas).

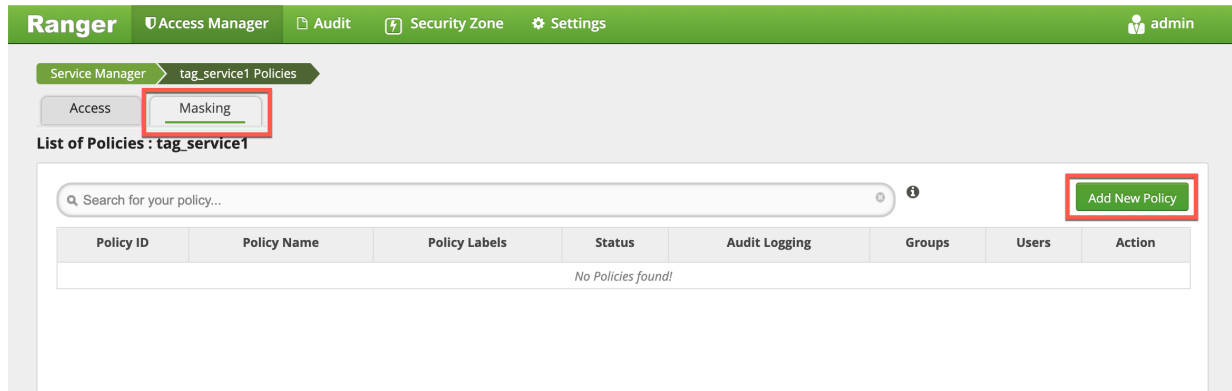
About this task

The following conditions apply when using Ranger column masking policies to mask data returned in Hive query results:

- A variety of masking types are available, such as show last 4 characters, show first 4 characters, Hash, Nullify, and date masks (show only year).
- You can specify a masking type for specific users, groups, and conditions.
- Wildcard matching is not supported.
- If there are multiple tag masking policies applied to the same Hive column, the masking policy with the lexicographically smallest policy-name is chosen for enforcement, E.G., policy "a" is enforced before policy "aa".
- Masks are evaluated in the order listed in the policy.
- An audit log entry is generated each time a masking policy is applied to a column.

Procedure

1. Select Access Manager > Tag Based Policies, then select a tag-based service.
2. Select the **Masking** tab, then click **Add New Policy**.



3. On the **Create Policy** page, add the following information for the column-masking filter:

Table 51: Policy Details

| Field | Description |
|---|---|
| Policy Type (required) | Set to Masking by default. |
| Policy Name (required) | Enter an appropriate policy name. This name cannot be duplicated across the system. The policy is enabled by default. |
| normal/override | Enables you to specify an override policy. When override is selected, the access permissions in the policy override the access permissions in existing policies. This feature can be used with Add Validity Period to create temporary access policies that override existing policies. |
| TAG (required) | Enter the applicable tag name, E.G., MASK. |
| Audit Logging | Audit Logging is set to Yes by default. Select No to turn off audit logging. |
| Description | Enter an optional description for the policy. |
| Add Validity Period | Specify a start and end time for the policy. |
| Policy Conditions (applied at the policy level) | <p>Click the + icon to add policy conditions. Currently "Accessed after expiry_date? (yes/no)" is the only available policy condition.</p> <p>"Accessed after expiry_date (yes/no)?": To set this condition, type yes in the text box, then click the check mark button to add the condition.</p> <p>Enter boolean expression: Available for allow or deny conditions on tag-based policies. For examples and details, see "Using Tag Attributes and Values in Ranger Tag-Based Policy Conditions".</p> <p>Click Save to save the policy condition.</p> |

Table 52: Mask Conditions

| Label | Description |
|---|--|
| Select Group | Specify the groups to which this policy applies. The public group contains all users, so granting access to the public group grants access to all users. |
| Select User | Specify one or more users to which this policy applies. |
| Policy Conditions (applied at the item level) | Click Add Conditions to add policy conditions. Currently "Accessed after expiry_date? (yes/no)" is the only available policy condition. "Accessed after expiry_date (yes/no)?:": To set this condition, type yes in the text box, then click the check mark button to add the condition. Enter boolean expression: Available for allow or deny conditions on tag-based policies. For examples and details, see "Using Tag Attributes and Values in Ranger Tag-Based Policy Conditions". |
| Access Types | Currently select is the only available access type for the hive component. |
| Select Masking Option | To create a row filter for the specified users and groups, click Select Masking Option, then select a masking type: <ul style="list-style-type: none"> • Redact – mask all alphabetic characters with "x" and all numeric characters with "n". • Partial mask: show last 4 – Show only the last four characters. • Partial mask: show first 4 – Show only the first four characters. • Hash – Replace all characters with a hash of entire cell value. • Nullify – Replace all characters with a NULL value. • Unmasked (retain original value) – No masking is applied. • Date: show only year – Show only the year portion of a date string and default the month and day to 01/01 • Custom – Specify a custom masked value or expression. Custom masking can use any valid Hive UDF (Hive that returns the same data type as the data type in the column being masked). Masking conditions are evaluated in the order listed in the policy. The condition at the top of the Masking Conditions list is applied first, then the second, then the third, and so on. |

Ranger Access Manager Audit Security Zone Settings admin

Service Manager tag_service1 Policies Create Policy

Create Policy

Policy Details :

Policy Type **Masking** Add Validity Period

Policy Name * enabled normal

Policy Label

TAG *

Description

Audit Logging **YES**

Policy Conditions +

No Conditions

Mask Conditions :

| Select Group | Select User | Policy Conditions | Access Types |
|---|-----------------------------------|--|----------------------------------|
| <input type="text" value="Select Group"/> | <input type="text" value="hive"/> | Add Conditions + | HIVE + |

Select Masking Option

- Redact
- Partial mask: show last 4
- Partial mask: show first 4
- Hash
- Nullify
- Unmasked (retain original value)
- Date: show only year
- Custom

HIVE: Unmasked (retain original value) ✕

Add Cancel

- You can use the Plus (+) symbols to add additional conditions. Conditions are evaluated in the order listed in the policy. The condition at the top of the list is applied first, then the second, then the third, and so on.
- Click Add to add the new policy.

Tag-Based Services and Policies

Ranger enables you to create tag-based services and add access policies to those services.

Adding a Tag-based Service

How to add a tag-based service to Ranger.

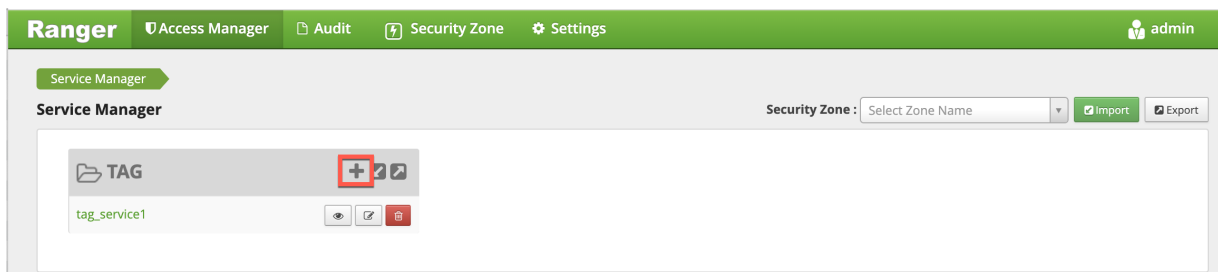
About this task

You can use the Service Manager for Tag-Based Policies page to create tag-based services and add tag-based access policies that can be applied to Hadoop resources. Using tag-based policies enables you to control access to resources across multiple Hadoop components without creating separate services and policies in each component. You can also use Ranger TagSync to synchronize the Ranger tag store with an external metadata service such as Apache Atlas.

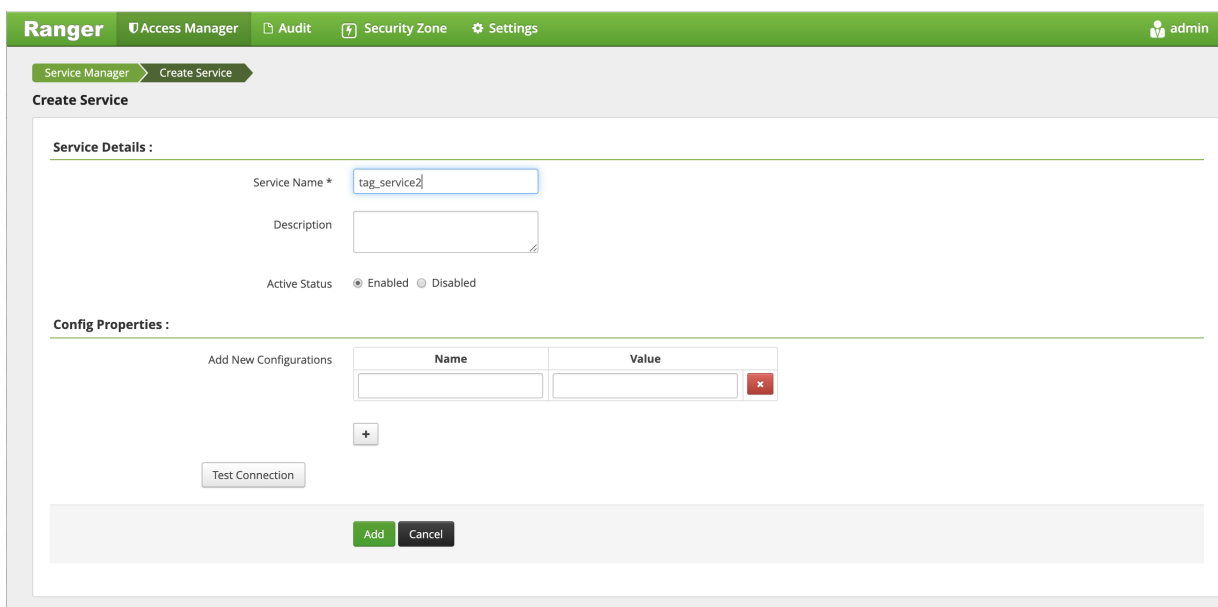
Procedure

1. Select Access Manager > Tag Based Policies, then click the Add icon

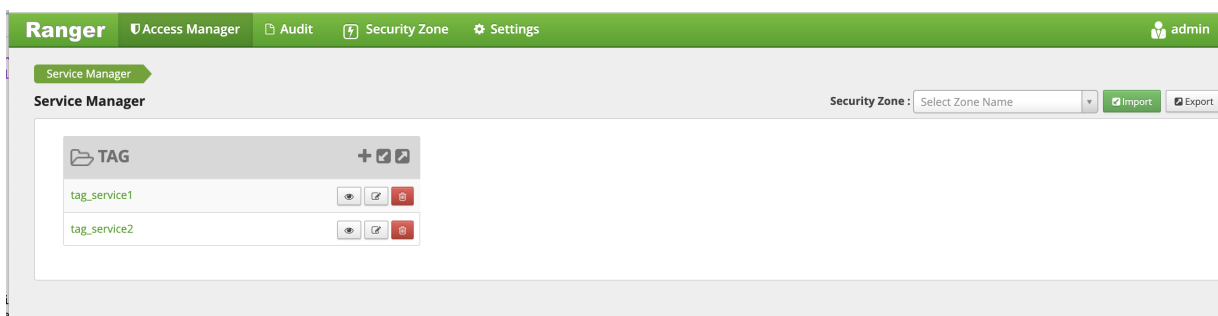
()
in the TAG box on the Service Manager page.



2. On the Create Service page, type in a service name and an optional description. The service is enabled by default, but you can disable it by selecting Disabled. To add the service, click **Add**.



3. The new tag service appears on the Service Manager page.

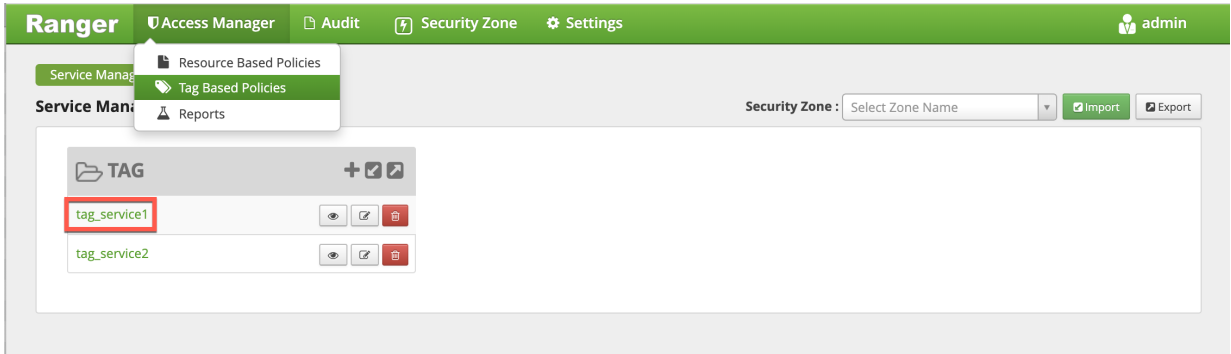


Adding Tag-Based Policies

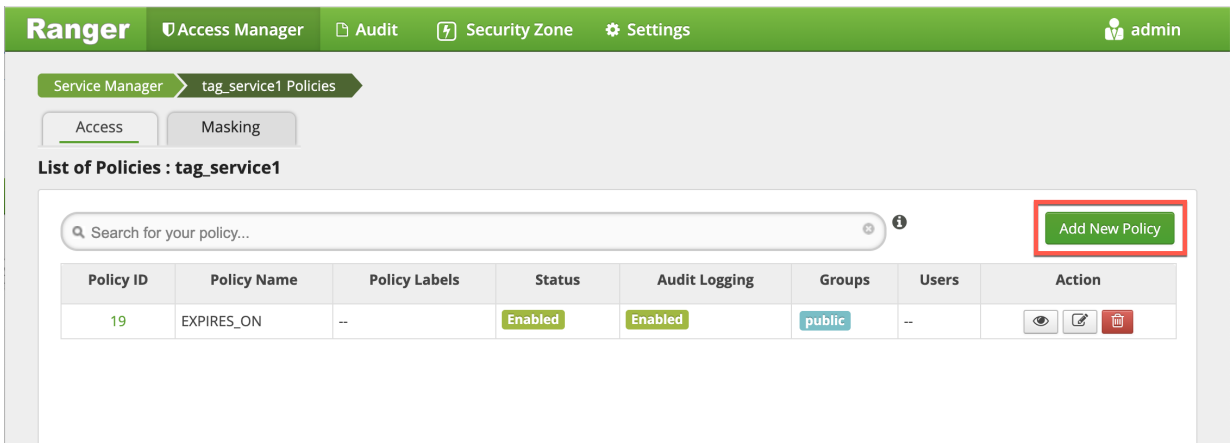
Tag-based policies enable you to control access to resources across multiple Hadoop components without creating separate services and policies in each component. You can also use Ranger TagSync to synchronize the Ranger tag store with an external metadata service such as Apache Atlas.

