

Cloudera Flow Management 2.1.7

Cloudera Flow Management Installation

Date published: 2019-06-26

Date modified: 2024-06-03

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.

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

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

CFM installation workflow.....	4
Scenario 1: NiFi and NiFi Registry on a Base cluster.....	5
Scenario 2: NiFi and NiFi Registry on a Compute cluster.....	5
Scenario 3: NiFi on a Compute cluster and NiFi Registry on a Base cluster.....	6
Preparing your CDP Private Cloud Base cluster.....	7
Installing the JDK.....	7
Installing and configuring a database for NiFi Registry.....	7
Installing MySQL.....	8
Configuring NiFi Registry Metadata Stores in MySQL.....	9
Installing PostgreSQL.....	9
Configuring NiFi Registry Metadata Stores in PostgreSQL.....	11
Installing Cloudera Manager and a CDP Private Cloud Base cluster.....	12
Installing the CFM parcel from the repository.....	13
Downloading Custom Service Descriptor files.....	14
Installing NiFi and NiFi Registry on your Base cluster.....	14
Adding the NiFi service.....	15
Adding the NiFi Registry service.....	16
Connecting NiFi to NiFi Registry.....	17
Adding users or groups to Ranger policies.....	18
Deselecting unwanted NiFi Registry dependencies.....	21
Installing NiFi and NiFi Registry on a Compute cluster.....	21
Creating a Shared Data Context.....	22
Adding the NiFi and NiFi Registry groups to Ranger in the Base cluster.....	23
Creating a Compute cluster.....	23
Adding the NiFi and NiFi Registry services to a Compute cluster.....	25
Connecting NiFi to NiFi Registry.....	26
Adding users or groups to Ranger policies.....	27
Deselecting unwanted NiFi Registry dependencies.....	29
Installing NiFi on a Compute cluster and NiFi Registry on a Base cluster.....	30
Installing NiFi Registry on your Base cluster.....	31
Creating a Shared Data Context.....	32
Adding the NiFi and NiFi Registry groups to Ranger in the Base cluster.....	33
Creating Compute cluster.....	33
Adding NiFi to the Compute cluster.....	34
Setting up NiFi node users in NiFi Registry.....	36
Connecting NiFi to NiFi Registry.....	36
Adding users or groups to Ranger policies.....	37
Deselecting unwanted NiFi Registry dependencies.....	39

CFM installation workflow

Before you start installing Cloudera Flow Management (CFM) on a CDP Private Cloud cluster, it is helpful to understand the steps involved. There are three main installation scenarios depending on the type of cluster or clusters on which you install the different components.

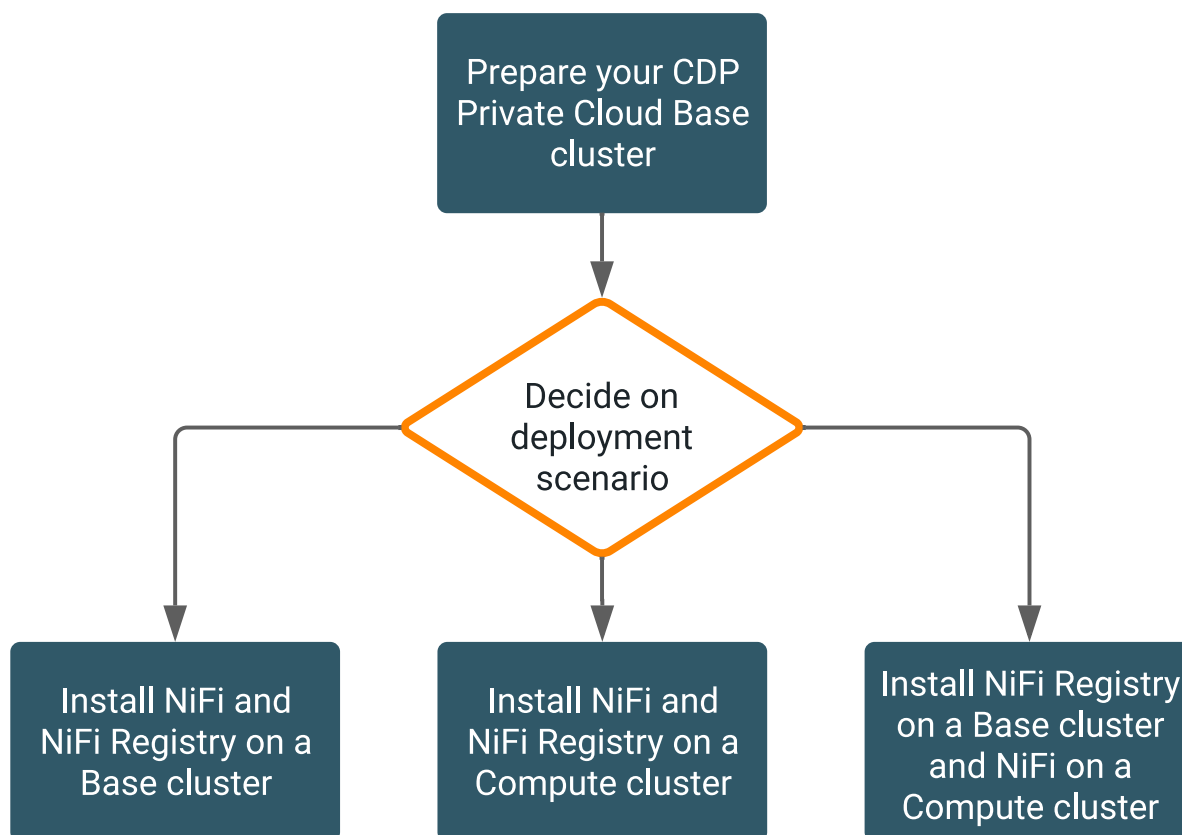
The Cloudera Flow Management (CFM) installation workflow consists of preparing your CDP Private Cloud Base cluster and then installing NiFi and NiFi Registry on a Base or on a Compute cluster, or on a mixed cluster layout. See the following diagram for the procedure steps, and review the description of each deployment scenario following the diagram.



Important: Running multiple NiFi instances on the same node is not recommended by Cloudera.



Tip: The following workflow diagram is clickable. Click each step to go directly to the relevant documentation.



[rect 237, 21, 402, 121 Link to documentation for preparing your CDP Private Cloud Base](#)

[rect 21, 338, 186, 440 Link to documentation for installing NiFi and NiFi Registry on a Base cluster](#)

[rect 237, 338, 402, 440 Link to documentation for installing NiFi and NiFi Registry on a Compute cluster](#)

[rect 453, 338, 619, 440 Link to documentation for installing NiFi on a Compute cluster and NiFi Registry on a Base cluster](#)

Scenario 1: NiFi and NiFi Registry on a Base cluster

Installing NiFi and NiFi Registry on your Base cluster is a simplified deployment during which you create only one cluster. You can do this if you are setting up in trial deployments or in simple production scenarios.

A Base cluster is a large cluster that contains SDX services like Ranger and Atlas, as well as compute and storage services like NiFi, Impala, and Hive. It is sometimes called a Regular cluster in Cloudera Manager, and it is also known as an SDX cluster, a Shared cluster, or a Data Lake cluster.

You can install NiFi Registry on the Base cluster. If you are using multiple NiFi services, you would likely want to share the NiFi Registry instance across all of these NiFi services.



Important: Cloudera recommends that you always install NiFi Registry on the Base cluster.

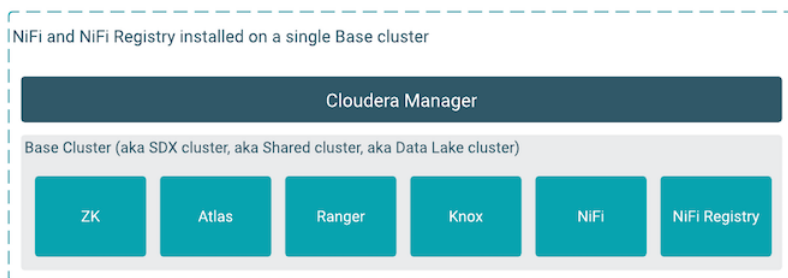
In simplified deployment scenarios, you can also install NiFi on the Base cluster to achieve a single-cluster setup. With this scenario, you may have many services on the Base cluster relying on the same Zookeeper quorum, which may not suit your use case depending on the expected workload.



