

Securing Flow Management Clusters

Date published: 2019-12-16

Date modified: 2023-06-27

The Cloudera logo is displayed in a bold, orange, sans-serif font. The word "CLOUDERA" is written in all caps, with a stylized 'E' that has a horizontal bar extending to the right.

Legal Notice

© Cloudera Inc. 2024. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

Flow Management security overview.....	4
User authorization.....	4
Assigning administrator level permissions.....	5
Assigning selective permissions to user.....	7
Assign the EnvironmentUser role.....	7
Add user to predefined Ranger access policies.....	8
Create custom access policy.....	9
Authorization example.....	14
Predefined Ranger access policies for Apache NiFi.....	14
Predefined Ranger access policies for Apache NiFi Registry.....	15

Flow Management security overview

Cloudera Data Platform (CDP) is a secure and governed hybrid data platform that offers a broad set of enterprise data services with the key data analytics and artificial intelligence functionality that enterprises require for any cloud, any analytics, any data. CDP Public Cloud is the cloud form factor of CDP. With the Cloudera Shared Data Experience (Cloudera SDX), CDP offers enterprise-grade security and governance. SDX combines enterprise-grade centralized security, governance, and management capabilities with shared metadata and a data catalog.

CDP provides the following default security features for flow management users and clusters:

- Single-sign on (SSO) authentication with Apache Knox
- Metadata management and governance capabilities with Apache Atlas
- Flow versioning and management with Apache NiFi Registry
- TLS encryption to secure communications over the network
- Fine-grained authorization to do a specific action and/or operation with Apache Ranger

For more information, see the *CDP Public Cloud Security Overview*.

Related Information

[CDP Public Cloud security overview](#)

User authorization

As a CDP administrator, after you create an environment, you must enable flow-management users to access the environment, utilize resources, and perform tasks in CDP. Assign the appropriate CDP resource role and one or more Ranger access policies based on the type of access the user requires. Before you begin, you must review the authorization workflow to understand the options and the process involved in authorizing users to access flow management resources.