Procedure

1. Select Access Manager> Tag Based Policies, then select a tag-based service.



2. On the List of Policies page, click Add New Policy.



The Create Policy page appears:

3. Enter information on the Create Policy page as follows:

Table 53: Policy Details

| Field | Description |
|-----------------|---|
| Policy Type | Set to Access by default. |
| Policy Name | Enter a unique policy name. This name cannot be duplicated across the system. This field is mandatory. |
| normal/override | Enables you to specify an override policy. When override is selected, the access permissions in the policy override the access permissions in existing policies. This feature can be used with Add Validity Period to create temporary access policies that override existing policies. |
| TAG | Enter the applicable tag name. |
| Description | (Optional) Describe the purpose of the policy. |

| Field | Description |
|---|---|
| Audit Logging | Specify whether this policy is audited. (De-select to disable auditing). |
| Policy Label | Specify a label for this policy. You can search reports and filter policies based on these labels. |
| Add Validity Period | Specify a start and end time for the policy. |
| Policy Conditions (applied at the policy level) | <p>Click the + icon to add policy conditions. Currently "Accessed after expiry_date? (yes/no)" is the only available policy condition.</p> <p>"Accessed after expiry_date (yes/no)?": To set this condition, type yes in the text box, then click the check mark button to add the condition.</p> <p>Enter boolean expression: Available for allow or deny conditions on tag-based policies. For examples and details, see "Using Tag Attributes and Values in Ranger Tag-Based Policy Conditions".</p> <p>Click Save to save the policy condition.</p> |

Table 54: Allow, Exclude from Allow, Deny, and Exclude from Deny Conditions

| Label | Description |
|---|---|
| Select Group | <p>Specify the group to which this policy applies. To designate the group as an Administrator for the chosen resource, specify Admin permissions. (Administrators can create child policies based on existing policies).</p> <p>The public group contains all users, so setting a condition for the public group applies to all users.</p> |
| Select User | Specify a particular user to which this policy applies (outside of an already-specified group) OR designate a particular user as Admin for this policy. (Administrators can create child policies based on existing policies). |
| Policy Conditions (applied at the item level) | <p>Click Add Conditions to add policy conditions. Currently "Accessed after expiry_date? (yes/no)" is the only available policy condition.</p> <p>"Accessed after expiry_date (yes/no)?": To set this condition, type yes in the text box, then click the check mark button to add the condition.</p> <p>Enter boolean expression: Available for allow or deny conditions on tag-based policies. For examples and details, see "Using Tag Attributes and Values in Ranger Tag-Based Policy Conditions".</p> |
| Component Permissions | Click Add Permissions to add or edit component conditions. To add component permissions, enter the component name in the text box, then use the check boxes to specify component permissions. Click the check mark button to add the chosen component conditions to the policy. |

- You can use the Plus (+) symbols to add additional conditions. Conditions are evaluated in the order listed in the policy. The condition at the top of the list is applied first, then the second, then the third, and so on.
- Click Add to add the new policy.

Related Information

[Using Tag Attributes and Values in Ranger Tag-Based Policy Conditions](#)

[Using Basic and Advanced Search](#)

Using Tag Attributes and Values in Ranger Tag-Based Policy Conditions

Enter boolean expression allows Ranger to use tag attributes and values when configuring tag-based policy Allow or Deny conditions. It allows admins to provide boolean expression(s) using tag attributes.

The policy condition is introduced in the tag service definition:

```
{
  "itemId":2,
  "name":"expression",
  "evaluator":
  "org.apache.ranger.plugin.conditionevaluator.RangerScriptConditionEvaluator",
  "evaluatorOptions" : {"engineName":"JavaScript",
  "ui.isMultiline":"true"},
  "label":"Enter boolean expression",
  "description": "Boolean expression"
}
```

The following variables can be referenced in the boolean expression:

- ctx: Context handler containing APIs to access metadata information from the request.
- tag: Information about the current tag.
- tagAttr: Map containing all the current tag attributes and corresponding values.

The following APIs available from the request:

- getUser(): Returns a string.
- getUserGroups(): Returns a set of strings containing groups.
- getClientIPAddress(): Returns a string containing client IP address.
- getAction(): Returns a string containing information about the action being requested.

For two scenarios:

- User “sam” needs to be denied a policy based on the IP address of the machine from where the resources are accessed.

Set the deny condition for user sam with the following boolean expression:

```
if ( tagAttr.get('ipAddr').equals(ctx.getClientIPAddress()) ) {
  ctx.result = true;
}
```

- Deny one particular user, “bob” from a group, “users”, only when this user is accessing resources from a particular IP defined as an tag attribute in Atlas.

Set the deny condition for group users with the following boolean expression:

```
if (tagAttr.get('ipAddr').equals(ctx.getClientIPAddress()) &&
  ctx.getUser().equals("bob")) {
  ctx.result=true;
}
```

Deny Conditions :

| Select Group | Select User | Policy Conditions | Component Permissions |
|-----------------------------------|-------------------------|----------------------------------|-----------------------|
| Select Group [X] users [X] bob | [X] sara Select User | expression: JavaScript Condition | deny [X] |
| | | expression: JavaScript Condition | deny [X] |

JavaScript Expression 1: `(tagAttr.get('ipAddr').equals(ctx.getClientIPAddress())) { ctx.result = true;}`

JavaScript Expression 2: `(tagAttr.get('ipAddr').equals(ctx.getClientIPAddress()) && ctx.getUser().equals('bob')) { ctx.result=true;}`

Adding a Tag-Based PII Policy

Example of how to add a PII tag-based policy. In this example we create a tag-based policy for objects tagged "PII" in Atlas. Access to objects tagged "PII" is allowed for members of the "audit" group. All other users (the "public" group) are denied access.

Procedure

1. Select Access Manager > Tag Based Policies, then select a tag-based service.

Ranger Access Manager

Service Manager > TAG

| | |
|--------------|----------------------|
| tag_service1 | [eye] [edit] [trash] |
| tag_service2 | [eye] [edit] [trash] |

2. On the List of Policies page, click Add New Policy.

Ranger Access Manager

Service Manager > tag_service1 Policies

Access Masking

List of Policies : tag_service1

Search for your policy...

Add New Policy

| Policy ID | Policy Name | Policy Labels | Status | Audit Logging | Groups | Users | Action |
|-----------|-------------|---------------|---------|---------------|--------|-------|----------------------|
| 19 | EXPIRES_ON | -- | Enabled | Enabled | public | -- | [eye] [edit] [trash] |

The Create Policy page appears:

Policy Details :

Policy Type **Access** Add Validity Period

Policy Name * enabled normal

Policy Label

TAG *

Description

Audit Logging **YES**

Policy Conditions +

No Conditions

Allow Conditions : hide ^

| Select Group | Select User | Policy Conditions | Component Permissions | |
|---|--|-------------------------------|--------------------------------|----------------|
| <input type="text" value="Select Group"/> | <input type="text" value="Select User"/> | Add Conditions + | Add Permissions + | × |

+

⚠ Exclude from Allow Conditions : show v

Deny Conditions : hide ^

| Select Group | Select User | Policy Conditions | Component Permissions | |
|---|--|-------------------------------|--------------------------------|----------------|
| <input type="text" value="Select Group"/> | <input type="text" value="Select User"/> | Add Conditions + | Add Permissions + | × |

3. Enter the following information on the Create Policy page:

Table 55: Policy Details

| Field | Description |
|---------------|--|
| Policy Type | Set to Access by default. |
| Policy Name | PII |
| TAG | PII |
| Audit Logging | YES |
| Description | Restrict access to resources with the PII tag. |

Table 56: Allow Conditions

| Label | Description |
|-----------------------|----------------------------------|
| Select Group | audit |
| Select User | <none> |
| Policy Conditions | <none> |
| Component Permissions | hive (select all permissions) |

Table 57: Deny Conditions

| Label | Description |
|-----------------------|----------------------------------|
| Select Group | public |
| Select User | <none> |
| Policy Conditions | <none> |
| Component Permissions | hive (select all permissions) |

Table 58: Exclude from Deny Conditions

| Label | Description |
|-----------------------|----------------------------------|
| Select Group | audit |
| Select User | <none> |
| Policy Conditions | <none> |
| Component Permissions | hive (select all permissions) |

Ranger
Access Manager
Audit
Security Zone
Settings
admin

Service Manager > tag_service1 Policies > Create Policy

Create Policy

Policy Details :

Policy Type: Access Add Validity Period

Policy Name *: enabled normal

Policy Label:

TAG *:

Description:

Audit Logging: YES

Policy Conditions +

No Conditions

Allow Conditions : hide ^

| Select Group | Select User | Policy Conditions | Component Permissions |
|--------------------------------------|--|--|---|
| <input type="text" value="x audit"/> | <input type="text" value="Select User"/> | Add Conditions + | HIVE ✎ ✕ |

⚠ Exclude from Allow Conditions : show v

Deny Conditions : hide ^

| Select Group | Select User | Policy Conditions | Component Permissions |
|---------------------------------------|--|--|---|
| <input type="text" value="x public"/> | <input type="text" value="Select User"/> | Add Conditions + | HIVE ✎ ✕ |

⚠ Exclude from Deny Conditions : hide ^

| Select Group | Select User | Policy Conditions | Component Permissions |
|--------------------------------------|--|--|---|
| <input type="text" value="x audit"/> | <input type="text" value="Select User"/> | Add Conditions + | HIVE ✎ ✕ |

In this example we used Allow Conditions to grant access to the "audit" group, and then used Deny Conditions to deny access to the "public" group. Because the "public" group includes all users, we then used Exclude from Deny Conditions to exclude the "audit" group, in effect reinstating the "audit" group's original Allow access condition.

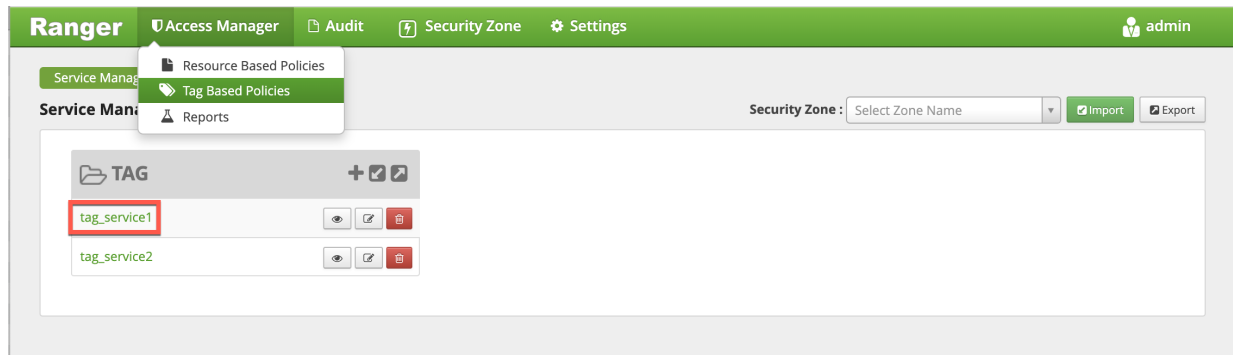
4. Click Add to add the new policy.

Default EXPIRES ON Tag Policy

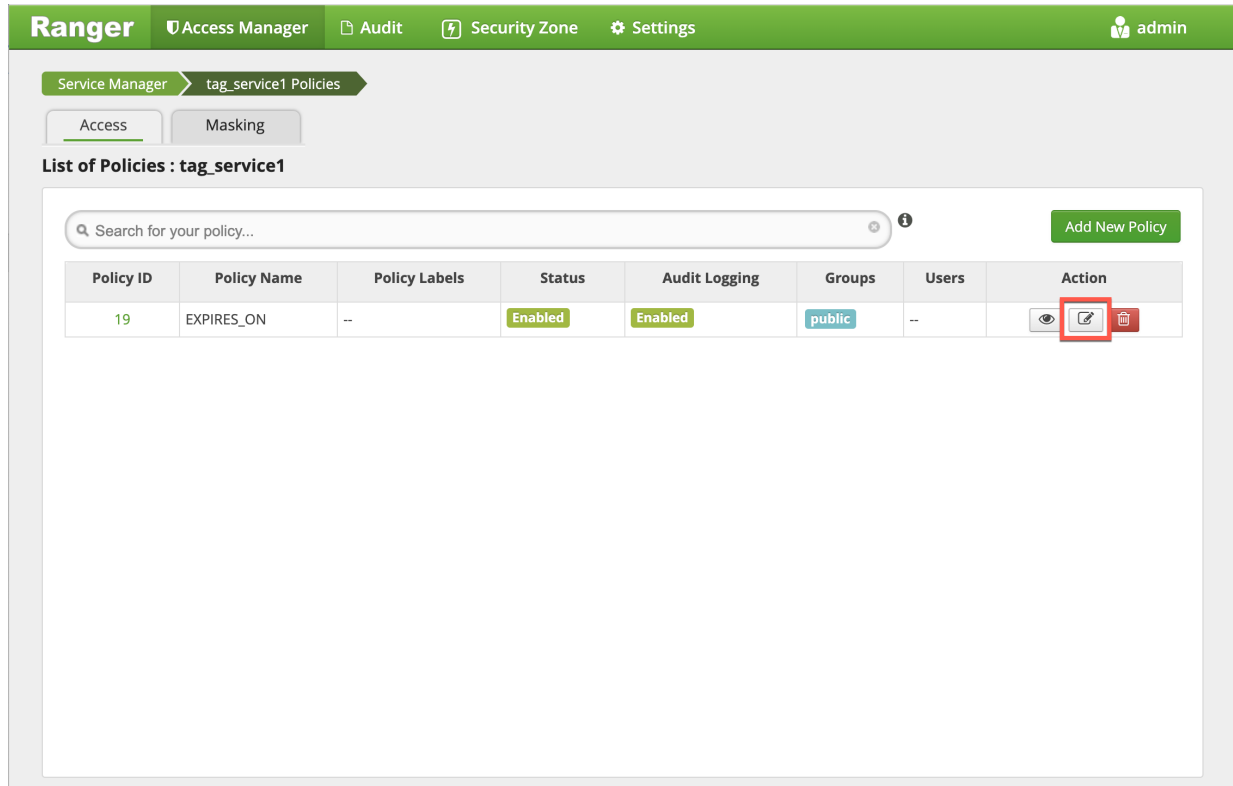
An EXPIRES_ON tag-based policy is created automatically when a tag service instance created. This default policy denies access to objects tagged with EXPIRES_ON after the expiry date specified in the Atlas tag attribute. You can use the following steps to review the default EXPIRES_ON policy.