Important: Starting with CFM 2.1.5, you are able to install multiple instances of the NiFi Service on the same cluster. These instances share the same Zookeeper quorum, so you would need to consider the expected workload.

For information on the installation steps, see [Installing NiFi and NiFi Registry on your Base cluster](#).

When you install both NiFi and NiFi Registry on your Base cluster, your cluster layout will look similar to the following:



Scenario 2: NiFi and NiFi Registry on a Compute cluster

Although it is not a recommended setup, it is possible to install both NiFi and NiFi Registry on a Compute cluster.

A Compute cluster consists of compute nodes only. The Compute cluster is then connected to a Base cluster with SDX services like Ranger and Atlas using a Shared Data Context.



Important: Cloudera recommends that you always install NiFi Registry on the Base cluster.

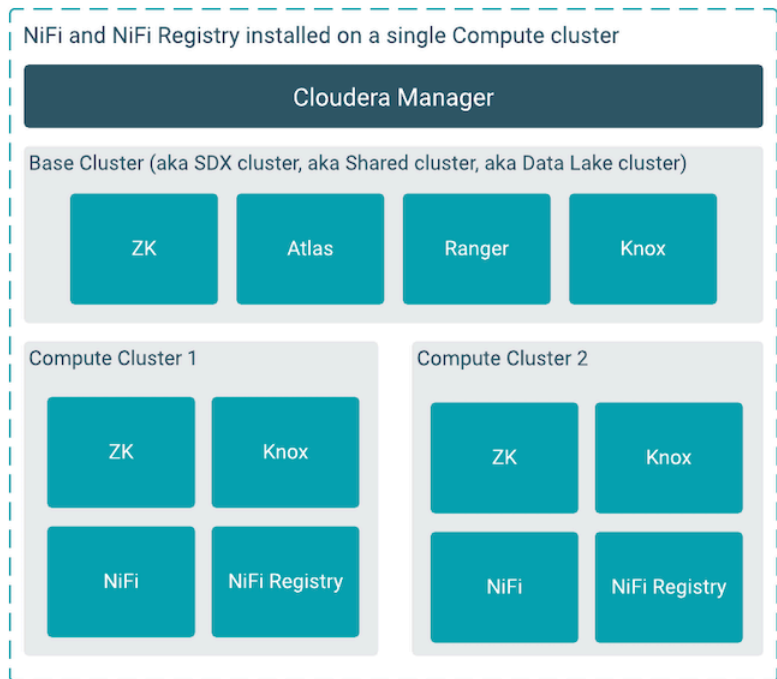
Installing NiFi on a Compute cluster is recommended when you want to:

- deploy multiple NiFi clusters for larger scale production deployments
- isolate your dataflows
- have different authorization models per dataflow

and in other similar deployment scenarios.

For information on the installation steps, see [Installing NiFi and NiFi Registry on a Compute cluster](#).

When you install both NiFi and NiFi Registry on a Compute cluster, your cluster layout will look similar to the following:



Scenario 3: NiFi on a Compute cluster and NiFi Registry on a Base cluster

For production deployments, it is recommended to install NiFi on one or more Compute clusters, and to install NiFi Registry on the Base cluster. This provides a dedicated Zookeeper quorum for each NiFi cluster. Use this setup when you want to share dataflows across multiple NiFi Compute clusters in your CDP environment.




Note: Cloudera recommends that you use this layout for production deployments.



Important: When you have multiple NiFi instances managed by the same Cloudera Manager instance, it is required to upgrade all NiFi clusters at the same time as it is currently not possible to have multiple versions of CSDs at the same time in Cloudera Manager.

For information on the installation steps, see [Installing NiFi on a Compute cluster and NiFi Registry on a Base cluster](#).

When you install NiFi on one or more Compute clusters and NiFi Registry on the Base cluster, your cluster layout will look similar to the following:



Preparing your CDP Private Cloud Base cluster

Before you begin installing NiFi and NiFi Registry, you need to prepare your CDP Private Cloud Base cluster.

Installing the JDK

You should install JAVA Development Kit (JDK) on each machine on which you will install NiFi, if your version of Cloudera Manager is not already configured with JDK.

Procedure

1. Download JDK from the appropriate website.
For the list of supported JDKs, see *System requirements*.
2. Run the installation command appropriate for your operating system:

For RHEL/CentOS:

```
yum install java-1.8.0-openjdk-devel
```

What to do next

Once you have installed the JDK, ensure that you have the required database installed and set up for NiFi Registry.

Related Information

[System requirements](#)

[Installing and configuring a database for NiFi Registry](#)

[Open JDK Download](#)

[Oracle JDK Download](#)

Installing and configuring a database for NiFi Registry

If you are installing NiFi Registry, you can install and configure an external MySQL or PostgreSQL database. By default, NiFi Registry is preconfigured to use an embedded H2 database.

H2 is an embedded database that is preconfigured in the default `nifi.registry.properties` file. The contents of the H2 database are stored in a file on your local file system. The H2 database location is specified as part of the JDBC URL property:

- NiFi Registry JDBC Url (nifi.registry.db.url) – jdbc:h2:./database/nifi-registry-primary

If you plan to use the H2 embedded database, you can skip the steps for installing and configuring an external database. Alternately, you may install a MySQL or PostgreSQL external database. To do this, review the following steps for either MySQL or PostgreSQL.

For a list of supported MySQL or PostgreSQL databases, see [Supported NiFi Registry databases](#).

Installing MySQL

Learn how to install a MySQL external database for NiFi Registry.

About this task

If you have already installed a PostgreSQL database, you may skip these steps. Both databases are not required.

For a list of supported MySQL databases, see [Supported NiFi Registry databases](#).

Before you begin

You have installed the JDK. For details, see *Installing the JDK*.

Procedure

1. Log in to the node on which you want to install NiFi Registry.
2. Install MySQL and the MySQL community server, and start the MySQL service:

```
yum localinstall https://dev.mysql.com/get/mysql[***MYSQL_VERSION***]-community-release-el[***RHEL_VERSION***]-[***PACKAGE_VERSION***].noarch.rpm
yum install mysql-community-server
systemctl start mysqld.service
```

You can replace the placeholders with the appropriate values for your specific use case:

- [***MYSQL_VERSION***]: Replace this with the desired MySQL version.
- [***RHEL_VERSION***]: Replace this with the Red Hat Enterprise Linux (RHEL) version you are using.
- [***PACKAGE_VERSION***]: Replace this with the specific package version.

For example, if you want to install the MySQL 8.0 Community Edition repository configuration package on a RHEL 8 system, you would use:

```
yum localinstall https://dev.mysql.com/get/mysql80-community-release-el8-8.noarch.rpm
```

3. Obtain the randomly generated MySQL root password:

```
grep 'A temporary password is generated for root@localhost' \
/var/log/mysqld.log |tail -1
```

4. Reset the MySQL root password.

Enter the following command:

```
/usr/bin/mysql_secure_installation
```

You are then prompted for the password you obtained in the previous step. MySQL then asks you to change the password.

What to do next

Once you have completed the MySQL installation, configure it for use with NiFi Registry.

Related Information

[Installing the JDK](#)

[Configuring NiFi Registry Metadata Stores in MySQL](#)

Configuring NiFi Registry Metadata Stores in MySQL

Learn how to configure a MySQL database for use with NiFi Registry.

About this task

MySQL provides the option to use an externally located database that supports high availability.

Before you begin

You have installed a MySQL database. For details, see *Installing MySQL*.

Procedure

1. Download the MySQL JDBC driver and place it somewhere accessible to NiFi Registry:

```
/path/to/drivers/mysql-connector-java-8.0.16.jar
```

2. Create a database inside MySQL (enter mysql shell using `mysql -u root -p`):

```
CREATE DATABASE nifi_registry;
```

3. Create a database user and grant privileges (for remote users, use `nifireg'@'<IP-ADDRESS>` or `nifireg'@'%'` for any remote host):

```
GRANT ALL PRIVILEGES ON nifi_registry.* TO 'nifireg'@'localhost' IDENTIFIED BY 'changeme';
```

4. After NiFi Registry service is installed, configure the database properties in Cloudera Manager:

- NiFi Registry JDBC Url (`nifi.registry.db.url`) – `jdbc:mysql://<MYSQL-HOSTNAME>/nifi_registry`
- NiFi Registry JDBC Driver (`nifi.registry.db.driver.class`) – `com.mysql.cj.jdbc.Driver`
- NiFi Registry H2 directory storage location (`nifi.registry.db.driver.directory`) – `/path/to/drivers`



Note:

The NiFi Registry H2 directory storage location field specifies the NiFi Registry database driver directory. The H2 database is used by default. Update this field when you are configuring it for an external database.