Resource roles

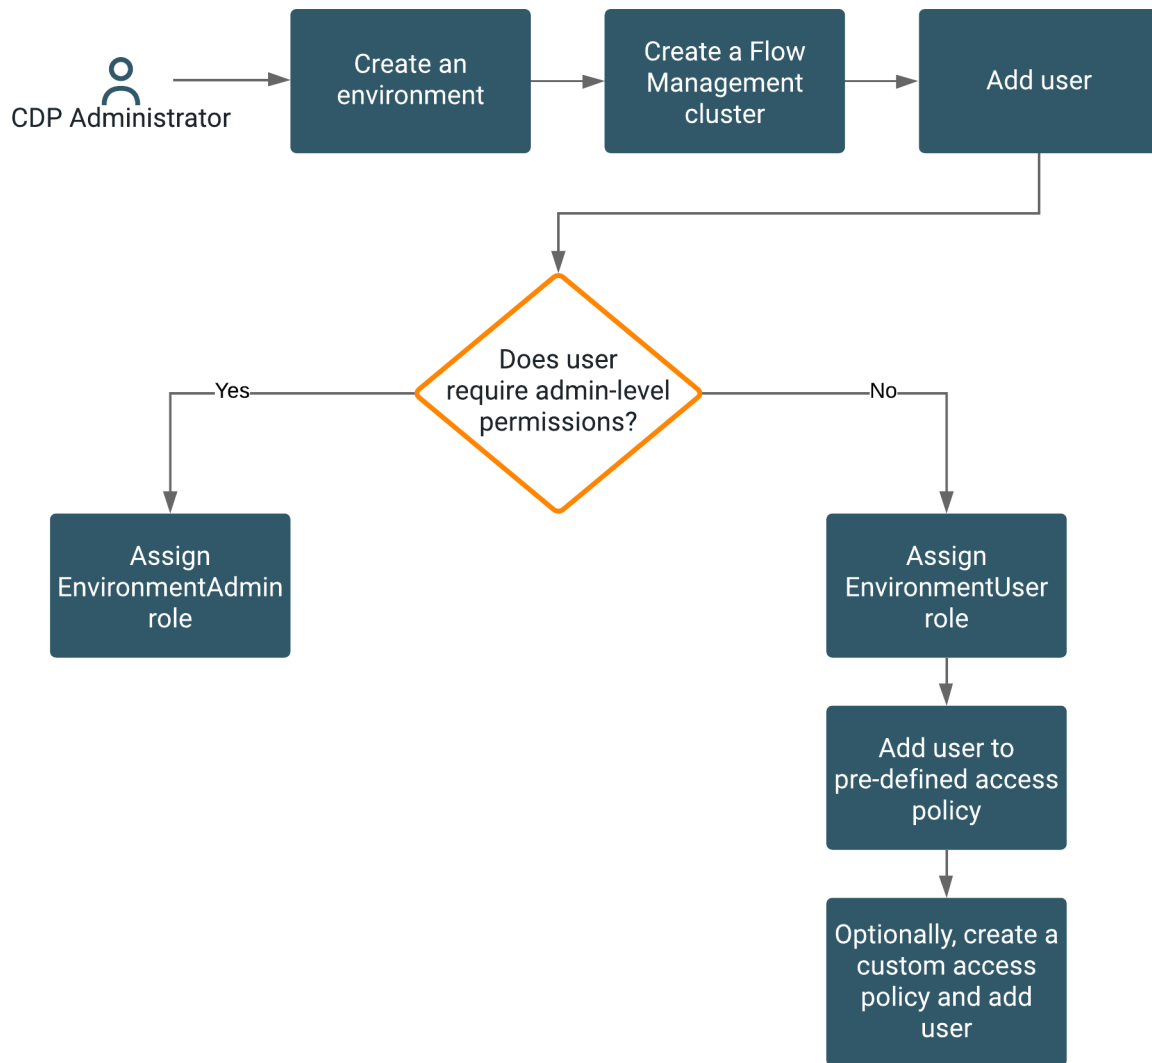
A CDP resource role grants a user or user-group permission to access and perform tasks on a resource. A Ranger access policy for flow management contains one or more access rights to NiFi or NiFi Registry resources.

To learn about CDP resource roles and Ranger policies, see the following documents:

- *Understanding roles and resource roles*
- *Ranger policies overview*

Authorization workflow

The following diagram shows the prerequisites and the workflow involved in authorizing a flow management user:



You must select one of the following options based on the privileges the user is entitled to:

- If the user requires administrator-level permissions to NiFi and NiFi Registry, assign the EnvironmentAdmin role.
- If the user requires selective permissions to NiFi and NiFi Registry, assign the EnvironmentAdmin role and add the user to the appropriate Ranger policies. Optionally, create a custom policy and add the user to it.

Related Information

[Understanding roles and resource roles](#)

[Ranger policies overview](#)

Assigning administrator level permissions

As a CDP administrator, assign the EnvironmentAdmin role to enable users to have administrator-level privileges to the environment. With the EnvironmentAdmin role, the user can access and manage environments, Flow Management clusters, and NiFi and NiFi Registry resources. Users can also authorize other users to access flow management resources.

Before you begin

Ensure that you meet the following prerequisites:

- You created a Flow Management cluster.
- You determined the permission level for each user.

Procedure

1. Click the **Environments** tab.
2. Locate the environment, and click the environment name.
3. Click **Actions Manage Access** .
The **Access** page appears.
4. Locate the user and click **Update Roles**.
The **Update Resource Role** page for the user appears.
5. Select the **EnvironmentAdmin** option.
6. Click **Update Roles**.
7. Go back to the **Environments** tab and locate the environment.
8. Click **Actions Synchronize Users to FreeIPA** .
The **Sync Users** window appears.
9. Click **Sync Users**.
This synchronizes the user to the FreeIPA identity management system to enable SSO.