Procedure

1. Select Access Manager > Tag Based Policies, then select a tag-based service.



2. On the List of Policies page, click the Edit icon for the default EXPIRES_ON policy.



The Edit Policy page appears:

Ranger Access Manager Audit Security Zone Settings admin

Service Manager > tag_service1 Policies > Edit Policy

Edit Policy

Policy Details :

Policy Type: **Access** + Add Validity Period

Policy ID: **19**

Policy Name *: EXPIRES_ON enabled normal

Policy Label:

TAG *:

Description:

Audit Logging: **YES**

Policy Conditions +

No Conditions

Allow Conditions : hide ^

| Select Group | Select User | Policy Conditions | Component Permissions |
|---|--|----------------------------------|--|
| <input type="text" value="Select Group"/> | <input type="text" value="Select User"/> | Add Conditions + | Add Permissions + ✖ |

⚠ Exclude from Allow Conditions : show v

Deny Conditions : hide ^

| Select Group | Select User | Policy Conditions | Component Permissions |
|-------------------------------------|--|----------------------------|---|
| <input type="text" value="public"/> | <input type="text" value="Select User"/> | accessed-after-expiry: yes | HDFS HBASE HIVE YARN KNOX STORM KMS SOLR KAFKA NIFI NIFI-REGISTRY ATLAS ✖ |

⚠ Exclude from Deny Conditions : hide ^

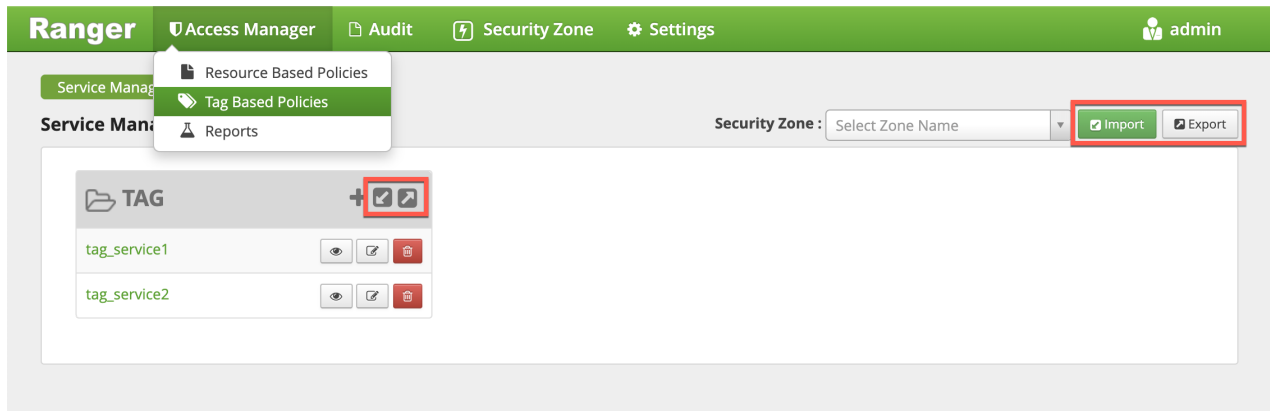
- We can see that the default EXPIRES_ON policy denies access to all users, and for all components, after the expiry date specified in the Atlas tag attribute.

Importing and Exporting Tag-Based Policies

You can export and import policies from the Ranger Admin UI for cluster resiliency (backups), during recovery operations, or when moving policies from test clusters to production clusters. You can import or export a specific subset of policies (such as those that pertain to specific resources or user/groups) or clone the entire repository (or multiple repositories) via the Ranger Admin UI.

Interfaces

You can import and export policies from the Tag Based Policies page:



You can also export policies from the Reports page:

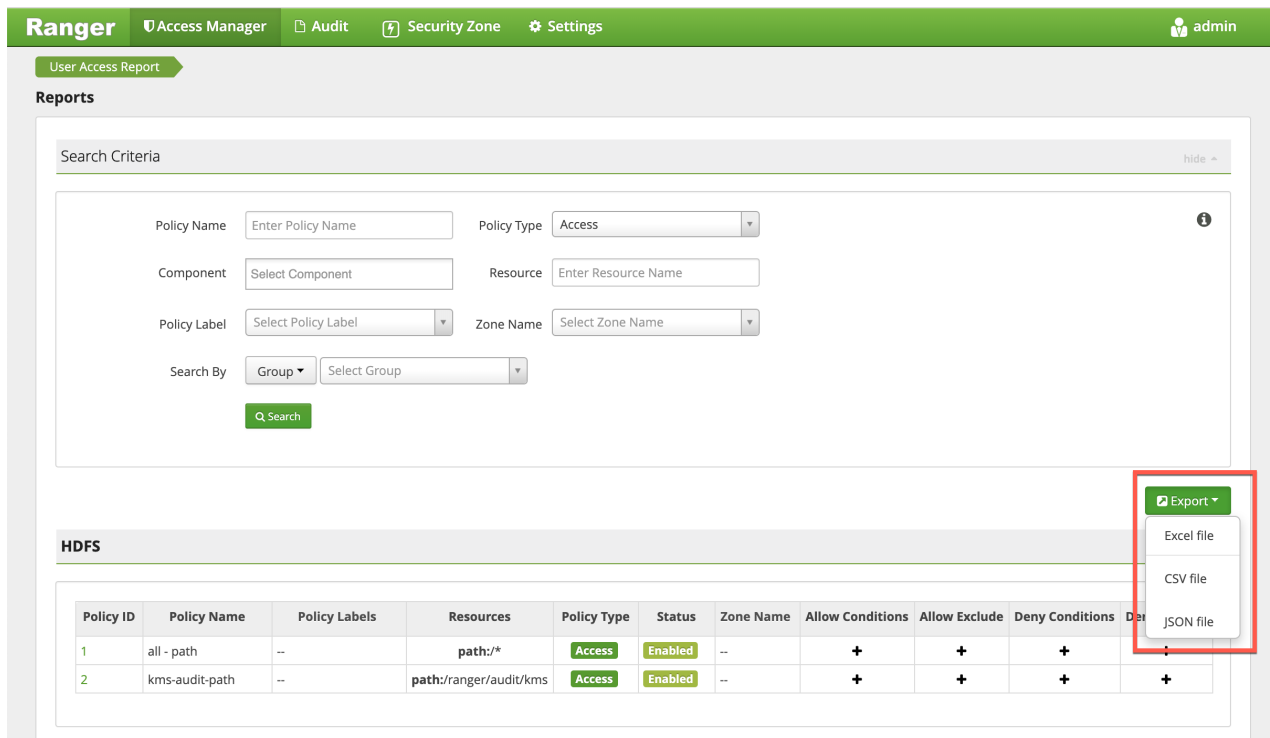


Table 59: Export Policy Options

| | Service Manager Page | Reports Page |
|-------------------------|----------------------|----------------------|
| Formats | JSON | JSON Excel CSV |
| Filtering Supported | No | Yes |
| Specific Service Export | Yes | Via filtering |

Filtering

When exporting from the Reports page, you can apply filters before saving the file.

Export Formats

You can export policies in the following formats:

- Excel

- JSON
- CSV

Note: CSV format is not supported for importing policies.

When you export policies from the Service Manager page, the policies are automatically downloaded in JSON format. If you wish to export in Excel or CSV format, export the policies from the Reports page dropdown menu.

Required User Roles

The Ranger admin user can import and export only Resource & Tag based policies. The credentials for this user are set in Ranger Configs > Advanced ranger-env in the fields labeled admin_username (default: admin/admin).

The Ranger KMS keyadmin user can import and export only KMS policies. The default credentials for this user are keyadmin/keyadmin.

Limitations

To successfully import policies, use the following database versions:

- MariaDB: 10.1.16+
- MySQL: 5.6.x+
- Oracle: 11gR2+
- PostgreSQL: 8.4+
- MS SQL: 2008 R2+

Partial policy import is not supported.

Related Information

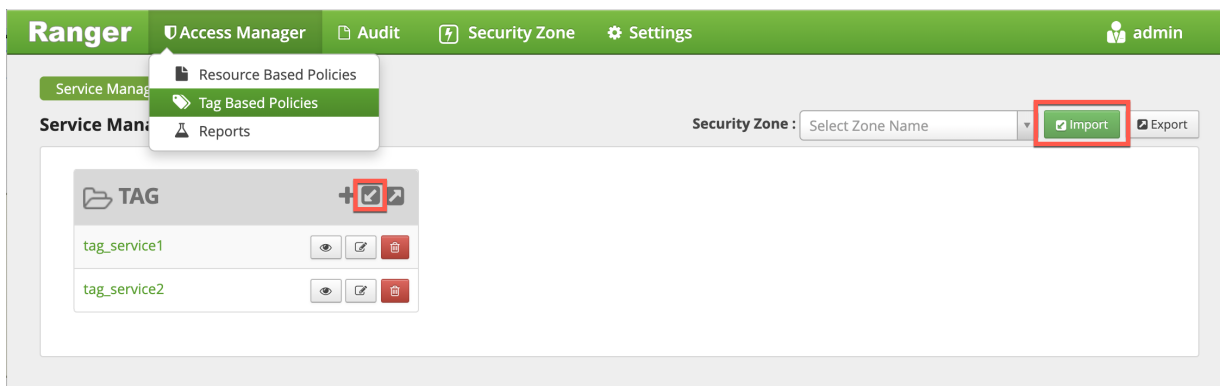
[Importing and Exporting Resource-Based Policies](#)

Import Tag Based Policies

How to import tag-based policies.

Procedure

1. On the Tag Based Policies page, click one of the Import icons:



2. Select the file to import.

You can only import policies in JSON format.

Import Policy ✕

Select File :

Select file

Override Policy :

Ranger_Policies_20190717_190622.json ✕

i All services gets listed on service destination when Zone destination is blank. When zone is selected at destination, then only services associated with that zone will be listed.

Specify Zone Mapping :

| Source | To | Destination |
|--------|----|---|
| | To | No zone selected ▼ |

Specify Service Mapping :

| Source | To | Destination |
|--|----|--|
| cm_hdfs ✕ ▼ | To | Select service name ▼ ✕ |

Cancel
Import

3. (Optional) Configure the import operation:

- a) The Override Policy option deletes all policies of the destination repositories.
- b) Zone Mapping – when no destination is selected, all services are imported. When a destination is selected, only the services associated with that security zone are imported.
- c) Service Mapping maps the downloaded file repository, i.e. source repository to destination repository. You can use the red x symbols to remove services from the import. Scroll down to view all service mappings.

Import Policy ✕

Specify Zone Mapping :

| Source | | Destination |
|--------|----|--------------------|
| | To | No zone selected ▼ |

Specify Service Mapping :

| Source | | Destination |
|--------------|----|-------------------------|
| cm_hdfs ✕ ▼ | To | Select service name ▼ ✕ |
| cm_hbase ✕ ▼ | To | Select service name ▼ ✕ |
| cm_yarn ✕ ▼ | To | Select service name ▼ ✕ |
| cm_hive ✕ ▼ | To | Select service name ▼ ✕ |
| cm_knox ✕ ▼ | To | Select service name ▼ ✕ |
| cm_storm ✕ ▼ | To | Select service name ▼ ✕ |

Cancel
Import

4. Click **Import**.

A confirmation message appears after the file is imported.

Related Information

[Export Tag-Based Policies](#)

Export Tag-Based Policies

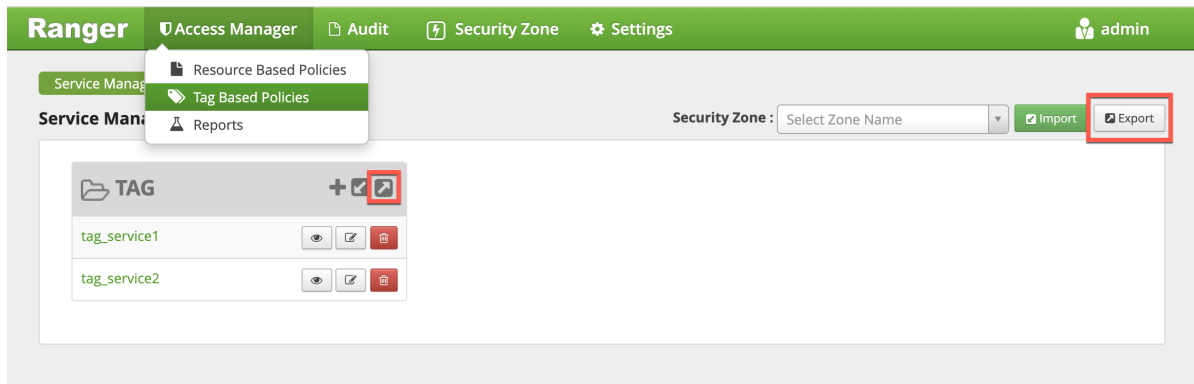
How to export all tag-based policies.

About this task

You can only export policies in JSON format from the Tag-based policies page. If you would like to export in Excel or CSV format, export the policies from the Reports page drop-down menu.


Procedure

- From the Access Manager>Tag Based Policies page:
 - a) Click the Export button or icon:



The Export Policy page appears.

- b) Remove components or specific services, then click Export.



- c) The file downloads in your browser as a JSON file.
- From the Reports page:
 - a) Filter **Component** to tag and click **Search**.
 - b) (Optional) Apply filters before exporting file.
 - c) Open the Export drop-down menu:

The screenshot shows the Ranger Access Manager interface. At the top, there are navigation tabs: Access Manager, Audit, Security Zone, and Settings. The user is logged in as 'admin'. The main section is titled 'User Access Report' and 'Reports'. Below this is a 'Search Criteria' form with fields for Policy Name, Policy Type (set to 'Access'), Component, Resource, Policy Label, Zone Name, and Search By (set to 'Group'). A 'Search' button is present. Below the search criteria is a table titled 'HDFS' with columns: Policy ID, Policy Name, Policy Labels, Resources, Policy Type, Status, Zone Name, Allow Conditions, Allow Exclude, Deny Conditions, and Deny. Two rows are visible in the table. An 'Export' dropdown menu is open on the right side of the table, showing options for 'Excel file', 'CSV file', and 'JSON file'. The 'Excel file' option is highlighted.

| Policy ID | Policy Name | Policy Labels | Resources | Policy Type | Status | Zone Name | Allow Conditions | Allow Exclude | Deny Conditions | Deny |
|-----------|----------------|---------------|-----------------------|-------------|---------|-----------|------------------|---------------|-----------------|------|
| 1 | all - path | -- | path:/* | Access | Enabled | -- | + | + | + | + |
| 2 | kms-audit-path | -- | path:ranger/audit/kms | Access | Enabled | -- | + | + | + | + |

- d) Select the file format.
The file downloads in your browser.

Related Information

[Import Tag Based Policies](#)

Create a Time-bound Policy

Ranger policy validity periods enable you to configure a policy to be effective for a specified time range. You can add a validity period to both resource-based and tag-based policies.

About this task

Time-bound policy use-case examples:

- To restrict access to sensitive financial information until the earnings release date.
- To block a certain user for a specific time period (e.g., a compromised user account being investigated needs to be put on "hold" from accessing resources in Hadoop services).
- To block a certain group for a specific time (e.g., excluding temporary employees from writing on resources during the holiday season).