- NiFi Registry Database Username (`nifi.registry.db.username`) – `nifireg`
- NiFi Registry Database Password (`nifi.registry.db.password`) – `changeme`

What to do next

When you have completed the NiFi Registry database configuration, move on to installing Cloudera Manager and your CDP Private Cloud Base Cluster.

Related Information

[Installing MySQL](#)

[Installing Cloudera Manager and a CDP Private Cloud Base cluster](#)

Installing PostgreSQL

Learn how to install a PostgreSQL external database for use with NiFi Registry.

About this task

If you have already installed a MySQL database, you may skip these steps. Both databases are not required.

For a list of supported PostgreSQL databases, see [Supported NiFi Registry databases](#).

Before you begin

You have installed the JDK. For details, see *Installing the JDK*.

Procedure

1. Install Red Hat Package Manager (RPM) according to the requirements of your operating system:

```
yum install https://yum.postgresql.org/[***POSTGRESQL_VERSION**]/redhat/rhel-[***RHEL_VERSION**]-x86_64/[***REPOSITORY_PACKAGE_NAME**]
```

You can replace the placeholders with the appropriate values for your specific use case:

- **[***POSTGRESQL_VERSION**]**: Replace this with the desired PostgreSQL version.
- **[***RHEL_VERSION**]**: Replace this with the Red Hat Enterprise Linux (RHEL) version you are using.
- **[***REPOSITORY_PACKAGE_NAME**]**: Replace this with the name of the repository configuration package for the specific PostgreSQL version and RHEL version you are targeting.

For example, if you want to add the PostgreSQL 12 repository to a RHEL 8 system, you would use:

```
yum install https://yum.postgresql.org/12/redhat/rhel-8-x86_64/pgdg-redhat12-12-4.noarch.rpm
```

2. Install the PostgreSQL version of your choice:

```
yum install postgresql[***POSTGRESQL_VERSION**]-server postgresql[***POSTGRESQL_VERSION**]-contrib postgresql[***POSTGRESQL_VERSION**]
```

For example, if you wanted to add the PostgreSQL 12, you would use:

```
yum install postgresql12-server postgresql12-contrib postgresql12
```

3. Initialize the database.

For CentOS 7, use the following syntax:

```
[***POSTGRESQL_BIN_PATH**]/postgresql[***POSTGRESQL_VERSION**]-setup initdb
```

For example, if you wanted to initialize PostgreSQL 12, you would use:

```
/usr/pgsql-12/bin/postgresql12-setup initdb
```

4. Start PostgreSQL.

```
systemctl enable postgresql-[***POSTGRESQL_VERSION**].service  
systemctl start postgresql-[***POSTGRESQL_VERSION**].service
```

For example, if you are using CentOS 7, use the following syntax:

```
systemctl enable postgresql-12.service  
systemctl start postgresql-12.service
```

5. Verify that you can log in:

```
sudo su postgres
```

```
psql
```

What to do next

Once you have installed PostgreSQL, configure the database for use with NiFi Registry.

Related Information

[Installing the JDK](#)

[Configuring NiFi Registry Metadata Stores in PostgreSQL](#)

Configuring NiFi Registry Metadata Stores in PostgreSQL

Learn how to configure a PostgreSQL database for use with NiFi Registry.

About this task

Postgres provides the option to use an externally located database that supports high availability.

Before you begin

You have installed a PostgreSQL database. For details, see *Installing PostgreSQL*.

Procedure

1. Download the Postgres JDBC driver and place it somewhere accessible to NiFi Registry:

```
/path/to/drivers/postgresql-42.2.2.jar
```

2. Create a database inside Postgres:

```
createdb nifireg
```

3. Create a database user and grant privileges:

```
psql nifireg
CREATE USER nifireg WITH PASSWORD 'changeme';
GRANT ALL PRIVILEGES ON DATABASE nifireg to nifireg;
\q
```

4. After NiFi Registry service is installed, configure the database properties in Cloudera Manager:

- NiFi Registry JDBC Url (nifi.registry.db.url) – jdbc:postgresql://<POSTGRES-HOSTNAME>/nifireg
- NiFi Registry JDBC Driver (nifi.registry.db.driver.class) – org.postgresql.Driver
- NiFi Registry H2 directory storage location (nifi.registry.db.driver.directory) – /path/to/drivers



Note:

The NiFi Registry H2 directory storage location field specifies the NiFi Registry database driver directory. The H2 database is used by default. Update this field when you are configuring it for an external database.

- NiFi Registry Database Username (nifi.registry.db.username) – nifireg
- NiFi Registry Database Password (nifi.registry.db.password) – changeme

What to do next

When you have completed the NiFi Registry database configuration, move on to installing Cloudera Manager and your CDP Private Cloud Base Cluster.

Related Information

[Installing PostgreSQL](#)

[Installing Cloudera Manager and a CDP Private Cloud Base cluster](#)

Installing Cloudera Manager and a CDP Private Cloud Base cluster

Before installing NiFi and NiFi Registry, you need to install CDP Private Cloud Base with Cloudera Manager and certain Runtime services.

About this task

You should follow the instructions in the *CDP Private Cloud Base Installation Guide* for complete information about installing Cloudera Manager and a CDP base cluster. At minimum, you should ensure that you perform the following steps.

Before you begin

- You have installed a JDK.
- If you want to use an external database, you have installed and configured it for NiFi Registry.

Procedure

1. Install CDP Private Cloud Base.

See the *CDP Private Cloud Base Installation Guide* for more information.

2. Enable Auto-TLS and Kerberos.

Cloudera recommends the following security configuration:

- Enable Auto-TLS. Unsecured NiFi clusters are not supported.
- Enable Kerberos. Kerberos is required if you are using Apache Ranger.
- Use Apache Atlas for dataset level lineage graphs.
- Use Apache Ranger to authorize NiFi and NiFi Registry users.

For details on security recommendations and options, see *CFM Security*.



Note:

Wildcard certificates are not supported.

- If NiFi or NiFi Registry is behind Knox, do not use wildcard certificates for Knox.
- Do not generate wildcard certificates for the NiFi nodes. For example, if two nodes, node1.nifi.apache.org and node2.nifi.apache.org, are assigned the same certificate with a CN or SAN entry of *.nifi.apache.org, this certificate will not be supported.

For more information on certificate requirements, see *TLS certificate requirements and recommendations*.

3. Install the following Runtime services, at minimum:

Dependency	Description
ZooKeeper	NiFi has a required dependency on ZooKeeper, and this service must be installed.
Atlas	Atlas is an optional dependency. It is used for data lineage across the components. If you plan to use Atlas, install it as part of the Base cluster prior the installation of NiFi and NiFi Registry
Ranger	Ranger is an optional dependency. Ranger is used to manage user access policies. If you plan to use Ranger, install it as part of the Base cluster prior the installation of NiFi and NiFi Registry.

4. If you want to install the Ranger service and store Ranger audit logs, you have two options.
 - Install the HDFS service for long term audit log archive and the Solr service for searching and indexing the audit logs from the last 30 days. This is the default.

To do this, select HDFS and Solr as dependencies when installing Ranger.
 - Install only the Solr service, which stores audit logs for 30 days. To do this, select the Solar dependency when installing Ranger, and make the following configurations:
 - Select the Core Configuration service instead of HDFS as a dependency when installing Ranger.
 - Deselect Ranger Plugin DFS Audit Enabled option during Ranger service installation.

What to do next

When you have completed the CDP Private Cloud Base cluster installation, add the CFM parcel and CSD files.

Related Information

[CDP Private Cloud Base Installation Guide](#)

[CFM Security](#)

[Installing the JDK](#)

[Installing and configuring a database for NiFi Registry](#)

[TLS certificate requirements and recommendations](#)

[Installing the CFM parcel from the repository](#)

Installing the CFM parcel from the repository

Install the CFM parcel to make NiFi and NiFi Registry available for installation. To install the CFM parcel to Cloudera Manager, you must update the Parcel URL, and then download, distribute and activate the CFM parcel.