Results

With the **EnvironmentAdmin** role and membership in the internal NiFi or NiFi Registry groups, the user has the ability to:

- Access and manage the environment and Flow Management clusters.
- Authorize users or groups by adding them to Ranger access policies.
- Modify or create conditions in predefined Ranger access policies.
- Create new Ranger access policies and create conditions that specify the desired level of access for each user or group.



Note: The **EnvironmentAdmin** role also gives a user the following privileges:

- Administrator rights to environments outside of NiFi.
- The right to modify predefined or custom Ranger access policies for any Ranger service in the environment.

For more information on roles, see *Understanding roles and resource roles*.

For more information about Ranger access policies, see *Ranger policies overview*.

To authorize flow management users who do not require administrator-level permission, add the users individually or as a group to specific Ranger access policies for selective access to NiFi and NiFi Registry resources. For more information, see *Assigning selective permissions to a user*.

Related Information

[Understanding roles and resource roles](#)

[Ranger Policies Overview](#)

[Assigning selective permissions to user](#)

Assigning selective permissions to user

As a CDP administrator, assign the EnvironmentUser role to enable a user to access the environment and Flow Management clusters. Then, based on the user's access requirements, add the user or a group of users to the appropriate Ranger access policies for NiFi and NiFi Registry.

Perform the following steps to authorize access for a new user:

1. Assign the EnvironmentUser role.
2. Add the user or group to the appropriate pre-defined Ranger access policies.
3. Create a custom Ranger access policy and add the user or group.

Assign the EnvironmentUser role

For flow management users who do not need administrator privileges, you must assign the EnvironmentUser role. This role enables users to set their password and access the environment.

Before you begin

Ensure that you meet the following prerequisites:

- You created a Flow Management cluster.
- You determined the permission level for each user.

Procedure

1. From the Cloudera Management console, go to the Environments tab.
2. Use the Search bar to find the environment.
3. Click the environment name.
4. Click **Actions** **Manage Access** .
The **Access** page appears.
5. Locate the user and click **Update Roles** .
The **Update Resource Role** page for the user appears.
6. Select the EnvironmentUser option.
7. Click **Update Roles** .
8. Go back to the **Environments** tab and locate the environment.
9. Click **Actions** **Synchronize Users to FreeIPA** .
The **Sync Users** window appears.
10. Click **Sync Users** .
This synchronizes the user to the FreeIPA identity management system to enable SSO.

Results

The user is added to the environment and can access the environment and Flow Management clusters. You can now add the user or a group of users to Ranger policies that allow access to NiFi and NiFi Registry resources.

What to do next

Complete the steps listed in *Add the user to predefined Ranger access policies*.

Related Information

[Add user to predefined Ranger access policies](#)

Add user to predefined Ranger access policies

After assigning the EnvironmentUser role to the new user, you must add the user to the appropriate predefined Ranger access policies for NiFi and NiFi Registry. These policies determine what the user can command, control, and observe in a NiFi dataflow or in the NiFi Registry.

About this task

Each predefined Ranger access policy confers specific rights to NiFi or NiFi Registry resources. When an authenticated flow-management user attempts to view or modify a NiFi or NiFi Registry resource, the system checks whether the user is associated with the specific Ranger access policy that confers the privileges to perform that action.

For more information, see:

- *Predefined Ranger access policies for NiFi resources*
- *Predefined Ranger access policies for NiFi Registry resources*