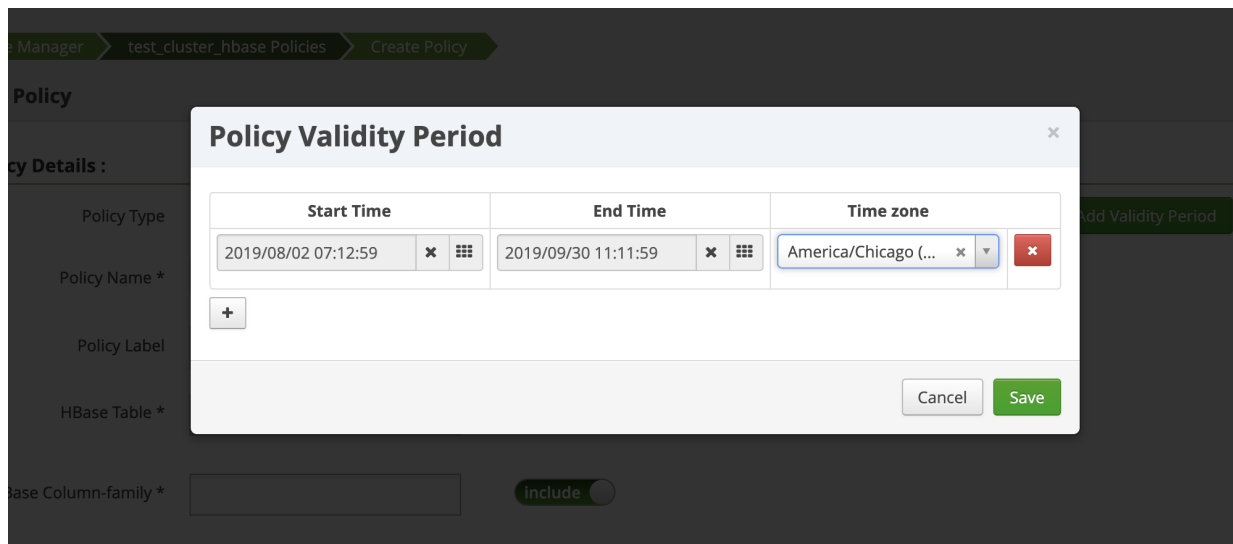


Note: The following procedure shows how to create a time-bound resource-based policy. The procedure is essentially the same for a tag-based resource policy.

Procedure

1. On the Ranger Service Manager page, select a service, then click Add New Policy.
2. Complete the fields on the **Create Policy** page.
3. Click **Add Validity Period**.

4. On the **Policy Validity Period** pop-up, specify a start time, end time, and time zone. To add additional validity periods, click the + symbol. Click Save to save the specified validity periods.



The screenshot shows a web interface for creating a policy. A modal dialog titled "Policy Validity Period" is open, allowing the user to define time-bound validity periods. The dialog contains a table with three columns: "Start Time", "End Time", and "Time zone".

| Start Time | End Time | Time zone |
|---------------------|---------------------|-----------------------|
| 2019/08/02 07:12:59 | 2019/09/30 11:11:59 | America/Chicago (...) |

Below the table is a "+" button to add more periods. At the bottom right of the dialog are "Cancel" and "Save" buttons. The background shows the "Policy Details" section with fields for Policy Name, Policy Label, HBase Table, and HBase Column-family, along with an "include" toggle.

5. If you would like the policy to override all other policies during its validity period, select **override**.

Ranger Access Manager Audit Security Zone Settings admin

Service Manager > test_cluster_hbase Policies > Create Policy

Create Policy

Policy Details :

Policy Type **Access** Add Validity Period

Policy Name * Temp employees override enabled **override**

Policy Label Policy Label

HBase Table * sales include

HBase Column-family * include

HBase Column * include

Description

Audit Logging **YES**

Allow Conditions : hide -

| Select Group | Select User | Permissions | Delegate Admin |
|----------------|-------------|-------------|--------------------------|
| temp_employees | Select User | Read | <input type="checkbox"/> |

Exclude from Allow Conditions : show -

Deny Conditions : show -

Add **Cancel**

6. Click **Add**.

Ranger Security Zones

Ranger security zones lets you organize service resources into multiple security zones.

Overview

Ranger Security Zones overview.

What is a Security Zone?

Lets you organize resource and tag-based services and policies into separate security zones. You can assign one or more administrators for each security zone. Security zone administrators can then create and update policies for their security zone.

For example, let us consider two security zones: "finance" and "sales":

- Security zone "finance" includes all content in a "finance" Hive database.
- Security zone "sales" includes all content in a "sales" Hive database.
- Sets of users and groups are designated as administrators in each security zone.
- Users are allowed to set up policies only in security zones in which they are administrators.
- Policies defined in a security zone are applicable only for resources of that zone.
- A zone can be extended to include resources from multiple services such as HDFS, Hive, HBase, Kafka, etc., allowing administrators of a zone to set up policies for resources owned by their organization across multiple services.

```
Zone: finance
service: prod_hdfs; path=/finance/*, /taxes/*
service: prod_hive; database=finance
service: prod_kafka; topic=FIN_*
service: test_hadoop; path=/finance/*, /taxes/*
Zone: sales
service: prod_hdfs; path=/sales/*
service: prod_hive; database=sales
service: prod_kafka; topic=SALES_*
```

- As shown above, resources can be specified using wildcards (FIN_*, SALES_*).
- A resource is not mappable to more than one security zone. Ranger does not allow creation of security zones that specify resources that match resources in another zone. For example, an attempt to update the "finance" zone above with the HDFS path /sales/finance/* is not be permitted, as this conflicts with the HDFS path /sales/* specified in the "sales" zone.
- A set of users and groups can be designated as administrators of a security zone. Administrators can create, update, and delete security policies for the resources in the security zone.
- A set of users and groups can be authorized to view audit logs of access to a security zone's resources. Other users are not allowed to view access-audit logs of the security zone resources.

Security Zone Administration

- Security zones can only be created, updated, or deleted by a user with the ROLE_SYS_ADMIN role in Ranger.
- Users can view, retrieve, and update policies only in security zones in which they have administrator privileges.

How are Security Zones Used in Authorization?

When a Ranger authorization plugin authorizes a resource access request, it first determines the zone in which the accessed resource resides. If the resource matches a security zone, only the policies of that security zone are used to authorize the access. If resource does not match any security zone, the policies in the default (unnamed) security zone are used to authorize the access.

Tag-based Policies in Security Zones

In a given service, each security zone can be configured to use tag-based policies from a specific security zone in a tag-service. This enables tag-based authorization policies to be used based on the security zone of the resource.

Audit Logs

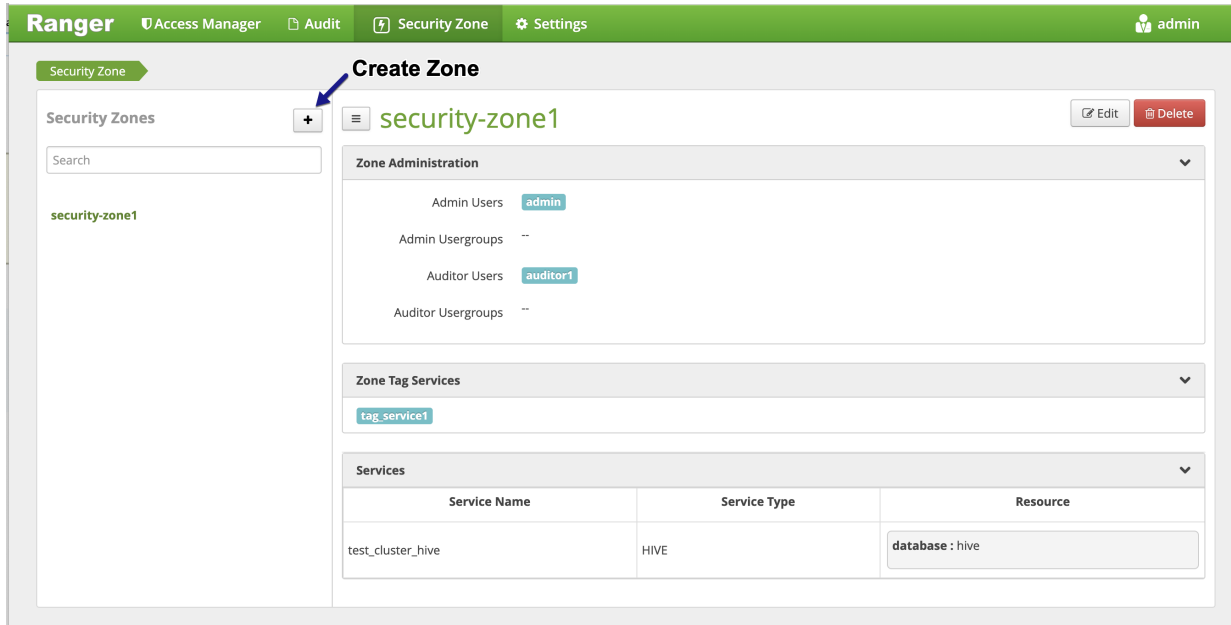
Audit logs generated by Ranger include the name of the security zone in which the accessed resource resides. Only users who have been assigned as an Admin or Auditor for the security zone are allowed to view the audit logs.

Adding a Ranger Security Zone

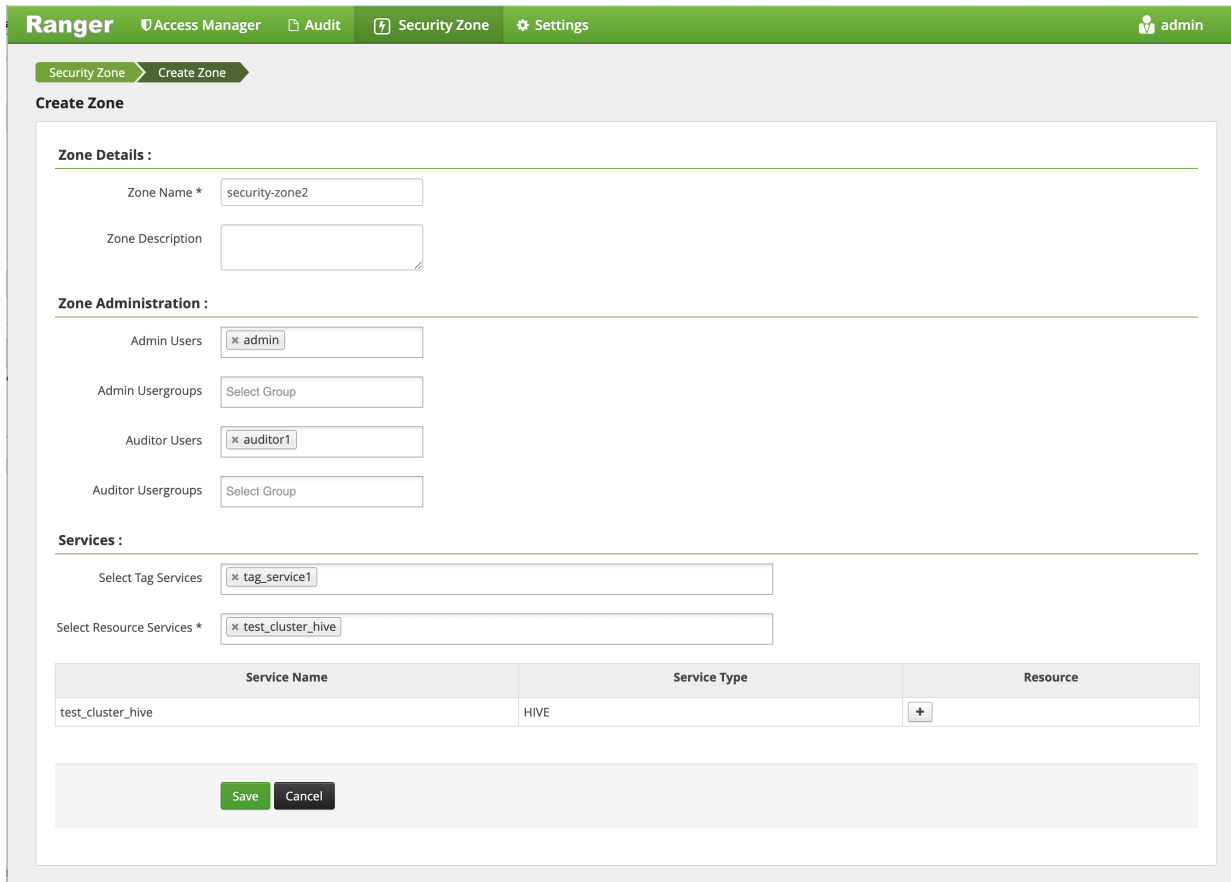
How to add a new Ranger Security Zone.

Procedure

1. Click Security Zone in the top menu.
The Security Zone page appears.
2. On the Security Zone page, click the + icon.



The Create Zone page appears.



- Complete the Create Zone page as follows:

Table 60: Zone Details

| Field | Description |
|------------------|--------------------------|
| Zone Name | The security zone name. |
| Zone Description | An optional description. |

Table 61: Zone Administration

| Field | Description |
|--------------------|--|
| Admin Users | The Admin users for the security zone. |
| Admin Usergroups | The Admin user groups for the security zone. |
| Auditor Users | The Auditor users for the security zone. |
| Auditor Usergroups | The Auditor user groups for the security zone. |

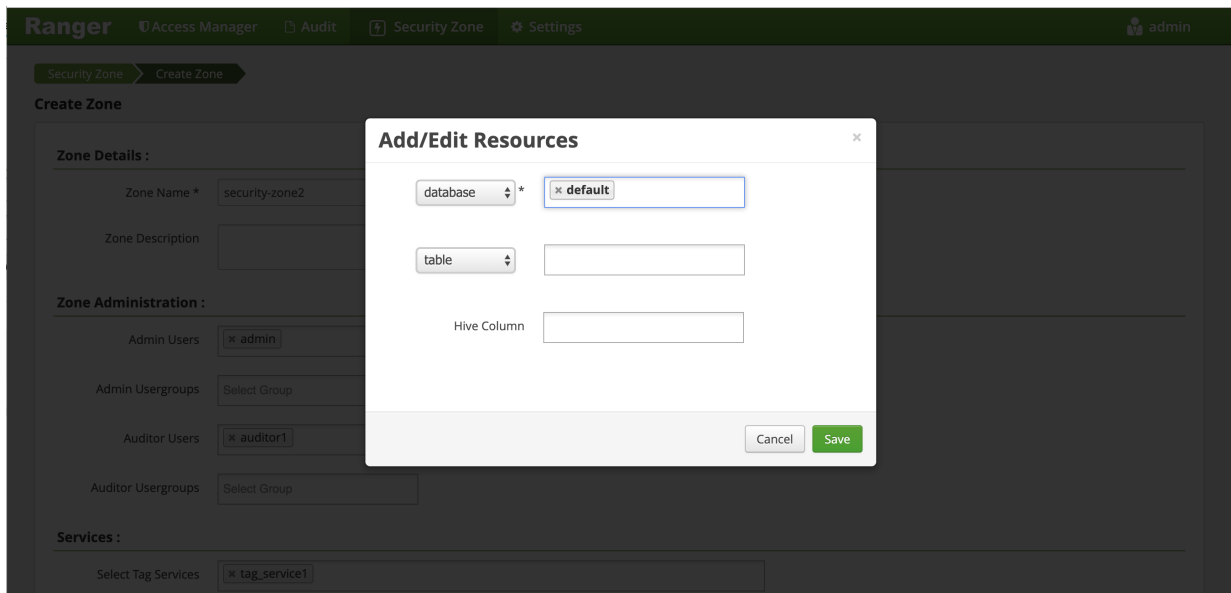
Table 62: Services

| Label | Description |
|--------------------------|---|
| Select Tag Services | Select tag-based services for the security zone. |
| Select Resource Services | Select resource-based services for the security zone. |

- Selected Services are listed in the Services table. To add resources for each selected service, click the + icon in the Resources column for the applicable service.

