

Using Parameter Context inheritance

Date published: 2019-06-26

Date modified: 2024-10-02



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

What is Parameter Context inheritance?	4
Example for configuring Parameter Context inheritance	4
Creating the basic Parameter Contexts	4
Setting up Parameter Context inheritance	6
Parameter overriding	8

What is Parameter Context inheritance?

With Inherited Parameter Contexts, you can make granular updates to Parameter Contexts and have the changes propagate to all Contexts that inherit from them.

Parameter Context is a named set of parameters with applied access policies in Apache NiFi. Using Parameter Contexts is a powerful way to make flows more portable. Paired with Process Groups, Parameter Contexts allow different values to be supplied to the same basic flow in multiple NiFi instances, or even within the same overall NiFi flow in a single instance. However, as the number of parameters grows, Parameter Contexts can become more difficult to maintain and can result in a monolithic Parameter Context. If you collect all parameters in a root Parameter Context, you are parameterizing your Process Groups with more parameters than they need, as some of the parameters included in the Parameter Context are applicable only to certain Process Groups.

The goal is to avoid parameter duplication and yet still compose groups of parameters as needed. Starting with Apache NiFi 1.15.0, Parameter Contexts support inheritance, so you can add inherited Parameter Contexts to an existing Parameter Context. This feature brings more flexibility and maintainability to NiFi's already powerful Parameter Context framework. Using inheritance makes it easier to manage common sets of parameters used across many flows and adjust specific parameters for a given instance as needed. Creating Parameter Contexts through inheritance provides the opportunity to use a structure of smaller, logically grouped Parameter Contexts and design your flows with these smaller, composable Contexts. Once you set up the inheritance hierarchy, you can make granular updates and have the changes propagate to all other Contexts that inherit from it.

For more information about Process Groups, see *Configuring a Process Group in the Apache NiFi User Guide*.

For more information about Parameter Contexts, see *Parameter Contexts in the Apache NiFi User Guide*.

Related Information

[Configuring a Process Group, Apache NiFi User Guide](#)

[Parameter Contexts, Apache NiFi User Guide](#)

Example for configuring Parameter Context inheritance

In this example you can see how Parameter Context inheritance is configured for two existing Process Groups (Source ABC Kafka to S3 and Source DEF Kafka to GCS) to avoid repeating the Kafka brokers or the common site-specific parameters, and yet still composing groups of parameters as needed. The small, composable Parameter Contexts are logically grouped through inheritance.

Creating the basic Parameter Contexts

Follow these steps to set up your individual Parameter Contexts and add the relevant parameters to them.

Procedure



















1. Create a Parameter Context for the common Kafka settings:
 - a) Select Parameter Contexts from the top-right Global menu in the NiFi UI.
The NiFi Parameter Contexts dialog opens.
 - b) Click the (+) button in the top-right corner of the NiFi Parameter Contexts dialog to add a new Parameter Context.
 - c) Add a name for your Parameter Context on the SETTINGS tab.
In this example the name is Kafka.
 - d) Switch to the PARAMETERS tab.
 - e) Click the (+) button to add the parameters you need for configuring your data flow.
 - f) Provide a name and a value in the Add Parameter dialog. You can also add a description of the parameter, if you want to.
In this example, the following parameters are added:
 - Name: Kafka Brokers / Value: localhost:9092
 - Name: Kafka Group ID / Value: MyGroup
 - g) Click APPLY to save the Parameter Context.
2. Create the remaining Parameter Contexts by repeating sub-steps a to g of Step 1. Use the following names and values:

Parameter Context	Parameter Name	Parameter Value
GCP Cloud Storage	GCP Project ID	my-project
	GCP Bucket	my-bucket
S3	S3 Bucket	my-bucket
	S3 Region	us-west-2
Site Properties	Site Identifier	1234
	Site Data Manager	MyDM
Source ABC	Kafka Topic	source-abc-topic
Source DEF	Kafka Topic	source-def-topic

Results

You are ready with the individual Parameter Contexts that you will use to create your inheritance hierarchy. Your Parameter Contexts should look like this:

NiFi Parameter Contexts

Name ▲				Description		
	GCP Cloud Storage			Configuration for GCP Cloud Storage.		
	Kafka			Contains common Kafka parameters.		
	S3			Configuration for S3.		
	Site Properties			Common properties for this site.		
	Source ABC			Source ABC's configuration.		
	Source DEF			Source DEF's configuration.		

Setting up Parameter Context inheritance

Follow these steps to configure inherited Parameter Contexts.

About this task

Now that you have the individual Parameter Contexts set up, you can begin grouping them together, and create higher-level Parameter Contexts to build out your inheritance hierarchy and use them with the two Process Groups (Source ABC Kafka to S3 and Source DEF Kafka to GCS). The goal is to have one Parameter Context for each Process Group, composed of smaller units.

In this example, site properties and Kafka parameters are needed for both Process Groups. Additionally, each Process Group will have a source-related and a destination-related Parameter Context.