Before you begin

- You have installed a JDK.
- You have installed and configured a database for use with NiFi Registry.
- You have installed a CDP Private Cloud Base cluster.

Procedure

1. Navigate to the Parcels page.
2. From the navigation bar, select **Hosts** **Parcels** and click **Parcel Repository & Network Settings**.
3. On the **Parcel Repository & Network Settings** page, click **+** to add an additional row in the **Remote Parcel Repository URLs** list.
4. Add the URL of the CFM parcel you want to install.

See *Download locations* for a complete list of parcels.

For example:

```
https://archive.cloudera.com/p/cfm2/2.1.5.1000/redhat7/yum/tars/parcel/CFM-2.1.5.1000-56-e17.parcel
```

5. Click **Save & Verify Configuration**.
6. Click **Close**.

The new CFM parcel displays with the set of parcels available for download on the **Parcels** page.

7. From the **Parcels** page, download, distribute, and activate the CFM parcel.

What to do next

When you have finished downloading, distributing and activating the CFM parcel, add the CSD files for use by Cloudera Manager Server. For details, see *Download the CFM Custom Service Descriptor files*.

Related Information

[Installing the JDK](#)

[Installing and configuring a database for NiFi Registry](#)

[Installing Cloudera Manager and a CDP Private Cloud Base cluster](#)

[Download locations](#)

[CFM installation workflow](#)

Downloading Custom Service Descriptor files

A Custom Service Descriptor (CSD) file contains configuration information needed to describe and manage a new service. Download the CSDs for each Cloudera Flow Management (CFM) Service, specify ownership and permissions, and then restart the Cloudera Manager Server.

Before you begin

You have installed the CFM parcel from the repository.

Procedure

1. Download the CSDs for NiFi and NiFi Registry and put them into `/opt/cloudera/csd`.

You can download the files from the CFM repository. See *Download Locations*.



Important:

Ensure that the CSD version is the same as the CFM version. Different CSD and CFM versions can cause problems when configuring and managing NiFi through Cloudera Manager.

2. Change the CSD owner. In `opt/cloudera/csd`, enter:

```
chown cloudera-scm:cloudera-scm ./*  
chmod 644 ./*
```

3. Restart the Cloudera Manager Server.

```
sudo service cloudera-scm-server restart
```

4. Restart the Cloudera Management Service.

On the Cloudera Manager Status tab, select Restart from the Cloudera Management Service drop-down.

What to do next

When you have finished adding the CSD files, you have completed your CDP Private Cloud Base installation. You can now proceed with installing, securing, and configuring CFM.

You can install NiFi and NiFi Registry on the Base cluster, on a Compute cluster, or as a hybrid setup. To decide which option is best for you, see *CFM installation workflow*. For the installation steps, see *Installing NiFi and NiFi Registry on your Base cluster*, *Installing NiFi and NiFi Registry on a Compute cluster*, or *Installing NiFi on a Compute cluster and NiFi Registry on a Base cluster* respectively.

Installing NiFi and NiFi Registry on your Base cluster

Learn the steps to add and configure NiFi and NiFi Registry on a CDP Private Cloud Base Base cluster.

**Note:**

Cloudera recommends that you always install NiFi Registry on your Base cluster.

Installing both NiFi and NiFi Registry on your Base cluster is a simplified setup that can be used for trial deployments or simple production scenarios.

Adding the NiFi service


Provides the steps for adding and configuring the NiFi service to a CDP Private Cloud Base cluster.

Before you begin

You have installed a CDP Private Cloud Base cluster.

Procedure

1. From Home > Status tab, select the drop-down to the right of your cluster, and select Add Service. Install one service at a time.



2. Specify that you want to add the NiFi service and click Continue to display the Add Service wizard.
3. Select NiFi dependencies and click Continue.
 - ZooKeeper is a required dependency.
 - Select Knox if you want a single entry point to securely access the services.
 - Cloudera recommends that you also select Ranger and Atlas dependencies.
4. Assign Roles to your NiFi service.

Select the hosts onto which you want to install your new NiFi roles. Click Continue.



Note: If you selected Knox as a dependency, then install Knox and NiFi on different nodes because Knox, by default, has a port conflict on 8443.

5. Review changes to your configuration. Click Continue. This will start the service installation.

Provide a value for your Initial Admin Identity. You may choose to further customize your NiFi configuration here.

**Note:**

You must add any group name or identity set by these three parameters to Ranger, in order for your installation install to succeed.

- Initial Admin Identity (nifi.initial.admin.identity)
 - NiFi proxy group (nifi.proxy.group)
 - Initial Admin Groups (nifi.initial.admin.groups)
6. Click Continue and Finish to complete the installation.

7. Open a browser and enter the URL to the NiFi UI.

The URL format for the NiFi UI is based on whether or not you selected Knox as a dependency during the installation:

- If you did not select Knox, the URL format is:

```
https://[***HOSTNAME***]:8443/nifi
```

- If you selected Knox, use the Knox URL as a single entry point to securely access all NiFi nodes and switch nodes if one fails. The format is:

```
https://[***KNOX-GATEWAY-HOSTNAME***]:[***KNOX-GATEWAY-PORT***]/gateway/cdp-proxy/nifi-app/nifi/
```

Results

Verify the new service is added properly by checking the health status for the new service. If the Health Status is Good, then the service added properly.

What to do next

When you have finished adding the NiFi service, proceed with adding the NiFi Registry service.

Related Information

[Preparing your CDP Private Cloud Base cluster](#)

[Adding the NiFi Registry service](#)

Adding the NiFi Registry service


Provides the steps for adding and configuring the NiFi Registry service.

Before you begin

You have added the NiFi Service to a cluster.

Procedure

1. From Home > Status tab, select the drop-down to the right of your cluster, and select Add a Service. Install one service at a time.



2. Specify that you want to add the NiFi Registry service and click Continue to display the Add Service wizard.
3. Select NiFi Registry dependencies and click Continue.

Cloudera recommends that you select NiFi and Ranger dependencies.



Note:

NiFi Registry does not require any dependencies. However, if you select No Optional Dependencies, some services are still selected as dependencies. You can deselect unwanted dependencies once you have finished the installation.

4. Assign Roles to your NiFi Registry service.

Select the host onto which you want to install your new NiFi Registry role. Click Continue.

5. Review changes to your configuration. Click Continue to start the service installation.

Provide a value for your Initial Admin Identity. You may choose to further customize your NiFi configuration here.

**Note:**

You must add any group name or identity set by these three parameters to Ranger, in order for your installation install to succeed.

- Initial Admin Identity (nifi.registry.initial.admin.identity)
- NiFi Registry proxy group (nifi.registry.proxy.group)
- Initial Admin Groups (nifi.registry.initial.admin.groups)

6. Click Continue and Finish to complete the installation.**7.** Open a browser and enter the URL to the NiFi Registry UI.

The URL format for the NiFi Registry UI is based on whether or not you selected Knox as a dependency during the installation:

- If you did not select Knox, the URL format is:

```
https://[***HOSTNAME***]:18443/nifi-registry
```

- If you selected Knox, use the Knox URL as a single entry point to securely access all NiFi Registry nodes and switch nodes if one fails. The format is:

```
https://[***KNOX-GATEWAY-HOSTNAME***]:[***KNOX-GATEWAY-PORT***]/gateway/  
cdp-proxy/nifi-registry-app/nifi-registry/
```

Results

Verify the new service is added properly by checking the health status for the new service. If the Health Status is Good, then the service added properly.

What to do next

When you have finished adding the NiFi Registry service, proceed by connecting NiFi to NiFi Registry.

Related Information

[Adding the NiFi service](#)

[Connecting NiFi to NiFi Registry](#)

Connecting NiFi to NiFi Registry

Learn about the steps to configure NiFi to identify the NiFi Registry so that you can support versioned flows in Cloudera Flow Management.

Before you begin

- You have added and configured NiFi and NiFi Registry.
- You have started both NiFi and NiFi Registry.



Procedure

1. On the NiFi UI, select Controller Settings from the Global Menu.
2. Select the Registry Clients tab.

3. Click the Add icon (+) to launch the Add Registry Client dialog.
4. Add the name and URL location of the NiFi Registry service you just created.