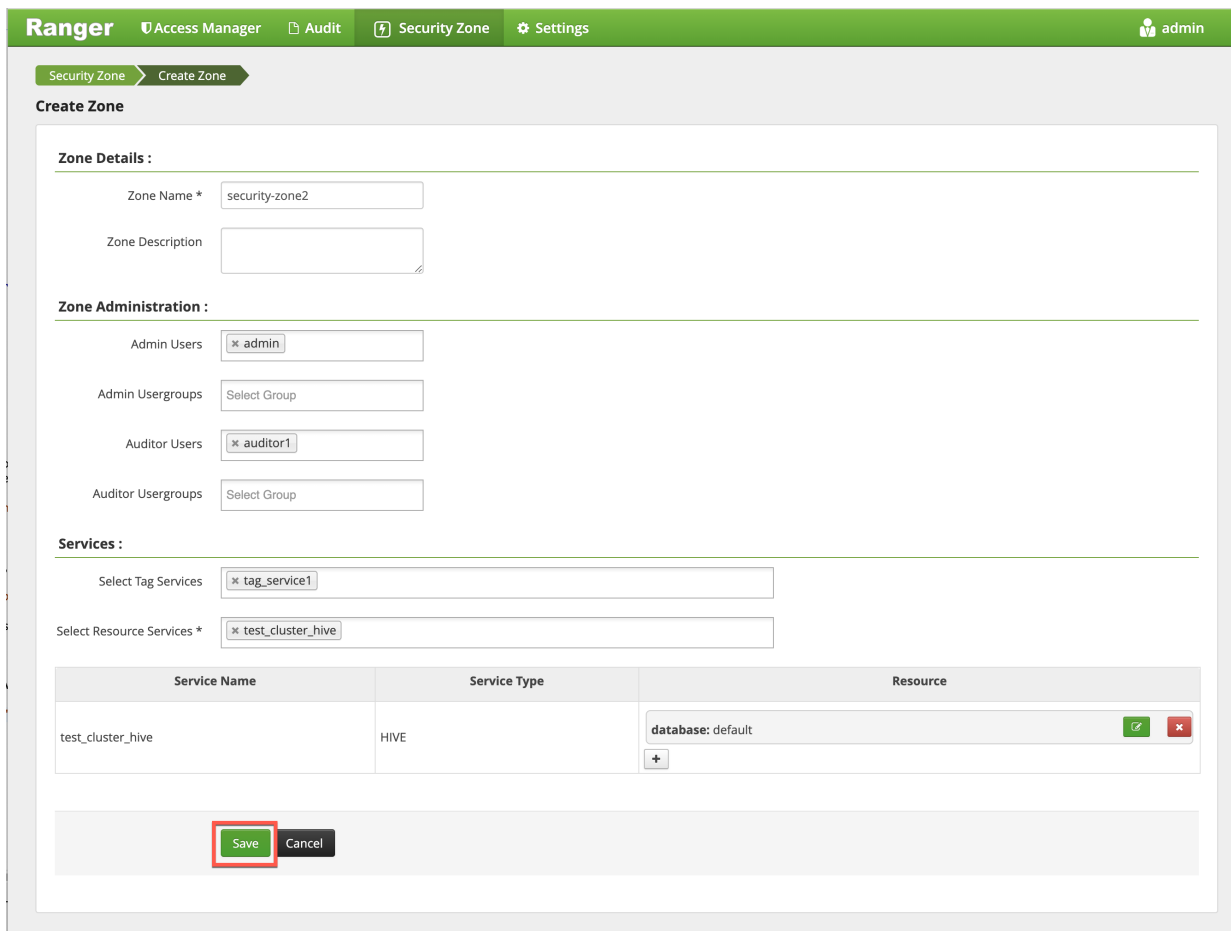
The screenshot shows the 'Create Zone' page in the Ranger interface. The page is divided into three main sections: Zone Details, Zone Administration, and Services. In the Zone Details section, the Zone Name is 'security-zone2' and the Zone Description is empty. In the Zone Administration section, Admin Users is 'admin', Admin Usergroups is 'Select Group', Auditor Users is 'auditor1', and Auditor Usergroups is 'Select Group'. In the Services section, Select Tag Services is 'tag_service1' and Select Resource Services is 'test_cluster_hive'. Below the Services section is a table with columns Service Name, Service Type, and Resource. The table contains one row: test_cluster_hive, HIVE, and a '+' icon. A blue arrow points to the '+' icon with the text 'Add Resources'. At the bottom of the page are 'Save' and 'Cancel' buttons.

- Use the Add/Edit Resources pop-up to specify resources for the service, then click Save.

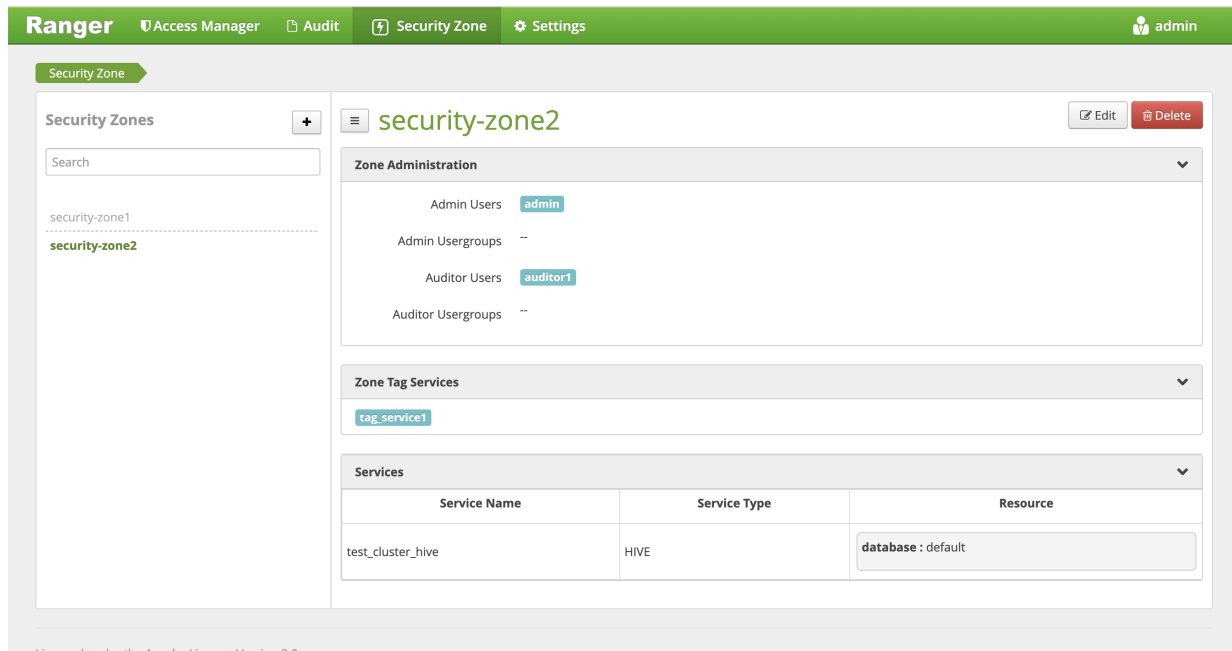


The resources are listed in the Resources column of the Services table.

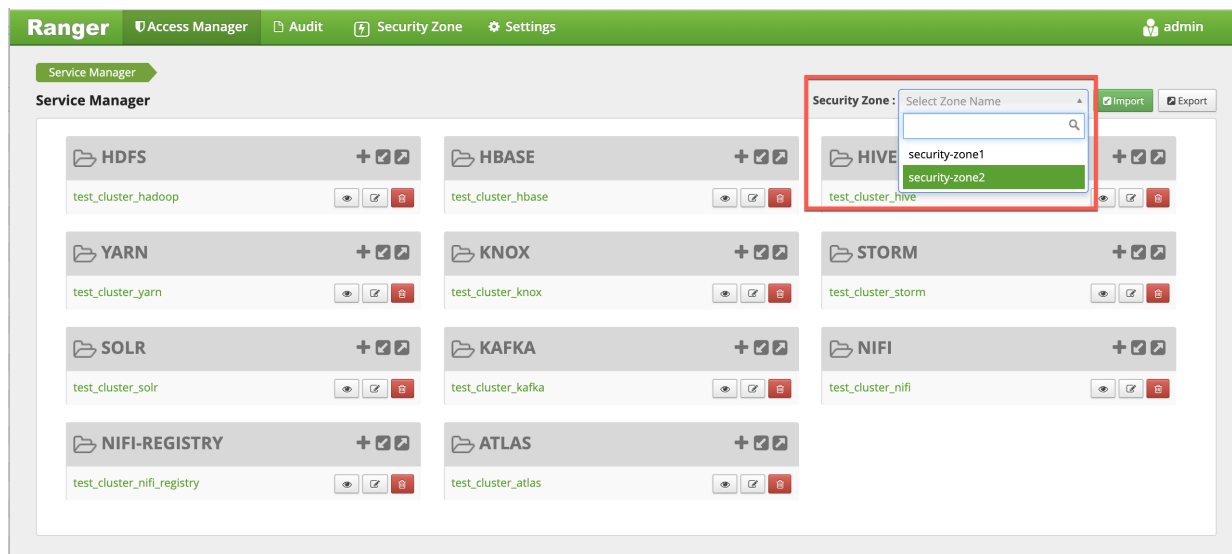
- Click Save at the bottom of the Create Zone page to save the new security zone.



- The new security zone is listed on the Security Zone page.



8. To edit a security zone, click the security zone name in the Security Zones list, then click Edit.
9. After security zones have been created, you can use the Security Zone selection box on the Service Manager page to display the services assigned to the selected security zone. A Zone Name column appears in the table on the Audit > Access page, and also in the Access Manager > Reports tables.



Administering Ranger Users, Groups, and Permissions

To view the list of users and groups who can access the Ranger portal or its services, select Settings > Users/Groups in the top menu.

The Users/Groups page lists:

- Internal users who can log in to the Ranger portal; created by the Ranger console Service Manager.
- External users who can access services controlled by the Ranger portal; created at other systems such as Active Directory, LDAP, or UNIX, and synched with those systems.

- Admin users who are the only users with permission to create users and services, run reports, and perform other administrative tasks. Admin users can also create child policies based on the original policy (base policy).
- On the Groups page, you can click the people icons in the Users column to view the members of the applicable group.

The screenshot shows the Ranger interface with the 'Users/Groups' section selected. The 'Groups' tab is active, displaying a 'Group List' table. The table has columns for Group Name, Group Source, Visibility, and Users. A search bar is at the top left, and 'Add New Group' and 'Set Visibility' buttons are at the top right.

| <input type="checkbox"/> | Group Name | Group Source | Visibility | Users |
|--------------------------|------------|--------------|------------|-------|
| <input type="checkbox"/> | public | Internal | Visible | |
| <input type="checkbox"/> | hadoop | External | Visible | |
| <input type="checkbox"/> | ranger | External | Visible | |
| <input type="checkbox"/> | hdfs | External | Visible | |
| <input type="checkbox"/> | polkitd | External | Visible | |
| <input type="checkbox"/> | nfsnobody | External | Visible | |
| <input type="checkbox"/> | spark | External | Visible | |
| <input type="checkbox"/> | jenkins | External | Visible | |

Add a User

How to add a new user to the user list in Ranger.

Procedure

1. Select **Settings > Users/Groups**.
The Users/Groups page appears.

The screenshot shows the Ranger interface with the 'Settings' menu open and 'Users/Groups' selected. The 'Users' tab is active, displaying a 'User List' table. The table has columns for User Name, Email Address, Role, User Source, Groups, and Visibility. A search bar is at the top left, and 'Add New User' and 'Set Visibility' buttons are at the top right.

| <input type="checkbox"/> | User Name | Email Address | Role | User Source | Groups | Visibility |
|--------------------------|-------------------|---------------|-------|-------------|---------------|------------|
| <input type="checkbox"/> | admin | | Admin | Internal | -- | Visible |
| <input type="checkbox"/> | rangerusersync | | Admin | Internal | -- | Visible |
| <input type="checkbox"/> | rangertagsync | | Admin | Internal | -- | Visible |
| <input type="checkbox"/> | yarn-ats | | User | External | hadoop | Visible |
| <input type="checkbox"/> | hive | | User | External | hadoop | Visible |
| <input type="checkbox"/> | infra-solr | | User | External | hadoop | Visible |
| <input type="checkbox"/> | atlas | | User | External | hadoop | Visible |
| <input type="checkbox"/> | ams | | User | External | hadoop | Visible |
| <input type="checkbox"/> | ranger | | User | External | hadoop ranger | Visible |
| <input type="checkbox"/> | activity_analyzer | | User | External | hadoop hdfs | Visible |
| <input type="checkbox"/> | polkitd | | User | External | polkitd | Visible |

2. Click **Add New User**.

The User Detail page appears.

Ranger Access Manager Audit Security Zone Settings admin

Users/Groups > User Create

User Detail

User Name * admin ⓘ

New Password * ⓘ

Password Confirm * ⓘ

First Name * ⓘ

Last Name ⓘ

Email Address ⓘ

Select Role * Admin ▾

Group *Please select* +

Save Cancel

3. Add the required user details, then click **Save**.
The user is immediately added to the list.

Edit a User

How to edit a user in Ranger.

Procedure

1. Select **Settings > Users/Groups**.
The Users/Groups page opens to the Users tab.

Ranger Access Manager Audit Security Zone Settings admin

Users/Groups

Users Groups

Group List

Search for your groups...

