Getting Started with Apache NiFi Registry

Date published:
Date modified:

https://docs.cloudera.com/

Legal Notice

Contents

Who is This Guide For?

This guide is written for users who have basic experience with NiFi but have little familiarity with the NiFi Registry. This guide is not intended to be an exhaustive instruction manual or a reference guide. The NiFi Registry User Guide and NiFi User Guide provide a great deal more information about using the Registry and integrating it with NiFi. This guide, in comparison, is intended to provide users with just the information needed in order to understand how to configure NiFi Registry, connect with NiFi and start using versioned NiFi dataflows.

Terminology Used in This Guide

In order to talk about NiFi Registry, there are a few key terms that readers should be familiar with:

Flow: A process group level NiFi dataflow that has been placed under version control and saved to the Registry.

Bucket: A container that stores and organizes flows.

Downloading and Installing NiFi Registry

You can download NiFi and NiFi Registry from the CFM repository. Review *Download from the CFM Repository* to identify the download link suitable for your operating system and operational objectives.

Starting NiFi Registry

Once NiFi Registry has been downloaded and installed as described above, it can be started by using the mechanism appropriate for your operating system.

For Linux/Unix/Mac OS X users

Use a Terminal window to navigate to the directory where NiFi Registry was installed. To run NiFi Registry in the foreground, run bin/nifi-registry.sh run. This will leave the application running until the user presses Ctrl-C. At that time, it will initiate shutdown of the application.

To run NiFi Registry in the background, instead run bin/nifi-registry.sh start. This will initiate the application to begin running. To check the status and see if NiFi Registry is currently running, execute the command bin/nifi-registry.sh status. NiFi Registry can be shutdown by executing the command bin/nifi-registry.sh stop.

Installing as a Service

To install the application as a service, navigate to the installation directory in a Terminal window and execute the command bin/nifi-registry.sh install to install the service with the default name nifi-registry. To specify a custom name for the service, execute the command with an optional second argument that is the name of the service. For example, to install NiFi Registry as a service with the name flow-registry, use the command bin/nifi-registry.sh install flow-registry.

Once installed, the service can be started and stopped using the appropriate commands, such as sudo service nifi-registry start and sudo service nifi-registry stop. Additionally, the running status can be checked via sudo service nifi-registry status.

I Started NiFi Registry. Now What?

Now that NiFi Registry has been started, we can bring up the User Interface (UI). To get started, open a web browser and navigate to http://localhost:18080/nifi-registry. The port can be changed by editing the nifi-registry.properties file in the NiFi Registry conf directory, but the default port is 18080.

This will bring up the Registry UI, which at this point is empty as there are no flow resources available to share yet:

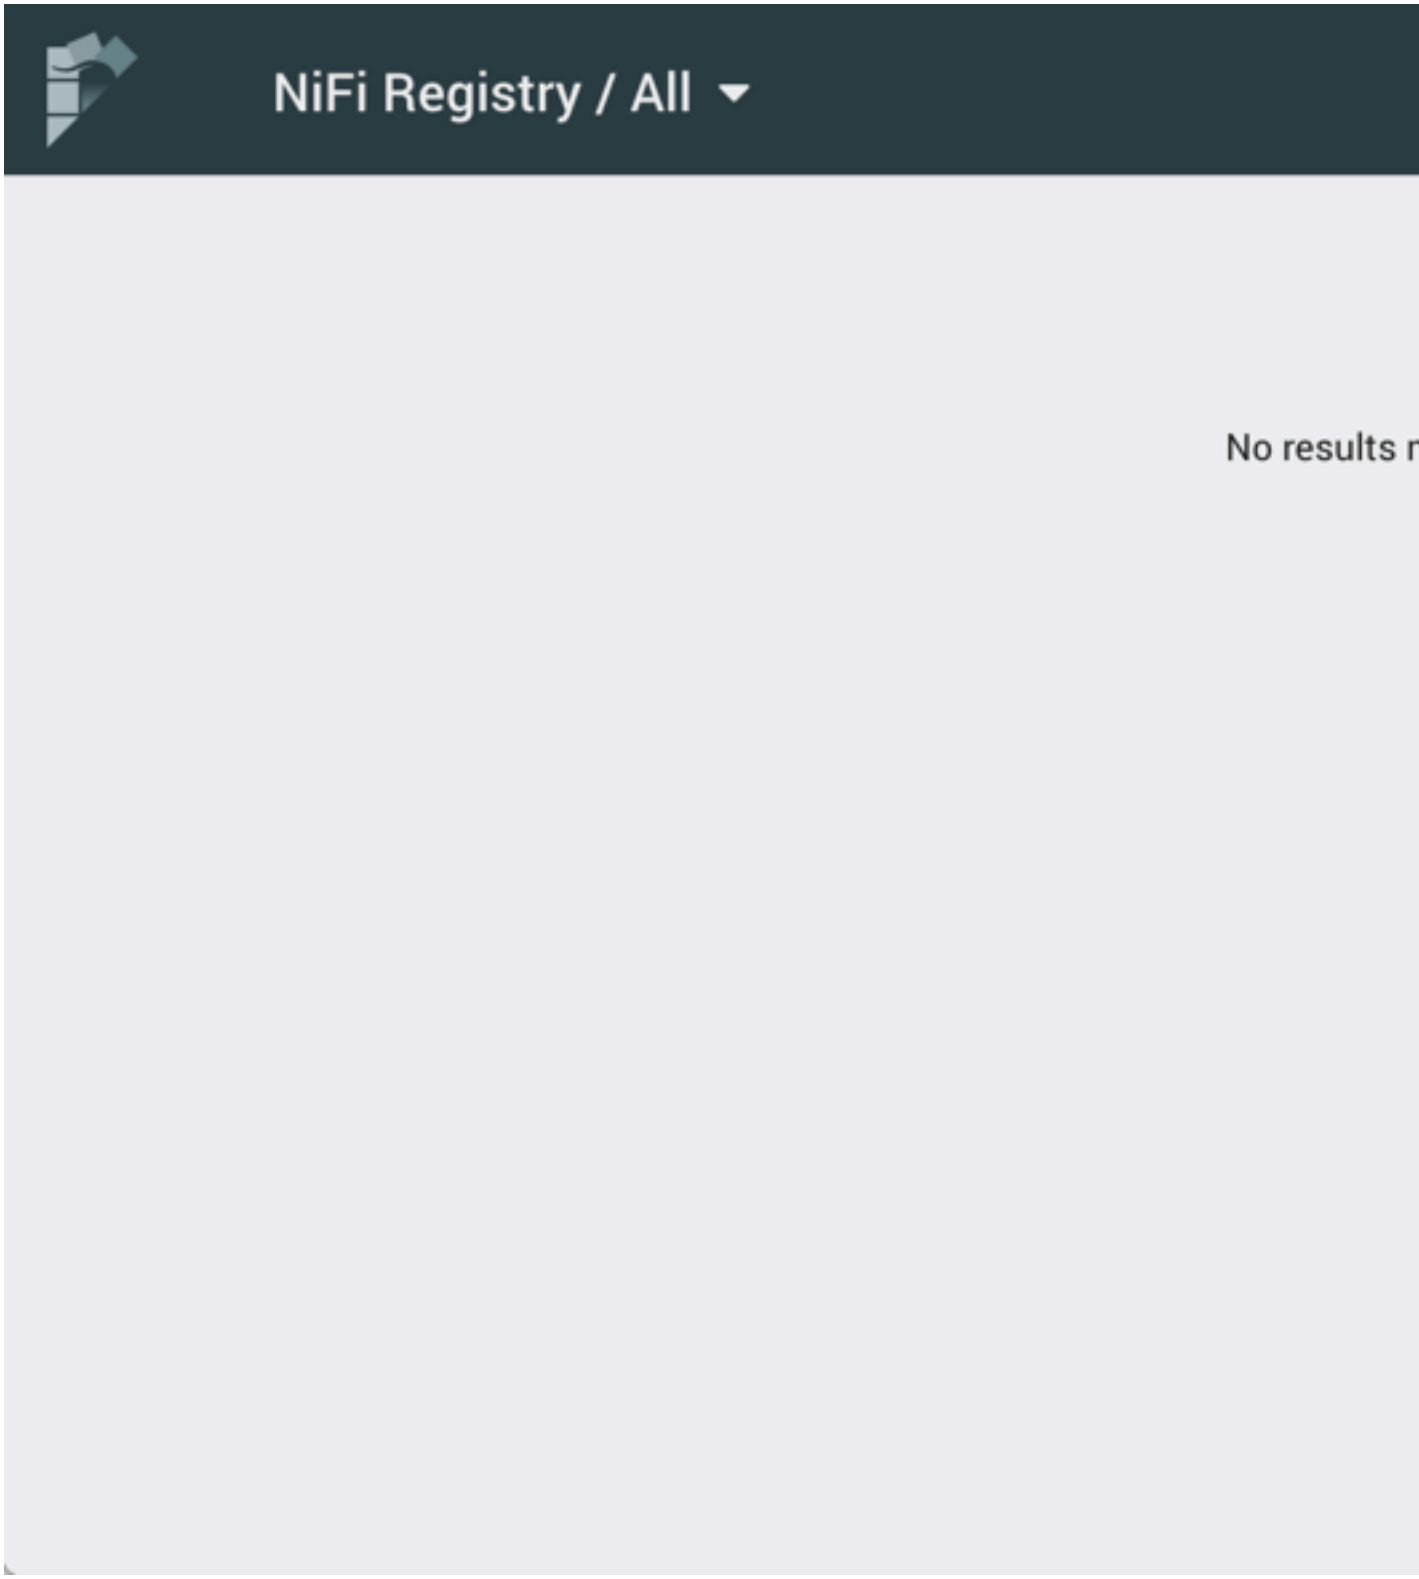

Create a Bucket

A bucket is needed in our registry to store and organize NiFi dataflows. To create one, select the Settings icon ()in the top right corner of the screen. In the Buckets window, select the "New Bucket" button.