Results

NiFi Settings

GENERAL	REPORTING TASK CONTROLLER SERVICES	REPORTING TASKS	REGISTRY CLIENTS
+			
Name ^	Location	Description	
NiFi Registry	https://cfm-node-3.cloudera.com:18433		 

What to do next

Once you have connected NiFi and NiFi Registry, add users or groups to Ranger policies, and if needed, deselect unwanted NiFi Registry dependencies.

Related Information

[Adding the NiFi Registry service](#)

[Adding users or groups to Ranger policies](#)

Adding users or groups to Ranger policies

Manage what users can access or control by adding users or user groups to Ranger policies.

About this task

Determine what the user can command, control, and observe in a NiFi dataflow or in NiFi Registry and accordingly add the user or a group of users to the appropriate predefined Ranger access policies.

Each predefined Ranger access policy confers specific rights to NiFi or NiFi Registry resources.

For more information, see:

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


Before you begin

- You have installed Ranger on your Base CDP Private Cloud Base cluster.
- You have added and connected the NiFi and NiFi Registry services.

Procedure

1. From the Base cluster, select Ranger from the list of services. Click Ranger Admin Web UI and log into Ranger.


The **Ranger Service Manager** page displays.



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.



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

Edit Policy

Policy Details :

Policy Type: **Access** Add Validity Period

Policy ID: **43**

Policy Name: Flow enabled normal

Policy Label:

NIFI Resource Identifier:

Description:

Audit Logging: **YES**

Allow Conditions : hide -

Select Role	Select Group	Select User	Permissions	Delegate Admin	
<input type="text" value="Select Roles"/>	<input type="text" value="nifi"/>	<input type="text" value="admin"/>	Read <input type="text"/>	<input type="checkbox"/>	<input type="button" value="x"/>
+					

Deny All Other Accesses: **False**

- 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

When you have completed the steps for adding users and groups to Ranger policies, review the steps to deselect unwanted NiFi Registry dependencies and determine whether this task applies to you.

Related Information

- [Predefined Ranger Access Policies for Apache NiFi](#)
- [Predefined Ranger Access Policies for Apache NiFi Registry](#)
- [Connecting NiFi to NiFi Registry](#)
- [Deselecting unwanted NiFi Registry dependencies](#)

Deselecting unwanted NiFi Registry dependencies

You might need to remove unwanted dependencies that have possibly been added to NiFi Registry during cluster installation.

About this task

This is an optional task. During your cluster installation, some dependencies may have been added to NiFi Registry. If you do not want these dependencies, follow these steps to remove them.

Before you begin

- You have installed a CDP Private Cloud Base cluster.
- You have added the NiFi Registry service.
- You have added users or user groups to Ranger policies.

Procedure

1. From Cloudera Manager, click the Clusters tab in the left-hand navigation
2. Click NiFi Registry in the list of services to display the NiFi Registry service page.
3. Select the Configuration tab.
4. Deselect any unwanted dependencies.

The screenshot shows the NiFi Registry Configuration page in Cloudera Manager. The page has a top navigation bar with 'Status', 'Instances', 'Configuration', 'Commands', 'Charts Library', 'Audits', 'NiFi Registry Web UI', and 'Quick Links'. A search bar is present at the top left. On the left side, there is a 'Filters' panel with sections for SCOPE, CATEGORY, and STATUS. The main content area displays a list of dependencies for the NiFi Registry service. The 'RANGER Service' dependency is highlighted with a blue box, and the 'SDX' checkbox is checked. Other dependencies include 'NiFi CA Service Service', 'KNOX Service', 'NiFi Service', and 'Enable Kerberos Authentication'.

Service	Dependency	Selected
NiFi CA Service Service	NiFi Registry (Service-Wide)	<input type="checkbox"/> none
RANGER Service	NiFi Registry (Service-Wide)	<input checked="" type="checkbox"/> SDX
KNOX Service	NiFi Registry (Service-Wide)	<input type="checkbox"/> none
NiFi Service	NiFi Registry (Service-Wide)	<input checked="" type="checkbox"/> NiFi
Enable Kerberos Authentication	NiFi Registry (Service-Wide)	<input checked="" type="checkbox"/> NiFi Registry (Service-Wide)

5. Click Save Changes. Restart the NiFi Registry service.

Results

You have completed your cluster and CFM installation.

Related Information

[Adding users or groups to Ranger policies](#)

Installing NiFi and NiFi Registry on a Compute cluster

Learn the steps to add and configure NiFi and NiFi Registry on a CDP Private Cloud Base Compute cluster.



Important: This is not a recommended setup. Cloudera recommends that you always install NiFi Registry on your Base cluster.

Creating a Shared Data Context

You need to create a Shared Data Context so that the data, metadata, and security services can be shared from the Base cluster. This can then be used to create Compute clusters.

Before you begin

- You have reviewed *CFM installation workflow* and have decided to install NiFi on a Compute cluster.
- You have prepared your CDP Private Cloud Base Base cluster.
- Your Base cluster has data, metadata, and security services to share.

Procedure

1. From Cloudera Manager Home, click on your Base cluster name to go the Detail page.
2. In the Data Contexts section, click Create.
3. In the Create Data Context dialog, provide a Data Context Name, specify your Base cluster, and click Create.

Example

Create Data Context ✕

A Data Context, part of [Cloudera SDX \(Shared Data Experience\)](#), allows you to share data, metadata, and security services from a Base Cluster. You can then use it to create **separate Compute Clusters**.

Data Context Name

Base Cluster

Only version 5.15 or higher is supported.

Data Services

  HDFS

Metadata Services

  Hive

Security Services

  Atlas

  Ranger

Cancel

Create

What to do next

When you have finished creating a Shared Data Context, proceed to adding the NiFi and NiFi Registry groups to Ranger in the Base cluster if needed, then to creating a Compute cluster.

Related Information

[CFM installation workflow](#)

[Preparing your CDP Private Cloud Base cluster](#)

[Adding the NiFi and NiFi Registry groups to Ranger in the Base cluster](#)

[Creating a Compute cluster](#)

Adding the NiFi and NiFi Registry groups to Ranger in the Base cluster

Unless the CFM parcel has been activated on the Base cluster, you need to add the NiFi and NiFi Registry groups in Ranger.

About this task

**Note:**

If the CFM parcel has been activated on the Base cluster, the nifi and nifi registry groups in Ranger are automatically created and you may skip this task.

If your Compute cluster uses a Shared Data Context that shares the Ranger service, you must create the nifi and nifi registry groups in Ranger on the Base cluster. You must create these groups before you install the NiFi and NiFi Registry service so that the required NiFi and NiFi Registry Ranger access policies are created.

Before you begin

- You have created a Shared Data Context.

Procedure

1. Go to Ranger.
2. Select Settings > Users/Groups/Roles
The **Users/Groups/Roles** page appears.
3. Select the Groups tab, then click Add New Group.
The **Group Detail** page appears.
4. For Group Name, enter nifi. Click Save.
5. Click Add New Group.
6. For Group Name, enter nifi registry. Click Save.

What to do next

When you have finished creating the Ranger groups, proceed to creating a Compute cluster.

Related Information

[Creating a Shared Data Context](#)

[Creating a Compute cluster](#)

Creating a Compute cluster

Provides the steps to create a CDP Private Cloud Base Compute cluster.

About this task

You should follow the instructions in the *CDP Private Cloud Base Installation Guide* for complete information about installing Cloudera Manager and a CDP base cluster. At minimum, you should ensure that you perform the following steps.

Before you begin

- You have prepared your CDP Private Cloud Base Base cluster.
- You have created a Shared Data Context in your Base cluster.

Procedure

1. Launch the Add Cluster – Installation wizard.

From your Base cluster detail page, select Add Compute Cluster from the Actions drop-down.

2. From the Cluster Basics step, add a cluster name, and select your Shared Data Context from the Data Context drop-down. Click Continue.

Add Cluster - Installation

Cluster Basics

Cluster Name

Compute Cluster Base Cluster

Data Context

Data Context

Data Services HDFS-1

Metadata Services None

Security Services RANGER-1 ATLAS-1

3. From the Specify Hosts step, select the hosts you want to belong to your Compute cluster. Click Continue.

4. From the Select Repository step, select the CFM parcel. If the CFM parcel has not been added yet, select Parcel Repositories & Network Settings and add the CFM Parcel URL

See *Install the CFM parcel from the repository* for more information.