Add New Group Set Visibility ▾

| <input type="checkbox"/> | Group Name | Group Source | Visibility | Users |
|--------------------------|------------|--------------|------------|-------|
| <input type="checkbox"/> | public | Internal | Visible | |
| <input type="checkbox"/> | hadoop | External | Visible | |
| <input type="checkbox"/> | ranger | External | Visible | |
| <input type="checkbox"/> | hdfs | External | Visible | |
| <input type="checkbox"/> | polkitd | External | Visible | |
| <input type="checkbox"/> | nfsnobody | External | Visible | |
| <input type="checkbox"/> | spark | External | Visible | |
| <input type="checkbox"/> | jenkins | External | Visible | |

2. Select a user profile to edit. To edit your own profile, select your user name, then click Profile.

The screenshot shows the Ranger Admin console interface. At the top, there is a navigation bar with 'Ranger', 'Access Manager', 'Audit', 'Security Zone', 'Settings', and a user profile 'admin'. Below this, there are tabs for 'Users/Groups', 'Users', and 'Groups'. A dropdown menu for the 'admin' user shows 'Profile' and 'Log Out' options. The main content area is titled 'User List' and contains a search bar, 'Add New User', 'Set Visibility', and a trash icon. A table lists users with columns for 'User Name', 'Email Address', 'Role', 'User Source', 'Groups', and 'Visibility'. The 'hive' user is highlighted with a blue arrow pointing to it, and the text 'Edit a user profile' is placed above the arrow. Another blue arrow points to the 'Profile' button in the top right, with the text 'Edit your own profile' placed above it.

| | User Name | Email Address | Role | User Source | Groups | Visibility |
|--------------------------|-------------------|---------------|-------|-------------|---------------|------------|
| <input type="checkbox"/> | admin | | Admin | Internal | -- | Visible |
| <input type="checkbox"/> | rangerusersync | | Admin | Internal | -- | Visible |
| <input type="checkbox"/> | rangertagsync | | Admin | Internal | -- | Visible |
| <input type="checkbox"/> | yarn-ah | | User | External | hadoop | Visible |
| <input type="checkbox"/> | hive | | User | External | hadoop | Visible |
| <input type="checkbox"/> | infra-solr | | User | External | hadoop | Visible |
| <input type="checkbox"/> | atlas | | User | External | hadoop | Visible |
| <input type="checkbox"/> | ams | | User | External | hadoop | Visible |
| <input type="checkbox"/> | ranger | | User | External | hadoop ranger | Visible |
| <input type="checkbox"/> | activity_analyzer | | User | External | hadoop hdfs | Visible |

The User Detail page appears.

The screenshot shows the 'User Detail' page in the Ranger Admin console. The navigation bar is the same as in the previous screenshot. The breadcrumb trail shows 'Users/Groups > User Edit'. The page title is 'User Detail'. There are two tabs: 'Basic Info' (active) and 'Change Password'. The 'Basic Info' tab contains the following form fields:

- User Name *: rangerusersync
- First Name *: rangerusersync
- Last Name: (empty)
- Email Address: (empty)
- Select Role *: Admin
- Group: Please select

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



Note:

You can only fully edit internal users. For external users, you can only edit the user role.

3. Edit the user details, then click **Save**.

Delete a User

How to delete a user in Ranger.

Before you begin

Only users with the "admin" role can delete a user.

Procedure

1. Select **Settings > Users/Groups**.

The Users/Groups page appears.

| <input type="checkbox"/> | User Name | Email Address | Role | User Source | Groups | Visibility |
|--------------------------|-------------------|---------------|-------|-------------|---------------|------------|
| <input type="checkbox"/> | admin | | Admin | Internal | -- | Visible |
| <input type="checkbox"/> | rangerusersync | | Admin | Internal | -- | Visible |
| <input type="checkbox"/> | rangertagsync | | Admin | Internal | -- | Visible |
| <input type="checkbox"/> | yarn-ats | | User | External | hadoop | Visible |
| <input type="checkbox"/> | hive | | User | External | hadoop | Visible |
| <input type="checkbox"/> | infra-solr | | User | External | hadoop | Visible |
| <input type="checkbox"/> | atlas | | User | External | hadoop | Visible |
| <input type="checkbox"/> | ams | | User | External | hadoop | Visible |
| <input type="checkbox"/> | ranger | | User | External | hadoop ranger | Visible |
| <input type="checkbox"/> | activity_analyzer | | User | External | hadoop hdfs | Visible |
| <input type="checkbox"/> | polkitd | | User | External | polkitd | Visible |

2. Select the check box of the user you want to delete, then click the Delete icon



at the right of the User List menu bar.

| <input type="checkbox"/> | User Name | Email Address | Role | User Source | Groups | Visibility |
|-------------------------------------|-------------------|---------------|-------|-------------|---------------|------------|
| <input type="checkbox"/> | admin | | Admin | Internal | -- | Visible |
| <input type="checkbox"/> | rangerusersync | | Admin | Internal | -- | Visible |
| <input checked="" type="checkbox"/> | rangertagsync | | Admin | Internal | -- | Visible |
| <input type="checkbox"/> | yarn-ats | | User | External | hadoop | Visible |
| <input type="checkbox"/> | hive | | User | External | hadoop | Visible |
| <input type="checkbox"/> | infra-solr | | User | External | hadoop | Visible |
| <input type="checkbox"/> | atlas | | User | External | hadoop | Visible |
| <input type="checkbox"/> | ams | | User | External | hadoop | Visible |
| <input type="checkbox"/> | ranger | | User | External | hadoop ranger | Visible |
| <input type="checkbox"/> | activity_analyzer | | User | External | hadoop hdfs | Visible |

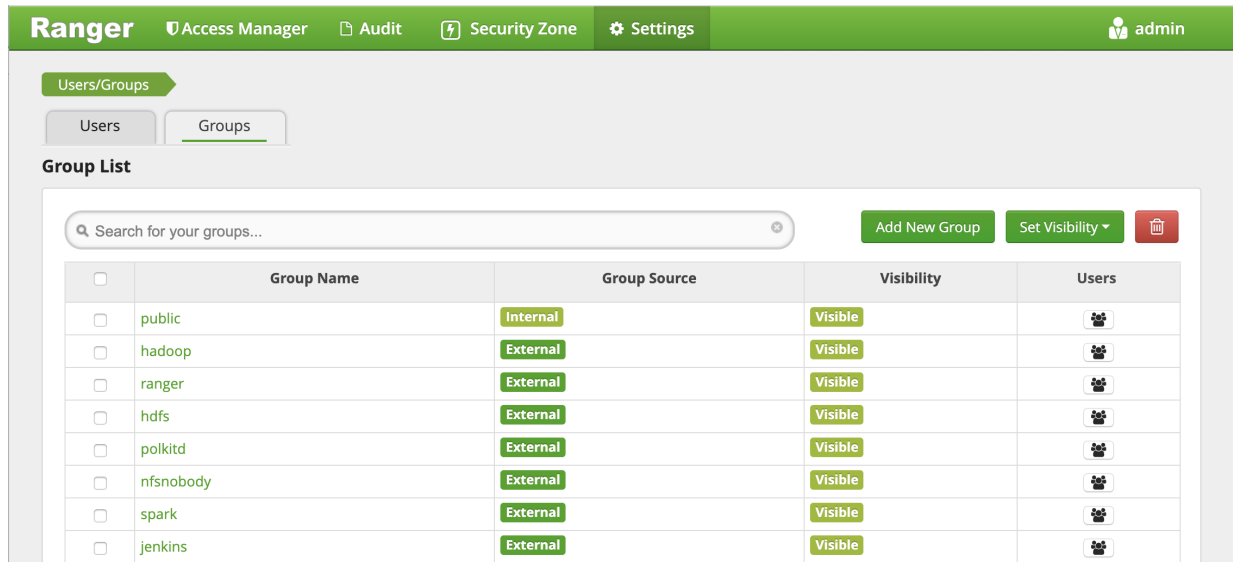
3. Click OK on the confirmation pop-up.

Add a Group

How to add a group in Ranger.

Procedure

1. Select **Settings > Users/Groups**, then click the **Groups** tab.
The Groups page appears.

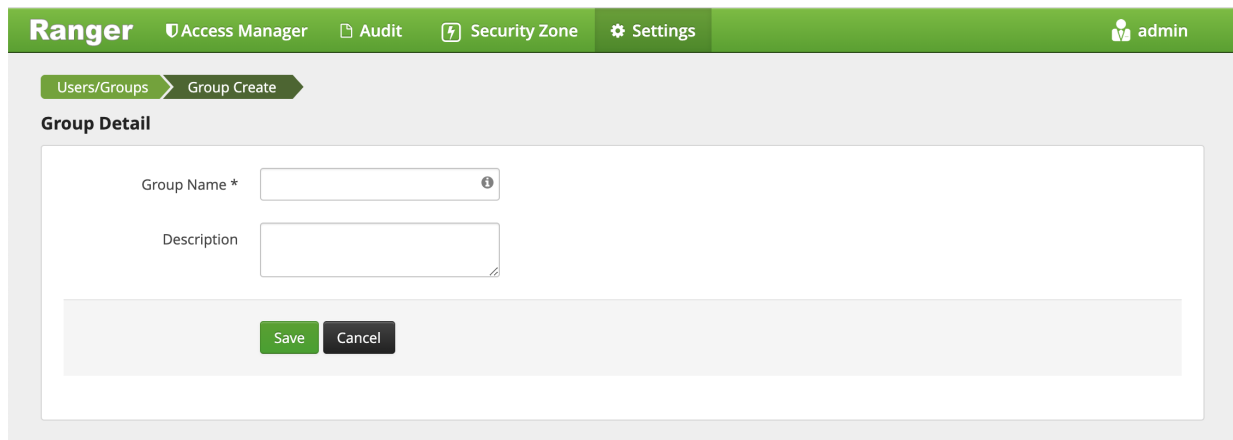


The screenshot shows the Ranger web interface. The top navigation bar includes 'Ranger', 'Access Manager', 'Audit', 'Security Zone', 'Settings', and a user profile 'admin'. The breadcrumb trail is 'Users/Groups' > 'Groups'. Below the breadcrumb are tabs for 'Users' and 'Groups'. The main content area is titled 'Group List' and contains a search bar, 'Add New Group', 'Set Visibility', and a delete icon. A table lists the following groups:

| <input type="checkbox"/> | Group Name | Group Source | Visibility | Users |
|--------------------------|------------|--------------|------------|-------|
| <input type="checkbox"/> | public | Internal | Visible | |
| <input type="checkbox"/> | hadoop | External | Visible | |
| <input type="checkbox"/> | ranger | External | Visible | |
| <input type="checkbox"/> | hdfs | External | Visible | |
| <input type="checkbox"/> | polkitd | External | Visible | |
| <input type="checkbox"/> | nfsnobody | External | Visible | |
| <input type="checkbox"/> | spark | External | Visible | |
| <input type="checkbox"/> | jenkins | External | Visible | |

2. Click **Add New Group**.

The Group Create page appears.



The screenshot shows the 'Group Create' page in the Ranger interface. The breadcrumb trail is 'Users/Groups' > 'Group Create'. The page is titled 'Group Detail' and contains a form with the following fields:

- Group Name * (required field)
- Description

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

3. Enter a unique name for the group and an optional description, then click **Save**.

Edit a Group

How to edit a group in Ranger.

Procedure

1. Select **Settings > Users/Groups**, then click the **Groups** tab.

The Groups page appears.

Ranger Access Manager Audit Security Zone Settings admin

Users/Groups

Users Groups

Group List

Search for your groups... Add New Group Set Visibility

| <input type="checkbox"/> | Group Name | Group Source | Visibility | Users |
|--------------------------|------------|--------------|------------|-------|
| <input type="checkbox"/> | public | Internal | Visible | |
| <input type="checkbox"/> | hadoop | External | Visible | |
| <input type="checkbox"/> | ranger | External | Visible | |
| <input type="checkbox"/> | hdfs | External | Visible | |
| <input type="checkbox"/> | polkitd | External | Visible | |
| <input type="checkbox"/> | nfsnobody | External | Visible | |
| <input type="checkbox"/> | spark | External | Visible | |
| <input type="checkbox"/> | jenkins | External | Visible | |

2. Select a group name to edit.

Ranger Access Manager Audit Security Zone Settings admin

Users/Groups

Users Groups

Group List

Search for your groups... Add New Group Set Visibility

| <input type="checkbox"/> | Group Name | Group Source | Visibility | Users |
|--------------------------|------------|--------------|------------|-------|
| <input type="checkbox"/> | public | Internal | Visible | |
| <input type="checkbox"/> | hadoop | External | Visible | |
| <input type="checkbox"/> | ranger | External | Visible | |
| <input type="checkbox"/> | hdfs | External | Visible | |
| <input type="checkbox"/> | polkitd | External | Visible | |
| <input type="checkbox"/> | nfsnobody | External | Visible | |
| <input type="checkbox"/> | spark | External | Visible | |
| <input type="checkbox"/> | jenkins | External | Visible | |
| <input type="checkbox"/> | users | External | Visible | |

3. The Group Edit page appears.

Ranger Access Manager Audit Security Zone Settings admin

Users/Groups Group Edit

Group Detail

Group Name * public

Description public group

Save Cancel

4. Edit the group details, then click **Save**.

Delete a Group

How to delete a group in Ranger.

Before you begin

Only users with the "admin" role can delete a group.

Procedure

1. Select **Settings > Users/Groups**, then click the **Groups** tab.

The Groups page appears.

| <input type="checkbox"/> | Group Name | Group Source | Visibility | Users |
|--------------------------|------------|--------------|------------|-------|
| <input type="checkbox"/> | public | Internal | Visible | |
| <input type="checkbox"/> | hadoop | External | Visible | |
| <input type="checkbox"/> | ranger | External | Visible | |
| <input type="checkbox"/> | hdfs | External | Visible | |
| <input type="checkbox"/> | polkitd | External | Visible | |
| <input type="checkbox"/> | nfsnobody | External | Visible | |
| <input type="checkbox"/> | spark | External | Visible | |
| <input type="checkbox"/> | jenkins | External | Visible | |

2. Select the check box of the group you want to delete, then click the Delete icon



at the right of the Group List menu bar.

The screenshot shows the Ranger web interface. At the top, there is a navigation bar with 'Ranger' and 'Access Manager', 'Audit', 'Security Zone', and 'Settings' tabs. The user 'admin' is logged in. Below the navigation bar, there are tabs for 'Users/Groups', 'Users', and 'Groups'. The 'Groups' tab is active, and the 'Group List' section is displayed. A search bar is at the top of the group list, and there are buttons for 'Add New Group', 'Set Visibility', and a delete icon (highlighted with a red box). The table below lists various groups with columns for Group Name, Group Source, Visibility, and Users.

| <input type="checkbox"/> | Group Name | Group Source | Visibility | Users |
|-------------------------------------|------------|--------------|------------|-------|
| <input type="checkbox"/> | public | Internal | Visible | |
| <input type="checkbox"/> | hadoop | External | Visible | |
| <input type="checkbox"/> | ranger | External | Visible | |
| <input type="checkbox"/> | hdfs | External | Visible | |
| <input checked="" type="checkbox"/> | polkitd | External | Visible | |
| <input type="checkbox"/> | nfsnobody | External | Visible | |
| <input type="checkbox"/> | spark | External | Visible | |
| <input type="checkbox"/> | jenkins | External | Visible | |
| <input type="checkbox"/> | users | External | Visible | |
| <input type="checkbox"/> | zeppelin | External | Visible | |
| <input type="checkbox"/> | cveetaet | External | Visible | |

3. Click OK on the confirmation pop-up.

What to do next

Users in a deleted group will be reassigned to no group. You can edit these users and reassign them to other groups.

