

Using custom processors

Date published: 2021-04-06

Date modified: 2024-06-03

CLOUDERA

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

Best practices for building custom components.....	4
Preparing cloud storage to deploy custom processors.....	5

Best practices for building custom components

Learn about general guidelines concerning the creation of custom NiFi archives (NARs).

The goal is to build your code once against a baseline version of Apache NiFi and it can be deployed to any flow you need using any version of Cloudera DataFlow powered by any NiFi deployment equal to or greater than the version it was built against, bar major version changes.

Apache NiFi extensions are packaged in NARs. A NAR allows several components and their dependencies to be packaged together into a single package. NiFi provides Maven archetypes for creating custom processor and controller service bundle project structures. For detailed information, see the [Maven Projects for Extensions](#) Apache NiFi wiki page.

- If directly specifying the nifi-nar-maven-plugin, ensure you use the most recent version when compiling your custom code.
- If inheriting the nifi-nar-maven-plugin, Cloudera recommends that the parent version of nifi-nar-bundles has the same major and minor version as the selected NiFi runtime version. For example, If the CFM Version is 1.18.0.2.3.7.1 (NiFi Major Version: 1, NiFi Minor Version: 18) the recommended compilation version is 1.18.0 (The first two version numbers must be equal to, or less than the CFM version).

The screenshot shows the NiFi Configuration page. The 'NiFi Runtime Version' section displays the 'CURRENT VERSION' as 'Latest Version (1.18.0.2.3.8.0-11)'. A modal dialog titled 'Change NiFi Runtime Version' is open, showing a dropdown menu with 'Version 1.18.0.2.3.7.1-1' selected. The dialog also includes a 'Cancel' button and an 'Apply' button.

```
<parent>
  <groupId>org.apache.nifi</groupId>
  <artifactId>nifi-nar-bundles</artifactId>
  <version>1.18.0.2.3.7.1-1</version>
</parent>
```

- Ensure your NAR pom only depends on another NAR pom for a controller service API NAR. Generally, do not extend from implementation NARs like nifi-standard-nar.
- Ensure your components jar pom marks API dependencies as provided in order to obtain them at runtime through the appropriate NAR dependency.

Preparing cloud storage to deploy custom processors

To use a custom Apache NiFi processor or controller service in one of your Cloudera DataFlow flow deployments, add the NiFi Archive (NAR) containing the custom processor or controller service to a cloud storage location for later use during a flow deployment.

Procedure

1. Create your cloud storage location.
2. Upload your NAR files to the cloud storage location.
3. Configure access to your cloud provider storage in one of two ways:
 - You have configured access to S3 buckets using ID Broker mapping.

If your environment is not RAZ-enabled, you can configure access to S3 buckets using ID Broker mapping.

a. Access IDBroker mappings.

1. To access IDBroker mappings in your environment, click **Actions Manage Access**.
 2. Choose the IDBroker Mappings tab where you can provide mappings for users or groups and click **Edit**.
- b. Add your Cloudera Workload User and the corresponding AWS role that provides write access to your folder in your S3 bucket to the **Current Mappings** section by clicking the blue + sign.



Note: You can get the AWS IAM role ARN from the Roles Summary page in AWS and can copy it into the IDBroker role field. The selected AWS IAM role must have a trust policy allowing IDBroker to assume this role.

c. Click **Save and Sync**.

- You have configured access to S3 buckets with a RAZ enabled environment.

It is a best practice to enable RAZ to control access to your object store buckets. This allows you to use your Cloudera credentials to access S3 buckets, increases auditability, and makes object store data ingest workflows portable across cloud providers.

- a. Ensure that Fine-grained access control is enabled for your Cloudera DataFlow environment.
- b. From the Ranger UI, navigate to the S3 repository.
- c. Create a policy to govern access to the S3 bucket and path used in your ingest workflow.



Tip:

The Path field must begin with a forward slash (/).

- d. Add the machine user that you have created for your ingest workflow to the policy you just created.

For more information, see *Creating Ranger policy to use in RAZ-enabled AWS environment*.

4. Note the workload user name and password, and cloud storage location to use in the Deployment Wizard.



Tip: If you want to provide a machine user as CDP Workload Username during flow deployment, make sure to note the full workload user name including the `srv_` prefix.

What to do next

Once you have added the NAR files to a cloud storage location, you are ready to launch the Deployment Wizard and deploy a flow.