The CDH parcel is added by default. Click Continue to download, distribute, and activate the CFM and CDH parcels.

5. From the Inspect Cluster step, run the Inspect Network Performance and Inspect Hosts inspections. Click Continue to select services to install.

What to do next

When you have finished creating your Compute cluster, proceed by adding the NiFi and NiFi Registry services.

Related Information

[CDP Private Cloud Base Installation Guide](#)

[Preparing your CDP Private Cloud Base cluster](#)

[Creating a Shared Data Context](#)

[Installing the CFM parcel from the repository](#)

[Adding the NiFi and NiFi Registry services to a Compute cluster](#)

Adding the NiFi and NiFi Registry services to a Compute cluster

Provides the steps for adding and configuring the NiFi and NiFi Registry services to a Compute cluster.

Before you begin

- You have created a Shared Data Context and a Compute cluster.
- You have added the nifi and nifiregistry groups to Ranger in the Base cluster.
- You have created a Compute cluster with CDH and CFM parcels installed.

Procedure

1. After your Compute cluster is created, the Add Cluster - Configuration wizard displays.
2. From the Select Services step, select Custom Services. Select NiFi and NiFi Registry. Click Continue.
3. From the Assign Roles step, select the hosts onto which you want to install your new NiFi roles and the host onto which you want to install your NiFi Registry role.

ZooKeeper is a required dependency of NiFi. Select the hosts onto which you want to install ZooKeeper. Click Continue.

4. From the Review Changes step, review changes to your configuration. For both NiFi and NiFi Registry, provide a value for your Initial Admin Identity. You may choose to further customize your NiFi and NiFi Registry configurations here.



Note:

You must add any group name or identity set by these parameters to Ranger, in order for your installation to succeed.

For NiFi:

- Initial Admin Identity (nifi.initial.admin.identity)
- NiFi proxy group (nifi.proxy.group)
- Initial Admin Groups (nifi.initial.admin.groups)

For NiFi Registry:

- Initial Admin Identity (nifi.registry.initial.admin.identity)
- NiFi Registry proxy group (nifi.registry.proxy.group)
- Initial Admin Groups (nifi.registry.initial.admin.groups)

Click Continue.

5. From the Configure Kerberos step, install Kerberos client libraries on all Compute cluster hosts as needed. Click Continue to initiate the Kerberos enablement.
6. When you have completed these steps for both NiFi and NiFi Registry, click Continue. This starts the NiFi and NiFi Registry services installation.
7. Click Continue and Finish to complete the installation.

8. Open a browser and enter the URL to the NiFi and NiFi Registry UI.

The URL format for the UIs is based on whether or not you selected Knox as a dependency during the installation:

- If you did not select Knox, the URL format is:

- NiFi:

```
https://[***HOSTNAME***]:8443/nifi
```

- NiFi Registry:

```
https://[***HOSTNAME***]:18443/nifi-registry
```

- If you selected Knox, use the Knox URL as a single entry point to securely access all NiFi and NiFi Registry nodes and switch nodes if one fails. The format is:

- NiFi:

```
https://[***KNOX-GATEWAY-HOSTNAME***]:[***KNOX-GATEWAY-PORT***]/gateway/cdp-proxy/nifi-app/nifi/
```

- NiFi Registry:

```
https://[***KNOX-GATEWAY-HOSTNAME***]:[***KNOX-GATEWAY-PORT***]/gateway/cdp-proxy/nifi-registry-app/nifi-registry/
```

Results

Verify the new services are added properly by checking the health status for the new services. If the Health Status is Good, then the services added properly.

What to do next

When you have finished adding the NiFi and NiFi Registry services, proceed by connecting NiFi to NiFi Registry.

Related Information

[Creating a Compute cluster](#)

[Connecting NiFi to NiFi Registry](#)

Connecting NiFi to NiFi Registry

Learn about the steps to configure NiFi to identify the NiFi Registry so that you can support versioned flows in Cloudera Flow Management.

Before you begin



- You have added and configured NiFi and NiFi Registry.
- You have started both NiFi and NiFi Registry.

Procedure

1. On the NiFi UI, select Controller Settings from the Global Menu.
2. Select the Registry Clients tab.
3. Click the Add icon (+) to launch the Add Registry Client dialog.
4. Add the name and URL location of the NiFi Registry service you just created.

Results

NiFi Settings

GENERAL	REPORTING TASK CONTROLLER SERVICES	REPORTING TASKS	REGISTRY CLIENTS
+			
Name ^	Location	Description	
NiFi Registry	https://cfm-node-3.cloudera.com:18433		 

What to do next

Once you have connected NiFi and NiFi Registry, add users or groups to Ranger policies, and if needed, deselect unwanted NiFi Registry dependencies.

Adding users or groups to Ranger policies

Manage what users can access or control by adding users or user groups to Ranger policies.

About this task

Determine what the user can command, control, and observe in a NiFi dataflow or in NiFi Registry and accordingly add the user or a group of users to the appropriate predefined Ranger access policies.

Each predefined Ranger access policy confers specific rights to NiFi or NiFi Registry resources.

For more information, see:

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


Before you begin

- You have installed Ranger on your Base CDP Private Cloud Base cluster.
- You have added and connected the NiFi and NiFi Registry services.

Procedure

1. From the Base cluster, select Ranger from the list of services. Click Ranger Admin Web UI and log into Ranger.


The **Ranger Service Manager** page displays.



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.



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

Policy Details :

Policy Type: **Access** [Add Validity Period](#)

Policy ID: **43**

Policy Name: **enabled** **normal**

Policy Label:

NiFi Resource Identifier:

Description:

Audit Logging: **YES**

Allow Conditions : hide -

Select Role	Select Group	Select User	Permissions	Delegate Admin	
<input type="text" value="Select Roles"/>	<input type="text" value="nifi"/>	<input type="text" value="admin"/>	Read <input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
+					

Deny All Other Accesses: **False**

Save

- 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

When you have completed the steps for adding users and groups to Ranger policies, review the steps to deselect unwanted NiFi Registry dependencies and determine whether this task applies to you.

Deselecting unwanted NiFi Registry dependencies

You might need to remove unwanted dependencies that have possibly been added to NiFi Registry during cluster installation.

About this task

This is an optional task. During your cluster installation, some dependencies may have been added to NiFi Registry. If you do not want these dependencies, follow these steps to remove them.

Before you begin

- You have installed a CDP Private Cloud Base cluster.
- You have added the NiFi Registry service.
- You have added users or user groups to Ranger policies.

Procedure

1. From Cloudera Manager, click the Clusters tab in the left-hand navigation
2. Click NiFi Registry in the list of services to display the NiFi Registry service page.
3. Select the Configuration tab.
4. Deselect any unwanted dependencies.

The screenshot shows the NiFi Registry Configuration page in Cloudera Manager. The page has a top navigation bar with 'Status', 'Instances', 'Configuration', 'Commands', 'Charts Library', 'Audits', 'NiFi Registry Web UI', and 'Quick Links'. Below the navigation bar is a search bar and a 'Filters' toggle. The main content area is divided into a left sidebar with 'Filters' and a main table of dependencies. The 'Filters' sidebar shows 'SCOPE' (NiFi Registry (Service-Wide) 21, Gateway 3, NiFi Registry 272) and 'CATEGORY' (Advanced 24, Logs 4, Main 210, Monitoring 16, Performance 1, Resource Management 5, Security 7, Stacks Collection 5). The main table lists dependencies with checkboxes and a 'Show All Descriptions' link. The 'RANGER Service' dependency is highlighted with a blue box, and the 'SDX' checkbox is checked. Other dependencies include 'NiFi CA Service Service' (none), 'KNOX Service' (none), 'NiFi Service' (NiFi), and 'Enable Kerberos Authentication' (NiFi Registry (Service-Wide)).

5. Click Save Changes. Restart the NiFi Registry service.

Results

You have completed your cluster and CFM installation.

Installing NiFi on a Compute cluster and NiFi Registry on a Base cluster

Learn the steps to add and configure NiFi on a CDP Private Cloud Base Compute cluster and NiFi Registry on the Base cluster.