Procedure

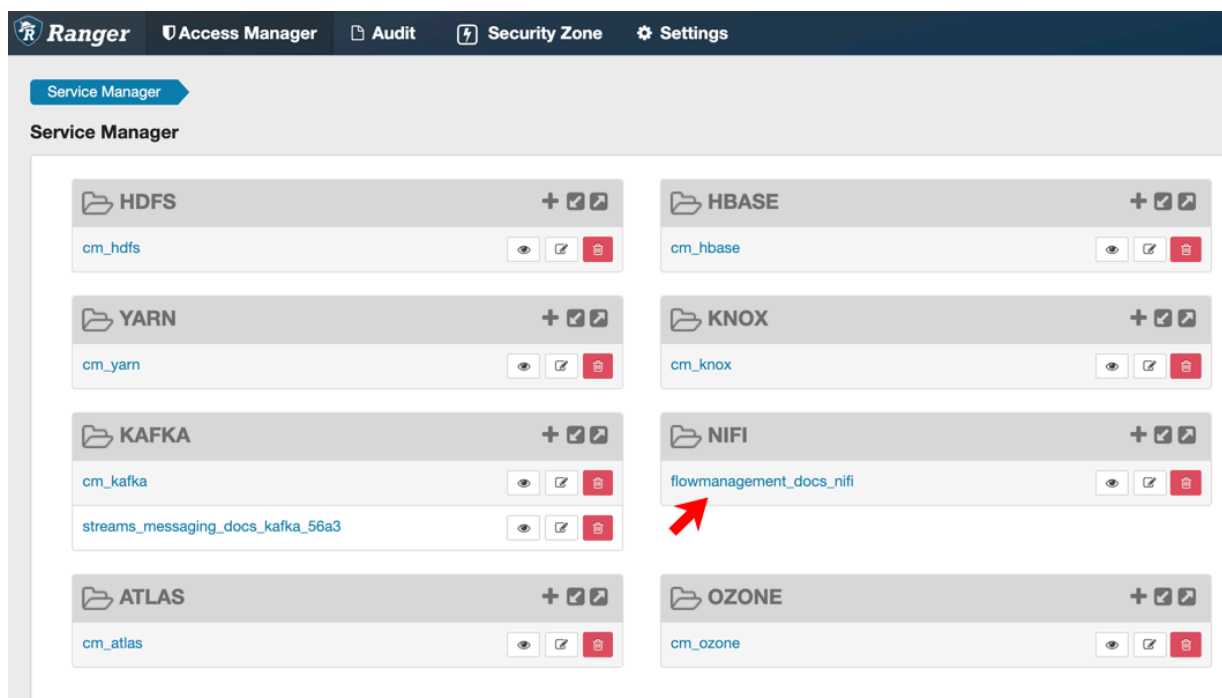
1. From the base cluster with Ranger, click the Ranger icon.

The **Ranger Service Manager** page appears. Each cluster in the environment is listed under its respective service. For example, the NiFi clusters in the environment are listed under NiFi.

2. Select a cluster from either the NiFi or NiFi Registry section.

The **List of Policies** page appears.

The following image shows the list of predefined policies for NiFi:



- Click the ID for a policy.
The **Edit Policy** page appears.

The following image shows the list of predefined policies for NiFi:

Policy ID	Policy Name	Policy Labels	Status	Audit Logging	Roles	Groups	Users	Action
52	all - nifi-resource	--	Enabled	Enabled	--	_c_ranger_admins_a44480d	rangerlookup	
53	Restricted Components	--	Enabled	Enabled	--	c_nifi_admins_a44480d	--	
54	Tenants	--	Enabled	Enabled	--	_c_nifi_admins_a44480d	--	
55	Controller	--	Enabled	Enabled	--	_c_nifi_admins_a44480d	--	
56	Flow	--	Enabled	Enabled	--	c_nifi_admins_a44480d	--	
57	Policies	--	Enabled	Enabled	--	_c_nifi_admins_a44480d	--	
58	Proxies	--	Enabled	Enabled	--	nifi	--	
66	Root Process Group	--	Enabled	Enabled	--	c_nifi_admins_a44480d	--	
67	Root Group Data	--	Enabled	Enabled	--	nifi _c_nifi_admins_a44480d	--	

- In the Allow Conditions section, add the user or the user group to the Select User field.
- Click Save.

Results

The user now has the NiFi and NiFi Registry rights according to the policies you added the user or user group to. These rights are inherited down the hierarchy unless there is a more specific policy on a component.