Related Information

[Edit a User](#)

Add/Edit Permissions

How to add or edit a user or group in Ranger.

Procedure

1. Select **Settings > Permissions**.
The Permissions page appears.

| Modules | Groups | Users | Action |
|-------------------------|--------|--|--------|
| Resource Based Policies | | admin rangerusersync keyadmin rangertagsync + More.. | |
| Users/Groups | | admin rangerusersync rangertagsync keyadmin + More.. | |
| Reports | | admin rangerusersync keyadmin rangertagsync + More.. | |
| Audit | | admin rangerusersync rangertagsync keyadmin + More.. | |
| Key Manager | | keyadmin | |
| Tag Based Policies | | admin rangerusersync rangertagsync amb_ranger_admin + More.. | |
| Security Zone | | admin rangerusersync rangertagsync yarn-ats + More.. | |

2. Click the Edit icon



for the permission you would like to edit.
The Edit Permission page appears.

3. Edit the permission settings, then click **Save**.

You can select multiple users and groups using the + icons.

Administering Ranger Reports

You can use the Reports page to help manage policies more efficiently as the number of policies increases. This page lists all resource-based and tag-based policies.

The screenshot shows the Ranger web interface. At the top, there's a navigation bar with 'Ranger', 'Access Manager', 'Audit', 'Security Zone', and 'Settings'. A user profile 'admin' is visible. Below the navigation bar, there are tabs for 'Resource Based Policies', 'Tag Based Policies', and 'Reports'. The 'Reports' tab is active, showing a search criteria section with fields for Policy Name, Policy Type (Access), Component, Resource, Policy Label, Zone Name, and Search By (Group). A 'Search' button is present. Below the search criteria, there are two tables: one for HDFS and one for HBASE. The HDFS table has columns: Policy ID, Policy Name, Policy Labels, Resources, Policy Type, Status, Zone Name, Allow Conditions, Allow Exclude, and Deny Conditions. The HBASE table has columns: Policy ID, Policy Name, Policy Labels, Resources, Policy Type, Status, Zone Name, Allow Conditions, and Allow Exclude.

| Policy ID | Policy Name | Policy Labels | Resources | Policy Type | Status | Zone Name | Allow Conditions | Allow Exclude | Deny Conditions |
|-----------|----------------|---------------|------------------------|-------------|---------|-----------|------------------|---------------|-----------------|
| 1 | all - path | -- | path:/* | Access | Enabled | -- | + | + | + |
| 2 | kms-audit-path | -- | path:/ranger/audit/kms | Access | Enabled | -- | + | + | + |

| Policy ID | Policy Name | Policy Labels | Resources | Policy Type | Status | Zone Name | Allow Conditions | Allow Exclude |
|-----------|------------------------------------|---------------|--|-------------|---------|-----------|------------------|---------------|
| 3 | all - table, column-family, col... | -- | column-family:* column:* table:* | Access | Enabled | -- | + | + |
| 4 | Service Check User Policy fo... | -- | column-family:* column:* table:ambarismoketest | Access | Enabled | -- | + | + |

View Ranger Reports

How to view reports for Ranger policies.

To view reports on one or more policies, select **Access Manager > Reports**.

- To view Allow Condition details for each policy, click the



icon in the Allow Conditions column. You can use the same method to view details for other policy conditions (Allow Exclude, Deny Conditions, etc.).

- To edit a policy from the Reports page, click the Policy ID.

Search Criteria

Policy Name: Policy Type:

Component: Resource:

Policy Label: Zone Name:

Search By:

HDFS

| Policy ID | Policy Name | Policy Labels | Resources | Policy Type | Status | Zone Name | Allow Conditions | Allow Exclude | Deny Conditio |
|-----------|----------------|---------------|------------------------|-------------|---------|-----------|------------------|---------------|---------------|
| 1 | all - path | -- | path:/* | Access | Enabled | -- | + | + | + |
| 2 | kms-audit-path | -- | path:/ranger/audit/kms | Access | Enabled | -- | + | + | + |

HBASE

| Policy ID | Policy Name | Policy Labels | Resources | Policy Type | Status | Zone Name | Allow Conditions | Allow Ex |
|-----------|------------------------------------|---------------|--|-------------|---------|-----------|------------------|----------|
| 3 | all - table, column-family, col... | -- | column-family:* column:* table:* | Access | Enabled | -- | + | + |
| 4 | Service Check User Policy fo... | -- | column-family:* column:* table:ambarismoketest | Access | Enabled | -- | + | + |

Search Ranger Reports

Reference information for searching Ranger reports on one or more policies.

You can search based on:

- Policy Name – The policy name.
- Policy Type – The policy type (Access, Masking, or Row Level Filter).
- Policy Label – The policy label.
- Component – The policy resource or tag component.
- Resource – The resource path used when creating the policy.
- Zone Name – The security zone name.
- Group, Username – The group or user name assigned to the policy.

The screenshot shows the Ranger Reports interface. At the top, there is a navigation bar with 'Ranger' and 'Access Manager', 'Audit', 'Security Zone', and 'Settings' tabs. A user profile 'admin' is visible in the top right. A dropdown menu is open under 'Access Manager', showing 'Resource Based Policies', 'Tag Based Policies', and 'Reports'. The 'Reports' section is active, displaying a 'Search Criteria' form with fields for Policy Name, Policy Type (Access), Component, Resource, Policy Label, Zone Name, and Search By (Group). A 'Search' button is present. Below the search form is an 'Export' button. The main content area is divided into two sections: 'HDFS' and 'HBASE', each with a 'hide' link. The 'HDFS' section contains a table with 10 columns: Policy ID, Policy Name, Policy Labels, Resources, Policy Type, Status, Zone Name, Allow Conditions, Allow Exclude, and Deny Conditions. It lists two policies: Policy 1 (all - path) and Policy 2 (kms-audit-path). The 'HBASE' section contains a similar table with 10 columns, listing two policies: Policy 3 (all - table, column-family, col...) and Policy 4 (Service Check User Policy fo...).

Search Criteria hide ↗

Policy Name: Policy Type: ⓘ

Component: Resource:

Policy Label: Zone Name:

Search By:

HDFS hide ↗

| Policy ID | Policy Name | Policy Labels | Resources | Policy Type | Status | Zone Name | Allow Conditions | Allow Exclude | Deny Conditions |
|-----------|----------------|---------------|------------------------|-------------|---------|-----------|------------------|---------------|-----------------|
| 1 | all - path | -- | path:/* | Access | Enabled | -- | + | + | + |
| 2 | kms-audit-path | -- | path:/ranger/audit/kms | Access | Enabled | -- | + | + | + |

HBASE hide ↗

| Policy ID | Policy Name | Policy Labels | Resources | Policy Type | Status | Zone Name | Allow Conditions | Allow Exclude | Deny Conditions |
|-----------|------------------------------------|---------------|--|-------------|---------|-----------|------------------|---------------|-----------------|
| 3 | all - table, column-family, col... | -- | column-family:* column:* table:* | Access | Enabled | -- | + | + | + |
| 4 | Service Check User Policy fo... | -- | column-family:* column:* table:ambarismoketest | Access | Enabled | -- | + | + | + |

Export Reports

Reference information for exporting Ranger reports on one or more policies.

You can export a list of reports in three file formats:

- CSV file
- Excel file
- JSON

The screenshot shows the Apache Ranger User Access Report interface. At the top, there is a navigation bar with 'Ranger', 'Access Manager', 'Audit', 'Security Zone', and 'Settings'. The user is logged in as 'admin'. Below the navigation bar, there is a 'User Access Report' section with a 'reports' tab. The main content area is divided into sections for 'HDFS' and 'HBASE'. Each section contains a table of policies. The 'HDFS' section has two rows of policies, and the 'HBASE' section has one row. An 'Export' button is located to the right of the HDFS table, and a dropdown menu is open, showing options for 'Excel file', 'CSV file', and 'JSON file'.

Search Criteria

Policy Name: Policy Type:

Component: Resource:

Policy Label: Zone Name:

Search By:

HDFS

| Policy ID | Policy Name | Policy Labels | Resources | Policy Type | Status | Zone Name | Allow Conditions | Allow Exclude | De |
|-----------|----------------|---------------|------------------------|-------------|---------|-----------|------------------|---------------|----|
| 1 | all - path | -- | path:/* | Access | Enabled | -- | + | + | + |
| 2 | kms-audit-path | -- | path:/ranger/audit/kms | Access | Enabled | -- | + | + | + |

HBASE

| Policy ID | Policy Name | Policy Labels | Resources | Policy Type | Status | Zone Name | Allow Conditions | Allow Ex |
|-----------|------------------------------------|---------------|--|-------------|---------|-----------|------------------|----------|
| 3 | all - table, column-family, col... | -- | column-family:* column:* table:* | Access | Enabled | -- | + | + |

For more information on exporting policies from the reports page, see links below.

Related Information

[Export Tag-Based Policies](#)

[Export Resource-Based Policies for a Specific Service](#)

[Export All Resource-Based Policies for All Services](#)

Adding a New Component to Apache Ranger

How to add a new component to Apache Ranger.

Apache Ranger has three main components:

- Admin Tool -- Provides web interface & REST API for managing security policies.
- Custom Authorization Module for components -- Provides custom authorization within the (Hadoop) component to enforce the policies defined in Admin Tool.
- UserGroup synchronizer -- Enables the user/group information in Apache Ranger to synchronize with the Enterprise user/group information stored in LDAP or Active Directory.

In order to support new component authorization using Apache Ranger, the component details need to be added to Apache Ranger as follows:

- Add component details to the Admin Tool.
- Develop a custom authorization module for the new component.

Adding Component Details to the Admin Tool

The Apache Ranger Admin tool supports policy management via both a web interface (UI) and support for a (public) REST API. In order to support a new component in both the UI and the Server, the Admin Tool must be modified.

Required UI changes to support the new component:

1. Add a new component template to the Service Manager page (console home page):

Show new component on the Service Manager page i.e home page[#!/policymanager]. Apache Ranger needs to add table template to Service Manager page and make changes in corresponding JS files. Ranger also needs to create a new service type enum to distinguish the component for which the service/policy is created/updated.

For example: Add a table template to PolicyManagerLayout_tmpl.html file to view the new component on the Access Manager page and make changes in the PolicyManagerLayout.js file related to the new component, such as passing Knox service collection data to the PolicyManagerLayout_tmpl template. Also create a new service type enum (for example, ASSET_KNOX) in the XAEnums.js file.

2. Add new configuration information to the Service Form:

Add new configuration fields to Service Form [AssetForm.js] as per new component configuration information. This will cause the display of new configuration fields in the corresponding service Create/Update page. Please note that the AssetForm.js is a common file for every component to create/update the service.

For example: Add new field(configuration) information to AssetForm.js and AssetForm_tmpl.js.

3. Add a new Policy Listing page:

Add a new policy listing page for the new component in the View Policy list. For example: Create a new KnoxTableLayout.js file and add JS-related changes as per the old component[HiveTableLayout.js] to the View Policy listing. Also create a template page, KnoxTableLayout_tmpl.html.

4. Add a new Policy Create/Update page:

Add a Policy Create/Update page for the new component. Also add a policy form JS file and its template to handle all policy form-related actions for the new component. For example: Create a new KnoxPolicyCreate.js file for Create/Update Knox Policy. Create a KnoxPolicyForm.js file to add Knox policy fields information. Also create a corresponding KnoxPolicyForm_tmpl.html template.

5. Other file changes, as needed:

Make changes in existing common files as per our new component like Router.js, Controller.js, XAUtils.js, FormInputList.js, UserPermissionList.js, XAEnums.js, etc.

Required server changes for the new component:

Let's assume that Apache Ranger has three components supported in their portal and we want to introduce one new component, Knox:

1. Create New Service Type

If Apache Ranger is introducing new component i.e Knox, then they will add one new service type for Knox. i.e serviceType = "Knox". On the basis of service type, while creating/updating service/policy, Apache Ranger will distinguish for which component this service/policy is created/updated.

2. Add new required parameters in existing objects and populate objects

For Policy Creation/Update of any component (i.e HDFS, Hive, Hbase), Apache Ranger uses only one common object, `VXPolicy`. The same goes for the Service Creation/Update of any component: Apache Ranger uses only one common object `VXService`. As Apache Ranger has three components, it will have all the required parameters of all of those three components in `VXPolicy/VXService`. But for Knox, Apache Ranger requires some different parameters which are not there in previous components. Thus, it will add only required parameters

into `VXPolicy/VXService` object. When a user sends a request to the Knox create/update policy, they will only send the parameters that are required for Knox to create/update the VXPolicy object.

After adding new parameters into VXPolixy/VXService, Apache Ranger populates the newly-added parameters in corresponding services, so that it can map those objects with Entity Object.

3. Add newly-added fields (into database table) related parameters into entity object and populate them

As Apache Ranger is using JPA-EclipseLink for database mapping into java, it is necessary to update the Entity object. For example, if for Knox policy Apache Ranger has added two new fields (`topology` and `service`) into db table `x_resource`, it will also have to update the entity object of table (i.e `XXResource`), since it is altering table structure.

After updating the entity object Apache Ranger will populate newly-added parameters in corresponding services (i.e XResourceService), so that it can communicate with the client using the updated entity object.

4. Change middleware code business logic

After adding and populating newly required parameters for new component, Apache Ranger will have to write business logic into file `AssetMgr`, where it may also need to do some minor changes. For example, if it wants to create a default policy while creating the Service, then on the basis of serviceType, Apache Ranger will create one default policy for the given service. Everything else will work fine, as it is common for all components.

Required database changes for the new component:

For service and policy management, Apache Ranger includes the following tables:

- x_asset (for service)
- x_resource (for service)

As written above, if Apache Ranger is introducing new component then it is not required to create individual table in database for each component. Apache Ranger has common tables for all components.

If Apache Ranger has three components and wants to introduce a fourth one, then it will add required fields into these two tables and will map accordingly with java object. For example, for Knox, Apache Ranger will add two fields (`topology`, `service`) into `x_resource`. After this, it will be able to perform CRUD operation of policy and service for our new component, and also for previous components.

Configuring Advanced Authorization Settings

How to customize the Ranger Advanced Settings when configuring authentication.