Note: Cloudera recommends that you use this layout for production deployments.

Installing NiFi Registry on your Base cluster


Provides the steps for adding and configuring the NiFi Registry service.

Before you begin

You have prepared your CDP Base cluster.

Procedure

1. From Home > Status tab, select the drop-down to the right of your cluster, and select Add a Service. Install one service at a time.



2. Specify that you want to add the NiFi Registry service and click Continue to display the Add Service wizard.
3. Select NiFi Registry dependencies and click Continue.

Cloudera recommends that you select NiFi and Ranger dependencies.



Note:

NiFi Registry does not require any dependencies. However, if you select No Optional Dependencies, some services are still selected as dependencies. You can deselect unwanted dependencies once you have finished the installation.

4. Assign Roles to your NiFi Registry service.
Select the host onto which you want to install your new NiFi Registry role. Click Continue.
5. Review changes to your configuration. Click Continue to start the service installation.

Provide a value for your Initial Admin Identity. You may choose to further customize your NiFi configuration here.



Note:

You must add any group name or identity set by these three parameters to Ranger, in order for your installation install to succeed.

- Initial Admin Identity (nifi.registry.initial.admin.identity)
- NiFi Registry proxy group (nifi.registry.proxy.group)
- Initial Admin Groups (nifi.registry.initial.admin.groups)

6. Click Continue and Finish to complete the installation.

Results

Verify the new service is added properly by checking the health status for the new service. If the Health Status is Good, then the service added properly.

What to do next

When you have finished adding the NiFi Registry service, proceed by creating a Shared Data Context for your Compute cluster.

Related Information

[Preparing your CDP Private Cloud Base cluster](#)

[Creating a Shared Data Context](#)

Creating a Shared Data Context

You need to create a Shared Data Context so that the data, metadata, and security services can be shared from the Base cluster. This can then be used to create Compute clusters.

Before you begin

- You have reviewed *CFM installation workflow* and have decided to install NiFi on a Compute cluster.
- You have installed NiFi Registry on your Base cluster.
- Your Base cluster has data, metadata, and security services to share.

Procedure

1. From Cloudera Manager Home, click on your Base cluster name to go the Detail page.
2. In the Data Contexts section, click Create.
3. In the Create Data Context dialog, provide a Data Context Name, specify your Base cluster, and click Create.

Example

Create Data Context ✕

A Data Context, part of [Cloudera SDX \(Shared Data Experience\)](#), allows you to share data, metadata, and security services from a Base Cluster. You can then use it to create **separate Compute Clusters**.

Data Context Name

Base Cluster

Only version 5.15 or higher is supported.

Data Services

  HDFS

Metadata Services

  Hive

Security Services

  Atlas

  Ranger

Cancel

Create

What to do next

When you have finished creating a Shared Data Context, proceed to adding NiFi and NiFi Registry groups to Ranger.

Related Information

[CFM installation workflow](#)

[Installing NiFi Registry on your Base cluster](#)

[Adding the NiFi and NiFi Registry groups to Ranger in the Base cluster](#)

Adding the NiFi and NiFi Registry groups to Ranger in the Base cluster

Unless the CFM parcel has been activated on the Base cluster, you need to add the NiFi and NiFi Registry groups in Ranger.

About this task

**Note:**

If the CFM parcel has been activated on the Base cluster, the nifi and nifi registry groups in Ranger are automatically created and you may skip this task.

If your Compute cluster uses a Shared Data Context that shares the Ranger service, you must create the nifi and nifi registry groups in Ranger on the Base cluster. You must create these groups before you install the NiFi and NiFi Registry service so that the required NiFi and NiFi Registry Ranger access policies are created.

Before you begin

- You have created a Shared Data Context.

Procedure

1. Go to Ranger.
2. Select Settings > Users/Groups/Roles
The **Users/Groups/Roles** page appears.
3. Select the Groups tab, then click Add New Group.
The **Group Detail** page appears.
4. For Group Name, enter nifi. Click Save.
5. Click Add New Group.
6. For Group Name, enter nifi registry. Click Save.

What to do next

When you have finished creating the Ranger groups, proceed to creating a Compute cluster.

Creating Compute cluster

Provides the steps to create a CDP Private Cloud Base Compute cluster.

About this task

You should follow the instructions in the *CDP Installation Guide* for complete information about installing Cloudera Manager and a CDP base cluster. At minimum, you should ensure that you perform the following steps.

Before you begin

- You have prepared your CDP Private Cloud Base Base cluster.
- You have created a Shared Data Context in your Base cluster.
- You have added NiFi and NiFi Registry groups to Ranger in your Base cluster.

Procedure

1. Launch the Add Cluster – Installation wizard.
From your Base cluster detail page, select Add Compute Cluster from the Actions drop-down.

- From the Cluster Basics step, add a cluster name, and select your Shared Data Context from the Data Context drop-down. Click Continue.

Add Cluster - Installation

- From the Specify Hosts step, select the hosts you want to belong to your Compute cluster. Click Continue.
- From the Select Repository step, select the CFM parcel. If the CFM parcel has not been added yet, select Parcel Repositories & Network Settings and add the CFM Parcel URL
See *Install the CFM parcel from the repository* for more information.
The CDH parcel is added by default. Click Continue to download, distribute, and activate the CFM and CDH parcels.
- From the Inspect Cluster step, run the Inspect Network Performance and Inspect Hosts inspections. Click Continue to select services to install.

What to do next

When you have finished creating your Compute cluster, proceed by adding NiFi to your Compute cluster.

Related Information

[CDP Private Cloud Base Installation Guide](#)

[Adding the NiFi and NiFi Registry groups to Ranger in the Base cluster](#)

[Adding NiFi to the Compute cluster](#)

Adding NiFi to the Compute cluster

Provides the steps for adding and configuring the NiFi service to a Compute cluster.

Before you begin

- You have created a Shared Data Context and a Compute cluster.

- You have added the NiFi and NiFi Registry groups to Ranger in the Base cluster.
- You have created a Compute cluster with CDH and CFM parcels installed.

Procedure

1. After your Compute cluster is created, the Add Cluster - Configuration wizard displays.
2. From the Select Services step, select Custom Services. Select NiFi and click Continue.
3. From the Assign Roles step, select the hosts onto which you want to install your new NiFi roles.

ZooKeeper is a required dependency of NiFi. Select the hosts onto which you want to install ZooKeeper. Click Continue.

4. From the Review Changes step, review changes to your configuration. Provide a value for your Initial Admin Identity. You may choose to further customize your NiFi configurations here.



Note:

You must add any group name or identity set by these parameters to Ranger, in order for your installation to succeed.

- Initial Admin Identity (nifi.initial.admin.identity)
- NiFi proxy group (nifi.proxy.group)
- Initial Admin Groups (nifi.initial.admin.groups)

Click Continue.

5. From the Configure Kerberos step, install Kerberos client libraries on all Compute cluster hosts as needed. Click Continue to initiate the Kerberos enablement.
6. When you have completed these steps, click Continue. This starts the NiFi service installation.
7. Click Continue and Finish to complete the installation.
8. Open a browser and enter the URL to the NiFi UI.

The URL format for the NiFi UI is based on whether or not you selected Knox as a dependency during the installation:

- If you did not select Knox, the URL format is:

```
https://[***HOSTNAME***]:8443/nifi
```

- If you selected Knox, use the Knox URL as a single entry point to securely access all NiFi nodes and switch nodes if one fails. The format is:

```
https://[***KNOX-GATEWAY-HOSTNAME***]:[***KNOX-GATEWAY-PORT***]/gateway/cdp-proxy/nifi-app/nifi/
```

Results

Verify the new services are added properly by checking the health status for the new services. If the Health Status is Good, then the services added properly.

What to do next

When you have finished adding the NiFi service, proceed by creating NiFi node users for NiFi Registry and Ranger.

Related Information

[Creating Compute cluster](#)

[Setting up NiFi node users in NiFi Registry](#)

Setting up NiFi node users in NiFi Registry

To be able to save versioned NiFi dataflows to NiFi Registry, you need to create NiFi node users for NiFi Registry and Ranger.

About this task

You must perform this task to be able to save versioned NiFi dataflows to NiFi Registry.