Before you begin

You have created two Process Groups in Apache NiFi:

- Source ABC Kafka to S3 that needs parameters for your Kafka brokers, Source ABC's topic, and for the S3 bucket
- Source DEF Kafka to GCS that needs parameters for the same Kafka brokers, Source DEF's topic, and for the GCP bucket
- In both Process Groups, you would like to use some common parameters tagging flowfiles with site-specific information

Procedure

1. Design the inheritance logic.

In this example, the following structure is used:

- Kafka and Site - inherits from Kafka and Site Properties
- ABC Kafka and Site - inherits from Kafka and Site and Source ABC
- DEF Kafka and Site - inherits from Kafka and Site and Source DEF
- ABC Kafka to S3 - inherits from ABC Kafka and Site and S3
- DEF Kafka to GCS - inherits from DEF Kafka and Site and GCP Cloud Storage

Based on the above structure, the hierarchy for the two top-level Parameter Contexts looks like this:

ABC Kafka to S3 Parameter Context

- ABC Kafka and Site Parameter Context
 - Kafka and Site Parameter Context
 - Kafka Parameter Context
 - Kafka Brokers parameter with value localhost:9092
 - Kafka Group ID parameter with value MyGroup
 - Site Properties Parameter Context
 - Site Identifier parameter with value 1234
 - Site Data Manager parameter with value MyDM
 - Source ABC Parameter Context
 - Kafka Topic parameter with value source-abc-topic
- S3 Parameter Context
 - S3 Bucket parameter with value my-bucket
 - S3 Region parameter with value us-west-2

DEF Kafka to GCS Parameter Context

- DEF Kafka and Site Parameter Context
 - Kafka and Site Parameter Context
 - Kafka Parameter Context
 - Kafka Brokers parameter with value localhost:9092
 - Kafka Group ID parameter with value MyGroup
 - Site Properties Parameter Context
 - Site Identifier parameter with value 1234
 - Site Data Manager parameter with value MyDM
 - Source DEF Parameter Context
 - Kafka Topic parameter with value source-def-topic
- GCP Cloud Storage Parameter Context
 - GCP Project ID parameter with value my-project
 - GCP Bucket parameter with value my-bucket

2. Configure Parameter Context inheritance in NiFi for ABC Kafka to S3.

- Create a new Parameter Context named Kafka and Site and go to the INHERITANCE tab.

You can see all of the individual Parameter Contexts.

- Drag Kafka and Site Properties to the right under Selected Parameter Context.
- Click APPLY, and then edit the Kafka and Site Parameter Context to see the new parameters.

All four parameters from the two Parameter Contexts are now available in the view, and each of them has an (arrow) icon. You can click one of these to navigate to the Parameter Context in which the parameter is defined and edit its value.



Note:

You can drag/drop from the Available Parameter Contexts to the required inherited order in the Selected Parameter Context column.

When you remove a Parameter Context from the Selected Parameter Context list, it returns to the Available Parameter Contexts that is sorted alphabetically.

- Repeat steps 1-3 to create ABC Kafka and Site Parameter Context inheriting all parameters from the Kafka and Site and the Source ABC Parameter Contexts.
- Finally, repeat steps 1-3 to create ABC Kafka to S3 Parameter Context inheriting all parameters from the ABC Kafka and Site and the S3 Parameter Contexts.

This is how ABC Kafka to S3 looks like after the full setup:

Update Parameter Context

SETTINGS

PARAMETERS

INHERITANCE

+

Name ▲	Value	
Kafka Brokers	localhost:9092	→
Kafka Group ID	MyGroup	→
Kafka Topic	source-abc-topic	→
S3 Bucket	my-bucket	→
S3 Region	us-west-2	→
Site Data Manager	MyDM	→
Site Identifier	1234	→

- Configure Parameter Context inheritance for DEF Kafka to GCS. Follow the steps you completed when creating the ABC Kafka to S3 Parameter Context.

Parameter overriding

When a Parameter Context inherits parameters from other Contexts, you can still manually override parameters with a different value, if needed.



In this example, another Process Group is added that uses the same parameters as ABC Kafka to S3, but with a Kafka Group ID of MyOtherGroup.

To configure it, you can override the original parameter value by creating a new Parameter Context that inherits from ABC Kafka to S3 and adding a Kafka Group ID parameter with the MyOtherGroup value directly to that new Context:

Update Parameter Context

SETTINGS	PARAMETERS	INHERITANCE
----------	------------	-------------

+

Name ▲	Value	
Kafka Group ID	MyOtherGroup	 
Kafka Brokers	localhost:9092	→
Kafka Topic	source-abc-topic	→
S3 Bucket	my-bucket	→
S3 Region	us-west-2	→
Site Data Manager	MyDM	→
Site Identifier	1234	→

In this case the Kafka Group ID with its value `MyOtherGroup` appears at the top of the parameter list, and not where it should be in alphabetical order. Direct parameters always appear at the top of the list so you can see them grouped together.

You can edit the manually set parameter directly from this Parameter Context. So in this case, you are editing the parameter in the new Context and not the inherited parameter.

Parameter overriding order

- Direct parameters always take precedence.
- Parameters inside directly inherited Parameter Contexts take precedence, from top to bottom in the Selected Parameter Context column of the INHERITANCE tab.
- This recursively repeats in a depth-first manner for as many layers of inherited Parameter Contexts exist.