The screenshot shows the Ranger Admin interface with the 'Advanced' tab selected. The settings are organized into three main sections:

- Admin Settings:**
 - Ranger Admin host: dw-weekly.field.hortonworks.com
 - Ranger Admin username for Ambari: amb_ranger_admin
 - Ranger Admin user's password for Ambari: (masked)
 - Location of Sql Connector Jar: {{driver_curl_target}}
- Ranger Settings:**
 - External URL: http://dw-weekly.field.hortonworks.com:6080
 - Authentication method:
 - LDAP
 - ACTIVE_DIRECTORY
 - UNIX
 - NONE
 - HTTP enabled:
- Unix Authentication Settings:**
 - Allow remote Login: true
 - Ranger UnixAuth service: /usr/share/hadoop...

Developing a Custom Authorization Module

In the Hadoop ecosystem, each component (i.e., Hive, HBase) has its own authorization implementation and ability to plug in a custom authorization module. To implement the centralized authorization and audit feature for a component, the component should support a customizable (or pluggable) authorization module.

The custom component Authorization Plugin should do the following:

- Provide authorization based on Policies defined in Policy Admin Tool
- Provide audit information based on the authorization decisions

Implementing Custom Component Authorization

To implement the custom component authorization plugin, the Ranger common agent framework provides the following functionalities:

- Ability to read all policies from Service Manager for a given service-id
- Ability to log audit information

When the custom authorization module is initialized, the module should do the following:

1. Initiate a REST API call to the “Policy Admin Tool” to retrieve all policies associated with the specific component.
2. Once the policies are available, it should:
 - be built into a custom data structure for enabling the authorization module.

- kick off the policy updater thread to refresh policies from “Policy Admin Tool” at a regular interval.

When the custom authorization module is called to perform authorization of a component action (such as READ action) on a specific component resource (such as /app folder), the authorization module will:

- Identify authorization decision - For each policy:policyList:
 - If (resource in policy <match> auth-requested-resource)
 - If (action-in-policy <match>action-requested)
 - If (current-user or current-user-groups or public-group <allowed> for the policy), Return access-allowed
- Identify auditing needs - For each policy:policyList
 - If (resource in policy <match> auth-requested-resource), return policy.isAuditEnabled()

Special Requirements for High Availability Environments

In a High Availability (HA) environment, the primary and secondary NameNodes must be configured as described in the HDP System Administration Guide.

To enable Ranger in the HDFS HA environment, the HDFS plugin must be set up in each NameNode, and then pointed to the same HDFS service set up in the Security Manager. Any policies created within that HDFS service are automatically synchronized to the primary and secondary NameNodes through the installed Apache Ranger plugin. That way, if the primary NameNode fails, the secondary NameNode takes over and the Ranger plugin at that NameNode begins to enforce the same policies for access control.

When creating the service, you must include the `fs.default.name` property, and it must be set to the full host name of the primary NameNode. If the primary NameNode fails during policy creation, you can then temporarily use the `fs.default.name` of the secondary NameNode in the service details to enable directory lookup for policy creation.

If, while the primary NameNode is down, you wish to create new policies, there is a slight difference in user experience when specifying the resource path. If everything is normal, this is a drop-down menu with selectable paths; however, if your cluster is running from the failover node, there will be no drop-down menu, and you will need to manually enter the path.

Primary NameNode failure does not affect the actual policy enforcement. In this setup for HA, access control is enforced during primary NameNode failure by the Ranger plugs at the secondary NameNodes.

For **Test Connection** to be successful for HBase and HDFS in a Ranger HA environment, complete the following: In `/etc/ranger/admin`, create a symbolic link between `hbase-site.xml` and `hdfs-site.xml`:

```
cd /etc/ranger/admin
ln -s /etc/hadoop/conf/hdfs-site.xml hdfs-site.xml
ln -s /etc/hbase/conf/hbase-site.xml hbase-site.xml
```

Configure Advanced Usersync Settings

To access Usersync settings, select the Advanced tab on the Customize Service page. Usersync pulls in users from UNIX, LDAP, or AD and populates Ranger's local user tables with these users.

About this task

Configure advanced User Sync settings for the following:

- Unix
- (Required) LDAP/AD
- (Optional) LDAP/AD
- Automatically Assign ADMIN KEYADMIN Role for External Users

Procedure

- Unix: If you are using UNIX authentication, the default values for the Advanced ranger-ugsync-site properties are the settings for UNIX authentication:

▼ **Advanced ranger-ugsync-site**

| | | | | |
|---|---|---|---|---|
| ranger.usersync.ldap. bindkeystore | <input type="text"/> | 🔒 | 🟢 | |
| ranger.usersync.ldap. ldapbindpassword | <input type="password" value="Type password"/> <input type="password" value="Retype Password"/> | 🔒 | | |
| ranger.usersync.group. memberattributename | <input type="text"/> | 🔒 | 🟢 | 🔄 |
| ranger.usersync.group. nameattribute | <input type="text"/> | 🔒 | 🟢 | 🔄 |
| ranger.usersync.group. objectclass | <input type="text"/> | 🔒 | 🟢 | 🔄 |
| ranger.usersync.group. searchbase | <input type="text"/> | 🔒 | 🟢 | 🔄 |
| ranger.usersync.group. searchenabled | <input type="text" value="false"/> | 🔒 | 🟢 | 🔄 |
| ranger.usersync.group. searchfilter | <input type="text"/> | 🔒 | 🟢 | 🔄 |
| ranger.usersync.group. searchscope | <input type="text"/> | 🔒 | 🟢 | 🔄 |
| ranger.usersync.group. usermapsyncenabled | <input type="text" value="false"/> | 🔒 | 🟢 | 🔄 |
| ranger.usersync.ldap. searchBase | <input type="text" value="dc=hadoop,dc=apache,dc=org"/> | 🔒 | 🟢 | 🔄 |
| ranger.usersync.source. impl.class | <input type="text" value="org.apache.ranger.unixusersync.process.UnixUserGroupBuilder"/> | 🔒 | 🟢 | 🔄 |
| ranger.usersync. credstore.filename | <input type="text" value="/usr/hdp/current/ranger-usersync/conf/ugsync.jceks"/> | 🔒 | 🟢 | 🔄 |
| ranger.usersync.enabled | <input type="text" value="true"/> | 🔒 | 🟢 | 🔄 |
| ranger.usersync. filesource.file | <input type="text" value="/tmp/usergroup.txt"/> | 🔒 | 🟢 | 🔄 |
| ranger.usersync. filesource.text.delimiter | <input type="text" value="."/> | 🔒 | 🟢 | 🔄 |
| ranger.usersync. keystore.file | <input type="text" value="/usr/hdp/current/ranger-usersync/conf/unixauthservice.jks"/> | 🔒 | 🟢 | 🔄 |

- (Required) LDAP/AD
 - a) LDAP Advanced ranger-ugsync-site Settings

Table 63: LDAP Advanced ranger-ugsync-site Settings

| Property Name | LDAP Value |
|-----------------------------------|---|
| ranger.usersync.ldap.bindkeystore | Set this to the same value as the ranger.usersync.credstore.filename property, i.e, the default value is /usr/hdp/current/ranger-usersync/conf/ugsync.jceks |
| ranger.usersync.ldap.bindalias | ranger.usersync.ldap.bindalias |
| ranger.usersync.source.impl.class | ldap |

b) AD Advanced ranger-ugsync-site Settings

Table 64: AD Advanced ranger-ugsync-site Settings

| Property Name | LDAP Value |
|-----------------------------------|------------|
| ranger.usersync.source.impl.class | ldap |

- (Optional) LDAP/AD. If you are using LDAP or Active Directory authentication, you may need to update the following properties, depending upon your specific deployment characteristics.

a) Advanced ranger-ugsync-site Settings for LDAP and AD

Table 65: Advanced ranger-ugsync-site Settings for LDAP and AD

| Property Name | LDAP ranger-ugsync-site Value | AD ranger-ugsync-site Value |
|---|---|---------------------------------------|
| ranger.usersync.ldap.url | ldap://127.0.0.1:389 | ldap://ad-conrowoller-hostname:389 |
| ranger.usersync.ldap.binddn | cn=ldadmin,ou=users,dc=example,dc=com | cn=adadmin,cn=Users,dc=example,dc=com |
| ranger.usersync.ldap.ldapbindpassword | secret | secret |
| ranger.usersync.ldap.searchBase | dc=example,dc=com | dc=example,dc=com |
| ranger.usersync.source.impl.class | org.apache.ranger.ladpusersync.process.LdapUserGroupBuilder | |
| ranger.usersync.ldap.user.searchbase | ou=users, dc=example, dc=com | dc=example,dc=com |
| ranger.usersync.ldap.user.searchscope | sub | sub |
| ranger.usersync.ldap.user.objectclass | person | person |
| ranger.usersync.ldap.user.searchfilter | Set to single empty space if no value. Do not leave it as "empty" | (objectcategory=person) |
| ranger.usersync.ldap.user.nameattribute | uid or cn | sAMAccountName |
| ranger.usersync.ldap.user.groupnameattribute | memberof,ismemberof | memberof,ismemberof |
| ranger.usersync.ldap.username.caseconversion | none | none |
| ranger.usersync.ldap.groupname.caseconversion | none | none |
| ranger.usersync.group.searchenabled * | false | false |
| ranger.usersync.group.usermapsyncenabled * | false | false |
| ranger.usersync.group.searchbase * | ou=groups, dc=example, dc=com | dc=example,dc=com |
| ranger.usersync.group.searchscope * | sub | sub |
| ranger.usersync.group.objectclass * | groupofnames | groupofnames |

| Property Name | LDAP ranger-ugsync-site Value | AD ranger-ugsync-site Value |
|--|-------------------------------|--|
| ranger.usersync.group.searchfilter * | needed for AD authentication | (member=CN={0}, OU=MyUsers, DC=AD-HDP, DC=COM) |
| ranger.usersync.group.nameattribute * | cn | cn |
| ranger.usersync.group.memberattributename * | member | member |
| ranger.usersync.pagedresultsenabled * | true | true |
| ranger.usersync.pagedresultssize * | 500 | 500 |
| ranger.usersync.user.searchenabled * | false | false |
| ranger.usersync.group.search.first.enabled * | false | false |

* Only applies when you want to filter out groups.

After you have finished specifying all of the settings on the Customize Services page, click Next at the bottom of the page to continue with the installation.

- Automatically Assign ADMIN KEYADMIN Role for External Users. You can use usersync to mark specific external users, or users in a specific external group, with ADMIN or KEYADMIN role within Ranger. This is useful in cases where internal users are not allowed to login to Ranger.

a) From Ambari>Ranger>Configs>Advanced>Custom ranger-ugsync-site, select **Add Property**.

b) Add the following properties:

- ranger.usersync.role.assignment.list.delimiter = &

The default value is &.

- ranger.usersync.users.groups.assignment.list.delimiter = :

The default value is :.

- ranger.usersync.username.groupname.assignment.list.delimiter = ,

The default value is ,.

- ranger.usersync.group.based.role.assignment.rules =

ROLE_SYS_ADMIN:u:userName1,userName2&ROLE_SYS_ADMIN:g:groupName1,groupName2&ROLE_KEY_AD

c) Click Add.

d) Restart Ranger.

```
ranger.usersync.role.assignment.list.delimiter = &
ranger.usersync.users.groups.assignment.list.delimiter = :
ranger.usersync.username.groupname.assignment.list.delimiter = ,
ranger.usersync.group.based.role.assignment.rules :
&ROLE_SYS_ADMIN:u:ldapuser_12,ldapuser2
```

Related Information

[Set Up Hadoop Group Mapping for LDAP/AD](#)

Configure User Sync LDAP SSL

How to configure LDAP SSL using self-signed certs in the default Ranger User Sync TrustStore.

Procedure

- The default location is /usr/hdp/current/ranger-usersync/conf/mytruststore.jks for the ranger.usersync.truststore.file property.
- Alternatively, copy and edit the self-signed ca certs.

3. Set the `ranger.usersync.truststore.file` property to that new cacert file.

```
cd /usr/hdp/<version>/ranger-usersync
service ranger-usersync stop
service ranger-usersync start
```

Where `cert.pem` has the LDAPS cert.

Set Up Database Users Without Sharing DBA Credentials

If you do not wish to provide system Database Administrator (DBA) account details to the Ambari Ranger installer, you can use the `dba_script.py` Python script to create Ranger DB database users without exposing DBA account information to the Ambari Ranger installer. You can then run the normal Ambari Ranger installation without specifying a DBA user name and password.

Procedure

1. Download the Ranger rpm using the yum install command: `yum install ranger-admin`.
2. You should see one file named `dba_script.py` in the `/usr/hdp/current/ranger-admin` directory.
3. Get the script reviewed internally and verify that your DBA is authorized to run the script.
4. Execute the script by running the following command: `python dba_script.py`.
5. Pass all values required in the argument. These should include db flavor, JDBC jar, db host, db name, db user, and other parameters.

- If you would prefer not to pass runtime arguments via the command prompt, you can update the `/usr/hdp/current/ranger-admin/install.properties` file and then run: `python dba_script.py -q`

When you specify the `-q` option, the script will read all required information from the `install.properties` file.

- You can use the `-d` option to run the script in "dry" mode. Running the script in dry mode causes the script to generate a database script.

```
python dba_script.py -d /tmp/generated-script.sql
```

- Anyone can run the script, but it is recommended that the system DBA run the script in dry mode. In either case, the system DBA should review the generated script, but should only make minor adjustments to the script, for example, change the location of a particular database file. No major changes should be made that substantially alter the script -- otherwise the Ranger install may fail.

The system DBA must then run the generated script.

6. Run the Ranger Ambari install procedure, but set Setup Database and Database User to No in the Ranger Admin section of the Customize Services page.


Updating Ranger Admin Passwords

For certain users, if you update the passwords on the Ranger Configs page, you must also update the passwords on the Configs page of each Ambari component that has the Ranger plugin enabled.

Individual Ambari component configurations are not automatically updated -- the service restart will fail if you do not update these passwords on each component.

- Ranger Admin user -- The credentials for this user are set in Configs > Advanced ranger-env in the fields labeled `admin_username` (default value: `admin`) and `admin_password` (default value: `admin`).
- Admin user used by Ambari to create repo/policies -- The user name for this user is set in Configs > Admin Settings in the field labeled Ranger Admin username for Ambari (default value: `amb_ranger_admin`). The password for this user is set in the field labeled Ranger Admin user's password for Ambari. This password is specified during the Ranger installation.

The following image shows the location of these settings on the Ranger Configs page:



Ranger Password Requirements

This topic lists password requirements for Ranger and Ranger KMS.

Ranger user password requirements:

- Minimum of 8 characters
- Must include at least one alphabetical and one numerical character
- Must not include the following unsupported special characters: " ' \ `

Ranger and Ranger KMS DB user password requirements:

- Must not include the following unsupported special characters: " ' \ `

Ranger database instance password requirements:

- Refer to the password requirements for the applicable database type (MySQL, PostgreSQL, Oracle, etc.)