Before you begin

- You have installed NiFi Registry on your base cluster.
- You have created a compute cluster and added the NiFi service.

Procedure

1. Perform this step if you want to set up NiFi node users through group identity.



Important: This step requires that you configure NiFi Registry to include the "file-user-group-provider" in the authorizers.xml first.

- a) Add a new user for each NiFi node (node user is derived from the NiFi nodes certificate DNs) in your compute cluster on the NiFi Registry UI.
 - b) Add "nifiregistry" group and then add the node users you have created in the previous step to this group on the NiFi Registry UI.
2. From Ranger, perform one of the following steps:
 - If you completed Step 1, verify that the "nifiregistry" group exists in Ranger. If it does not exist, add it.
 - If you did not perform Step 1, add a new user for each NiFi node (node user is derived from the NiFi nodes certificate DNs) in your compute cluster.
 3. Based on the path followed in Step 2, add either the "nifiregistry" group or all NiFi node users to the following NiFi Registry policies in Ranger:
 - /buckets policy with Read permissions
 - /proxy policy with Read, Write, and Delete permissions

What to do next

Once you have created NiFi node users for NiFi Registry and Ranger, proceed by connecting NiFi and NiFi Registry and complete any additional optional tasks.

Related Information

[Adding NiFi to the Compute cluster](#)

[Connecting NiFi to NiFi Registry](#)

Connecting NiFi to NiFi Registry

Learn about the steps to configure NiFi to identify the NiFi Registry so that you can support versioned flows in Cloudera Flow Management.

Before you begin



- You have added and configured NiFi and NiFi Registry.
- You have started both NiFi and NiFi Registry.

Procedure

1. On the NiFi UI, select Controller Settings from the Global Menu.
2. Select the Registry Clients tab.
3. Click the Add icon (+) to launch the Add Registry Client dialog.
4. Add the name and URL location of the NiFi Registry service you just created.

Results

NiFi Settings

GENERAL	REPORTING TASK CONTROLLER SERVICES	REPORTING TASKS	REGISTRY CLIENTS
+			
Name ^	Location	Description	
NiFi Registry	https://cfm-node-3.cloudera.com:18433		 

What to do next

Once you have connected NiFi and NiFi Registry, add users or groups to Ranger policies, and if needed, deselect unwanted NiFi Registry dependencies.

Adding users or groups to Ranger policies

Manage what users can access or control by adding users or user groups to Ranger policies.

About this task

Determine what the user can command, control, and observe in a NiFi dataflow or in NiFi Registry and accordingly add the user or a group of users to the appropriate predefined Ranger access policies.

Each predefined Ranger access policy confers specific rights to NiFi or NiFi Registry resources.

For more information, see:

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


Before you begin

- You have installed Ranger on your Base CDP Private Cloud Base cluster.
- You have added and connected the NiFi and NiFi Registry services.

Procedure

1. From the Base cluster, select Ranger from the list of services. Click Ranger Admin Web UI and log into Ranger.

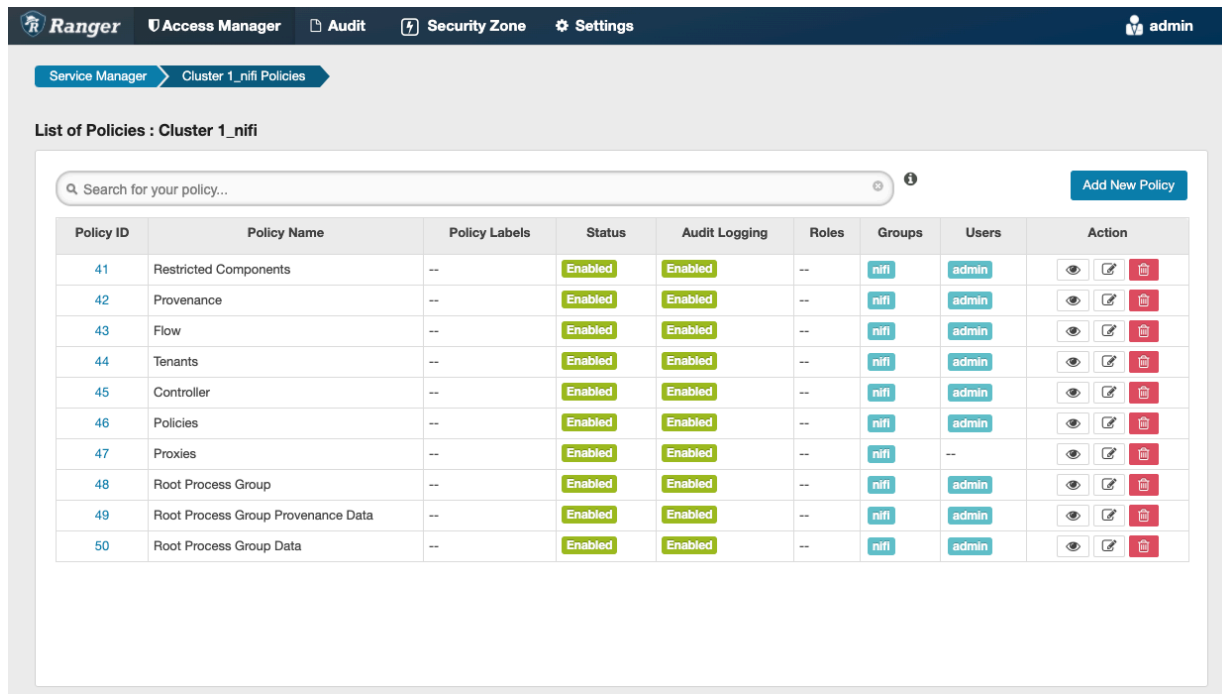
The **Ranger Service Manager** page displays.



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.



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

Policy Details :

Policy Type: **Access** Add Validity Period

Policy ID: **43**

Policy Name: Flow enabled normal

Policy Label:

NIFI Resource Identifier:

Description:

Audit Logging: **YES**

Allow Conditions : hide -

Select Role	Select Group	Select User	Permissions	Delegate Admin	
<input type="text" value="Select Roles"/>	<input type="text" value="nifi"/>	<input type="text" value="admin"/>	Read <input type="text"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
+					

Deny All Other Accesses: **False**

Save

- 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

When you have completed the steps for adding users and groups to Ranger policies, review the steps to deselect unwanted NiFi Registry dependencies and determine whether this task applies to you.

Deselecting unwanted NiFi Registry dependencies

You might need to remove unwanted dependencies that have possibly been added to NiFi Registry during cluster installation.

About this task

This is an optional task. During your cluster installation, some dependencies may have been added to NiFi Registry. If you do not want these dependencies, follow these steps to remove them.

Before you begin

- You have installed a CDP Private Cloud Base cluster.
- You have added the NiFi Registry service.
- You have added users or user groups to Ranger policies.

Procedure

1. From Cloudera Manager, click the Clusters tab in the left-hand navigation
2. Click NiFi Registry in the list of services to display the NiFi Registry service page.
3. Select the Configuration tab.
4. Deselect any unwanted dependencies.

The screenshot shows the NiFi Registry Configuration page in Cloudera Manager. The page has a top navigation bar with 'Status', 'Instances', 'Configuration', 'Commands', 'Charts Library', 'Audits', 'NiFi Registry Web UI', and 'Quick Links'. A search bar is present. On the left, there are filters for SCOPE, CATEGORY, and STATUS. The main content area shows a list of services and their dependencies. The 'RANGER Service' dependency for 'NiFi Registry (Service-Wide)' is highlighted with a blue box, and the 'SDX' checkbox is checked. Other dependencies include 'NiFi CA Service Service', 'KNOX Service', 'NiFi Service', and 'Enable Kerberos Authentication'.

Service	Dependency	Selected
NiFi CA Service Service	NiFi Registry (Service-Wide)	<input type="checkbox"/>
RANGER Service	NiFi Registry (Service-Wide)	<input checked="" type="checkbox"/>
KNOX Service	NiFi Registry (Service-Wide)	<input type="checkbox"/>
NiFi Service	NiFi Registry (Service-Wide)	<input checked="" type="checkbox"/>
Enable Kerberos Authentication	NiFi Registry (Service-Wide)	<input checked="" type="checkbox"/>

5. Click Save Changes. Restart the NiFi Registry service.

Results

You have completed your cluster and CFM installation.