

Configuring Flow Management clusters to hot load custom NARs

Date published: 2019-12-16

Date modified: 2021-10-25

The Cloudera logo is displayed in a bold, orange, sans-serif font. The word "CLOUDERA" is written in all caps, with the letter 'E' in "CLouDERA" featuring a unique design where the top bar is a horizontal line that extends to the right and then turns down to form the top of the letter 'R'.

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

Configuring Flow Management clusters to hot load custom NARs.....	4
--	----------

Configuring Flow Management clusters to hot load custom NARs

Hot loading custom NARs means that you can store the NARs for your NiFi custom processors in an external folder, and load them from that directory for use on the NiFi canvas. This makes it easy to edit custom processors and add new ones without having to restart your Flow Management cluster.

Before you begin

- Collect your Kerberos information.
 1. In Management Console, click Environments.
 2. From the Actions drop-down, click Manage Access.
 3. Create a Machine User with the EnvironmentUser Resource Role and set its Workload password.



Note:

It is best practice to have a dedicated Machine User in this role. This Machine User is in charge of accessing the object store resources for NiFi, having a dedicated user ensures that there is least privilege access and separation of duties.

4. Note down the user you create there and password, for your safety valves.
- You have created an ID Broker mapping.
 1. Access IDBroker mappings.
 - a. To access IDBroker mappings in your environment, click Actions > Manage Access.
 - b. Choose the IDBroker Mappings tab where you can provide mappings for users or groups and click Edit.
 2. Add your CDP Machine 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.

3. Click Save and Sync.

About this task

- You can configure NiFi to hot load custom NARs in all supported cloud providers.
- You can only access resources in the object store of the cloud provider in which your environment is running. But you can have a development and a production environment, for example, and configure both environments to access the same location to hot load custom NARs
- It is possible to define multiple locations if desired by repeating the configuration properties with a different “[*custom-name*]” for each location. But all locations should be in the object store of the underlying cloud provider of where the NiFi cluster is located.
- NiFi polls for new NARs in the configured location(s) every 5 minutes.

Procedure

1. In Management Console, select Data Hub Clusters, then go to the Flow Management cluster, then NiFi Service.
2. Click Configuration.
3. Search for Safety Valve and then Find: NiFi Node Advanced Configuration Snippet (Safety Valve) for staging/nifi.properties.xml
4. Create the following new properties:

```
Name: nifi.nar.library.provider.[*custom-name*].implementation
```

```
Example: nifi.nar.library.provider.narstorage1.implementation
Value: org.apache.nifi.nar.hadoop.HDFSNarProvider

Path to the configuration files
Name: nifi.nar.library.provider.[**custom-name**].resources
Example: nifi.nar.library.provider.narstorage1.resources
Value, by default: /etc/hadoop/conf.cloudera.core_settings/core-site.xml
Storage location for the NARs (the bucket name in S3 for example)
Name: nifi.nar.library.provider.[**custom-name**].storage.location
Example: nifi.nar.library.provider.narstorage1.storage.location
Value: s3a://my-custom-nars-bucket

Directory library with path to the NARs relative to the storage location
Name: nifi.nar.library.provider.[**custom-name**].source.directory
Example: nifi.nar.library.provider.narstorage1.source.directory
Value: /myNARs

Workload username and password of the machine user to use to access the re
sources
Name: nifi.nar.library.provider.[**custom-name**].kerberos.principal
Example: nifi.nar.library.provider.narstorage1.kerberos.principal
Value: srv_my_s3_nar_user

Name: nifi.nar.library.provider.[**custom-name**].kerberos.password
Example: nifi.nar.library.provider.narstorage1.kerberos.password
Value: myPassword123
```

**Note:**

Cloudera recommends that the workload user name and password that you configure above is strictly dedicated to accessing custom NARs, and that the only permissions granted to this user are to access the custom NAR object store location.

Alternately, you can also use the keytab file. However, this file must be distributed on all NiFi nodes. This becomes problematic when scaling up a NiFi cluster in DataHub. In addition, the file permissions must be properly changed to ensure the file is readable by the `nifi` user. However, if desired, the below configuration can be used:

```
Name: nifi.nar.library.provider.[**custom-name**].kerberos.keytab
Example: nifi.nar.library.provider.narstorage1.kerberos.keytab
Value: /path/to/my/file.keytab
```

5. Restart the NiFi service in the Cloudera Manager.

Results

Refresh the NiFi UI and then drag and drop a new processor to NiFi canvas and search for your custom processor.