What to do next

Complete the steps listed in *Create a custom access policy*.

Related Information

[Predefined Ranger access policies for Apache NiFi](#)

[Predefined Ranger access policies for Apache NiFi Registry](#)

[Create custom access policy](#)

Create custom access policy

If the user cannot access the component through an inherited Ranger access policy, you must create a custom Ranger access policy for the specific component and add the user to this policy. If all the users in a group require the same access, you can add the user group to the Ranger access policy.

About this task

A user might need access to specific NiFi or NiFi Registry resources such as a processor, processor group, remote process group, funnel, label, controller service, or bucket. Each custom Ranger access policy provides access to a specific component.

First determine which NiFi or NiFi Registry components a user needs access to. Then create a new policy for each component and add the user or user group to the new policy.

When you create a new policy, you must specify the ID of the component that the user requires access to.

**Note:**

If a user requires permission to view or modify data for a specific component, you must create a custom data access policy and add the user and the nifi group to that policy.

The nifi group is a dynamically-managed group that exists on all Flow Management Data Hub hosts and contains the identities of NiFi and Knox nodes. When you add the nifi group to the data policy for a specific component, you authorize the nodes to access data on behalf of the user.

Procedure

1. From the NiFi canvas, copy the ID of the process group, SSL Context Service, or controller service for reporting tasks that the user needs access to.
2. Locate the ID for a process group:
 - a) Click the process group.

The ID appears in the **Operate** pane.




- b) Copy the ID.

3. Locate the ID of the SSL Context Service:
 - a) Click the settings icon on the process group.
The **NiFi Flow Configuration** appears.
 - b) Click the **Controller Services** tab.
 - c) Click the **Settings** icon for the Default NiFi SSL Context Service.
The **Controller Service Details** window appears.
 - d) From the **Settings** tab, copy the ID from the Id field.

The screenshot displays the NiFi Flow Configuration interface. The 'CONTROLLER SERVICES' tab is active, showing a table with one entry: 'Default NiFi SSL Context Service'. A 'Controller Service Details' window is open, showing the 'SETTINGS' tab. The 'Id' field is highlighted with a red box, containing the value '7e8da72c-0171-1000-0000-000003ebdbf4'. The 'Name' field is 'Default NiFi SSL Context Service'. The 'Type' is 'StandardRestrictedSSLContextService 1.11.4.2.0.0-212'. The 'Bundle' is 'org.apache.nifi - nifi-ssl-context-service-nar'. The 'Supports Controller Service' section lists dependencies: 'RestrictedSSLContextService 1.11.4.2.0.0-212 from org.apache.nifi - nifi-standard-services-api-nar' and 'SSLContextService 1.11.4.2.0.0-212 from org.apache.nifi - nifi-standard-services-api-nar'. The 'Referencing Components' section is empty. The 'OK' button is visible at the bottom right of the details window. At the bottom left of the main interface, it says 'Last updated: 21:13:16 UTC'. At the bottom right, it says 'Listed services are available to all descendant Processors and services of this Process Group.'

4. Locate the ID of a controller service for reporting tasks:
 - a) Click the process group.
 - b) Click the menu on the top right of the UI and select Controller Settings.



The **NiFi Settings** page appears.

- c) Click the **Reporting Tasks Controller Services** tab.
- d) Click the Settings icon for the controller service.



The **Controller Service Details** page appears.

- e) From the **Settings** tab, copy the ID from the Id field.

Controller Service Details

SETTINGS
PROPERTIES
COMMENTS

Name
Default Reporting Task SSL Context Service

Referencing Components ⓘ
No referencing components.

Id
05ad168c-0171-1000-ffff-ffffe39561d8

Type
StandardSSLContextService 1.11.3.2.0.0.0-195

Bundle
org.apache.nifi - nifi-ssl-context-service-nar

Supports Controller Service

- SSLContextService 1.11.3.2.0.0.0-195 from org.apache.nifi - nifi-standard-services-api-nar

OK

5. Go back to the **Ranger List of Policies** page.
6. Click Add New Policy.

The screenshot shows the Ranger Access Manager interface. At the top, there's a navigation bar with 'Ranger', 'Access Manager', 'Audit', 'Security Zone', and 'Settings'. Below that, a breadcrumb trail shows 'Service Manager' > 'docs_flowm_nifi Policies'. The main heading is 'List of Policies : docs_flowm_nifi'. There's a search bar with the placeholder 'Search for your policy...'. To the right of the search bar, the 'Add New Policy' button is highlighted with a red box. Below the search bar is a table with the following columns: Policy ID, Policy Name, Policy Labels, Status, Audit Logging, Roles, Groups, Users, and Action. The table contains 10 rows of policies, all with 'Enabled' status and 'Enabled' audit logging. The 'Action' column for each row contains three icons: an eye (view), a pencil (edit), and a trash can (delete).

Policy ID	Policy Name	Policy Labels	Status	Audit Logging	Roles	Groups	Users	Action
52	all - nifi-resource	--	Enabled	Enabled	--	_c_ranger_admins_a44480d	rangerlookup	
53	Restricted Components	--	Enabled	Enabled	--	_c_nifi_admins_a44480d	--	
54	Tenants	--	Enabled	Enabled	--	_c_nifi_admins_a44480d	--	
55	Controller	--	Enabled	Enabled	--	_c_nifi_admins_a44480d	--	
56	Flow	--	Enabled	Enabled	--	_c_nifi_admins_a44480d	--	
57	Policies	--	Enabled	Enabled	--	_c_nifi_admins_a44480d	--	
58	Proxies	--	Enabled	Enabled	--	nifi	--	
66	Root Process Group	--	Enabled	Enabled	--	_c_nifi_admins_a44480d	--	
67	Root Group Data	--	Enabled	Enabled	--	nifi _c_nifi_admins_a44480d	--	

The **Create Policy** page appears.

7. Enter a unique name for the policy.
8. Optionally, enter a keyword in the Policy Label field to aid in searching for a policy.
9. Enter the resource descriptor and the resource ID in the NiFi Resource Identifier or NiFi Registry Resource Identifier field in the following format: <resource descriptor><resource ID>
 - To determine a NiFi resource descriptor, see *Predefined Ranger access policies for Apache NiFi*.
 - To determine a NiFi Registry resource descriptor, see *Predefined Ranger access policies for Apache NiFi Registry*.
10. Enter a description.

11. Add a user or a group.

Note: If a user requires permission to view or modify the data for a specific component, you must create a data policy with `/data/<component-type>/<component-UUID>` as the resource identifier. Then add the user and the nifi group to the policy to authorize the NiFi and Knox nodes to access data on behalf of the user.

12. Set the permission level for the user or group.**13.** Click Add.**Results**

The user or group of users can now access the component specified in the custom policy.

Related Information

[Predefined Ranger access policies for Apache NiFi](#)

[Predefined Ranger access policies for Apache NiFi Registry](#)

Authorization example

You can review an example to understand how you can enable a flow-management user to perform specific tasks like setting up version control for a flow, by assigning the appropriate Ranger policies.

UserA must be able to do the following tasks:

- Access the NiFi UI.
- Export a flow.
- View data queued in connections.
- View data flowing through.
- Use a NiFi SSLContextService to connect to SSL-enabled systems.
- Set up version control for a flow.

Complete the following steps to enable UserA to perform the required tasks:

1. Add UserA to the predefined Ranger access policy for NiFi, Flow. Set the permissions to Read.

The Flow policy gives the user the right to view the NiFi UI.

2. Create a Ranger access policy for NiFi with:

- Resource descriptor: `/data/process-groups/<ID of process-group>`
- Permission: Read and Write

Add UserA to this custom policy. The policy gives the user the right to export the data, view the data that is queued and flowing through the connections.

3. Create a Ranger access policy for NiFi with:

- Resource descriptor: `/controller-service/<ID of SSL Context Service>`
- Permission: Read

Add UserA to this custom policy. The policy gives the user the right to use the specified SSLContextService in their flows to connect to SSL-enabled systems.

4. Create a Ranger access policy for NiFi Registry with:

- Resource descriptor: `/buckets/<ID of bucket>`
- Permission: Read, Write, and Delete

Add UserA to this custom policy. The policy gives the user the right to set up version control for a flow.

Predefined Ranger access policies for Apache NiFi

You can review the predefined Ranger policies for NiFi to determine the appropriate policy to assign to a user.

The following table lists the predefined Ranger access policies for NiFi. If you create a custom policy, refer to the Resource Descriptor column in this table to enter the value in the NiFi Resource Identifier field on the **New Policy** page.



Important:


Do not rename the default policies as some cluster operations rely on these policy names.

Do not select the Delegate Admin checkbox.



Note: The NiFi and Knox nodes have permission to the following Ranger policies:

- Proxies; /proxy
- Root Group Data; /data/process-groups

Ranger Policy	Description	Resource Descriptor
Controller	Allows users to view and modify the controller including Reporting Tasks, Controller Services, Parameter Contexts and Nodes in the Cluster.	/controller
Flow	Allows users to view the NiFi UI.	/flow
Policies	Allows users to view the policies for all components.	/policies
Provenance	Allows users to submit a Provenance Search and request Event Lineage.	/provenance
Proxies	Allows NiFi and Knox hosts to proxy user requests. Does not apply to users or user groups.	/proxy
Restricted Components	<p>Allows users to create/modify restricted components assuming other permissions are sufficient.</p> <p>The restricted components may indicate the specific permissions that are required.</p> <p>Permissions can be granted for specific restrictions or be granted regardless of restrictions. If permission is granted regardless of restrictions, the user can create/modify all restricted components.</p> <p>Some examples of restricted components are ExecuteScript, List/FetchHDFS, and TailFile.</p>	<p>/restricted-components</p> <p>See the <i>NiFi Restricted Components</i> topic for information on the sub-policies.</p>
Root Group Data	<p>Allows users and the nifi group to view and delete data from the root group and down the hierarchy unless there is a more specific policy on a component.</p> <p> Note: The nifi group is a dynamically managed list of Knox and NiFi node identities. The group exists on all Data Hub Flow Management hosts.</p>	/data/process-groups/<uuid>
Root Group Provenance Data	Allows users to view provenance data.	/provenance-data/process-groups/
Root Process Group	<p>Allows users to view and modify the root process group including adding/removing processors to the canvas.</p> <p>This policy is inherited down the hierarchy unless there is a more specific policy on a component.</p>	/process-groups/<uuid>
Tenants	Allows users to view and modify user accounts and user groups.	/tenants

Predefined Ranger access policies for Apache NiFi Registry

You can review the predefined Ranger policies for NiFi Registry to determine the appropriate policy to assign to a user.

The following table lists the pre-defined Ranger access policies for NiFi Registry. If you create a custom policy, refer to the Resource Descriptor column in this table to enter the value in the NiFi Registry Resource Identifier field on the **New Policy** page.



Important:

Do not rename the default policies as some cluster operations rely on these policy names.

Do not select the Delegate Admin checkbox.



Note: The NiFi Registry and Knox nodes have permission to the Proxies (/proxy) Ranger policy.

Ranger Policy	Description	Resource Descriptor
Actuator	Allows users to access the Spring Boot Actuator end-points.	/actuator
Buckets	Allows users to view and modify all buckets.	/buckets
Policies	Allows users to view the policies for all components.	/policies
Proxies	Allows NiFi Registry and Knox hosts to proxy user requests. Does not apply to users or user groups.	/proxy
Swagger	Allows users to access the self-hosted Swagger UI.	/swagger
Tenants	Allows users to view and modify user accounts and user groups.	/tenants