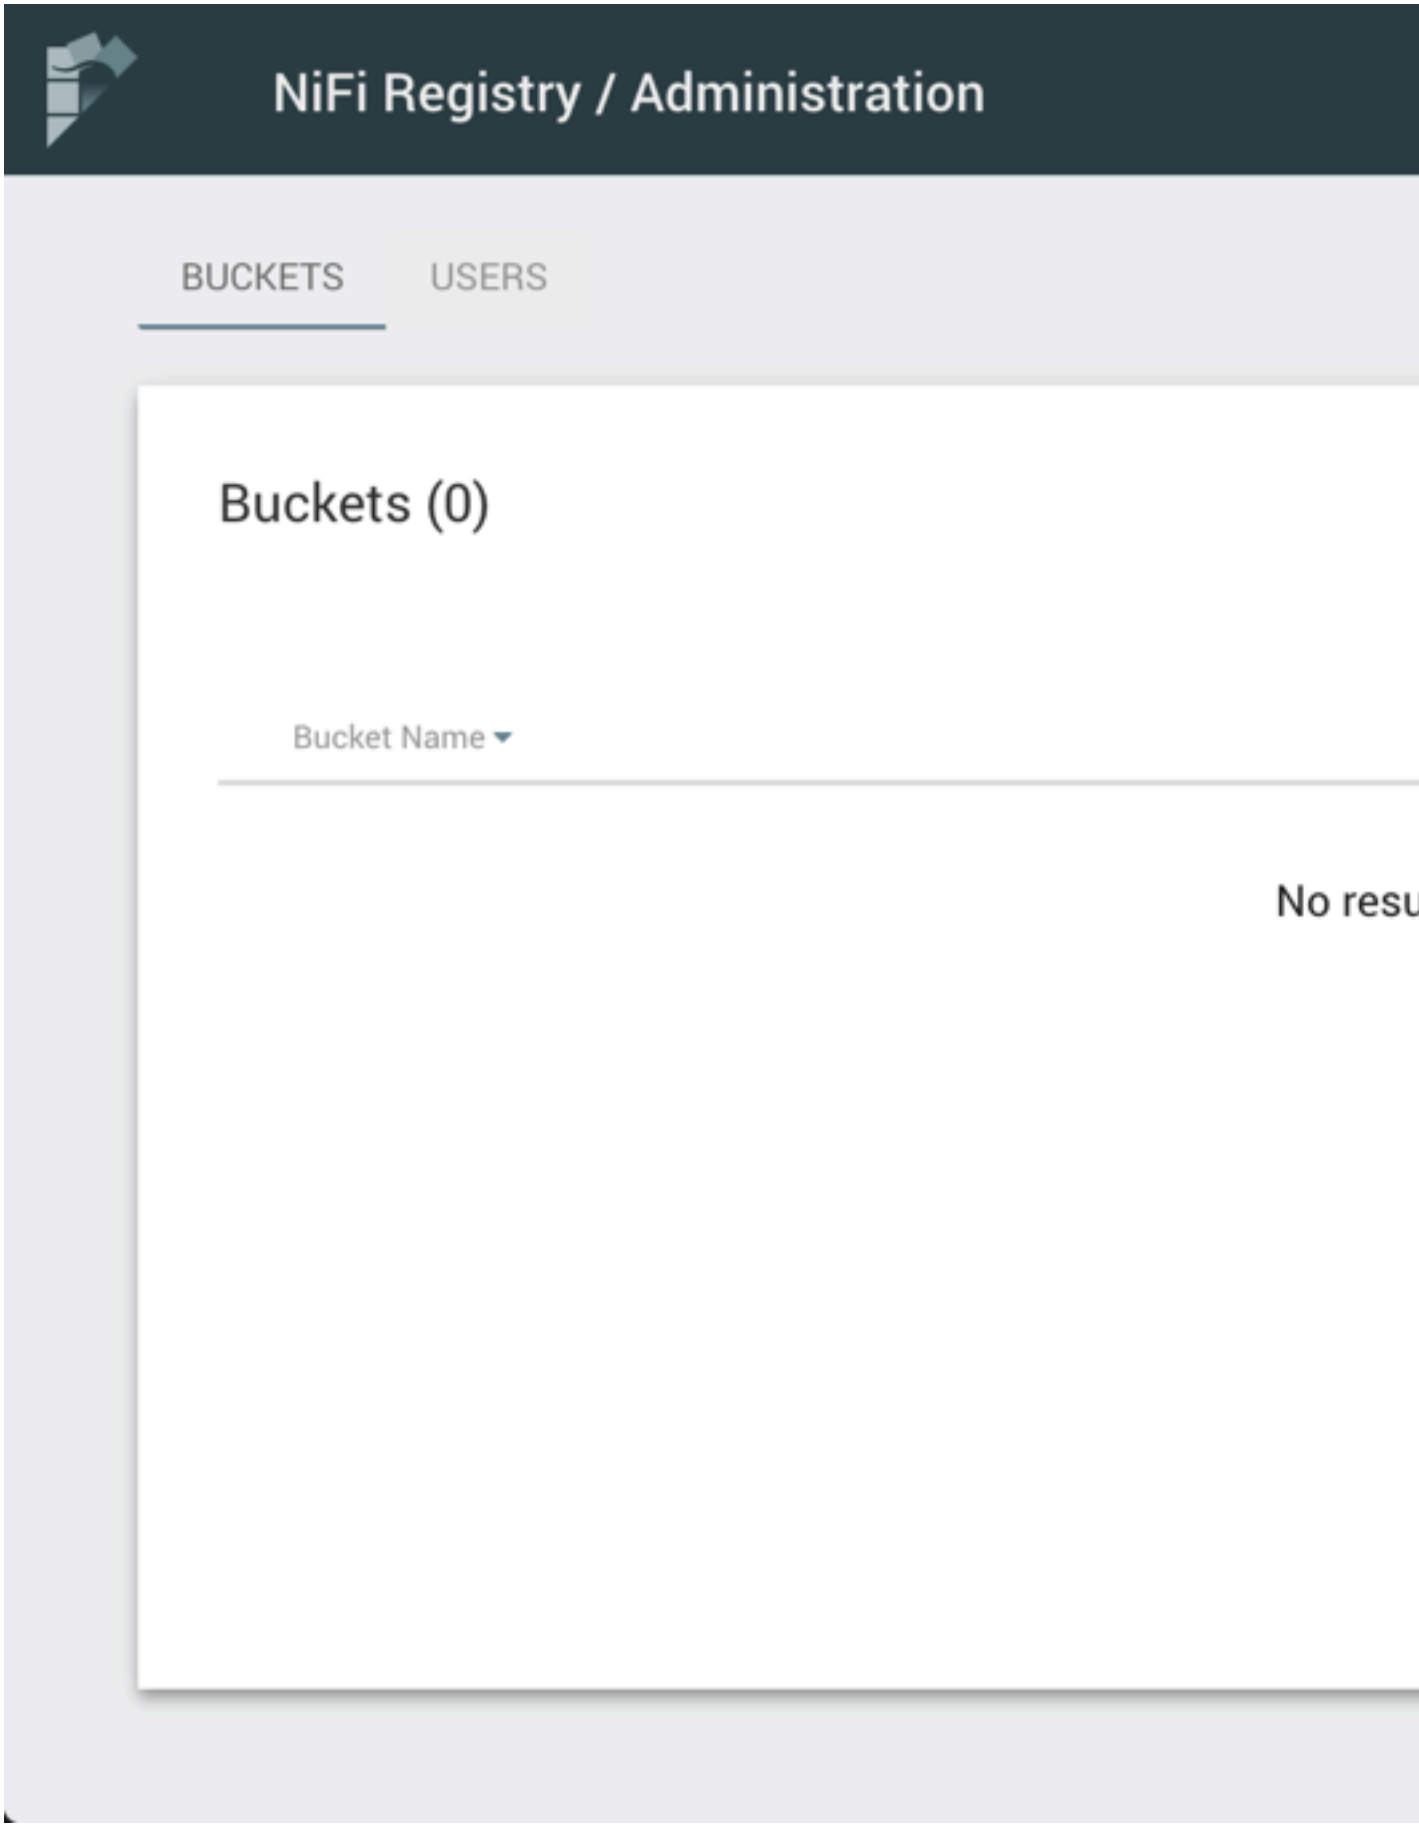
NiFi Registry / Administration
BUCKETS
USERS
Buckets (0)
Bucket Name
No resu

Enter the bucket name "Test" and select the "Create" button.

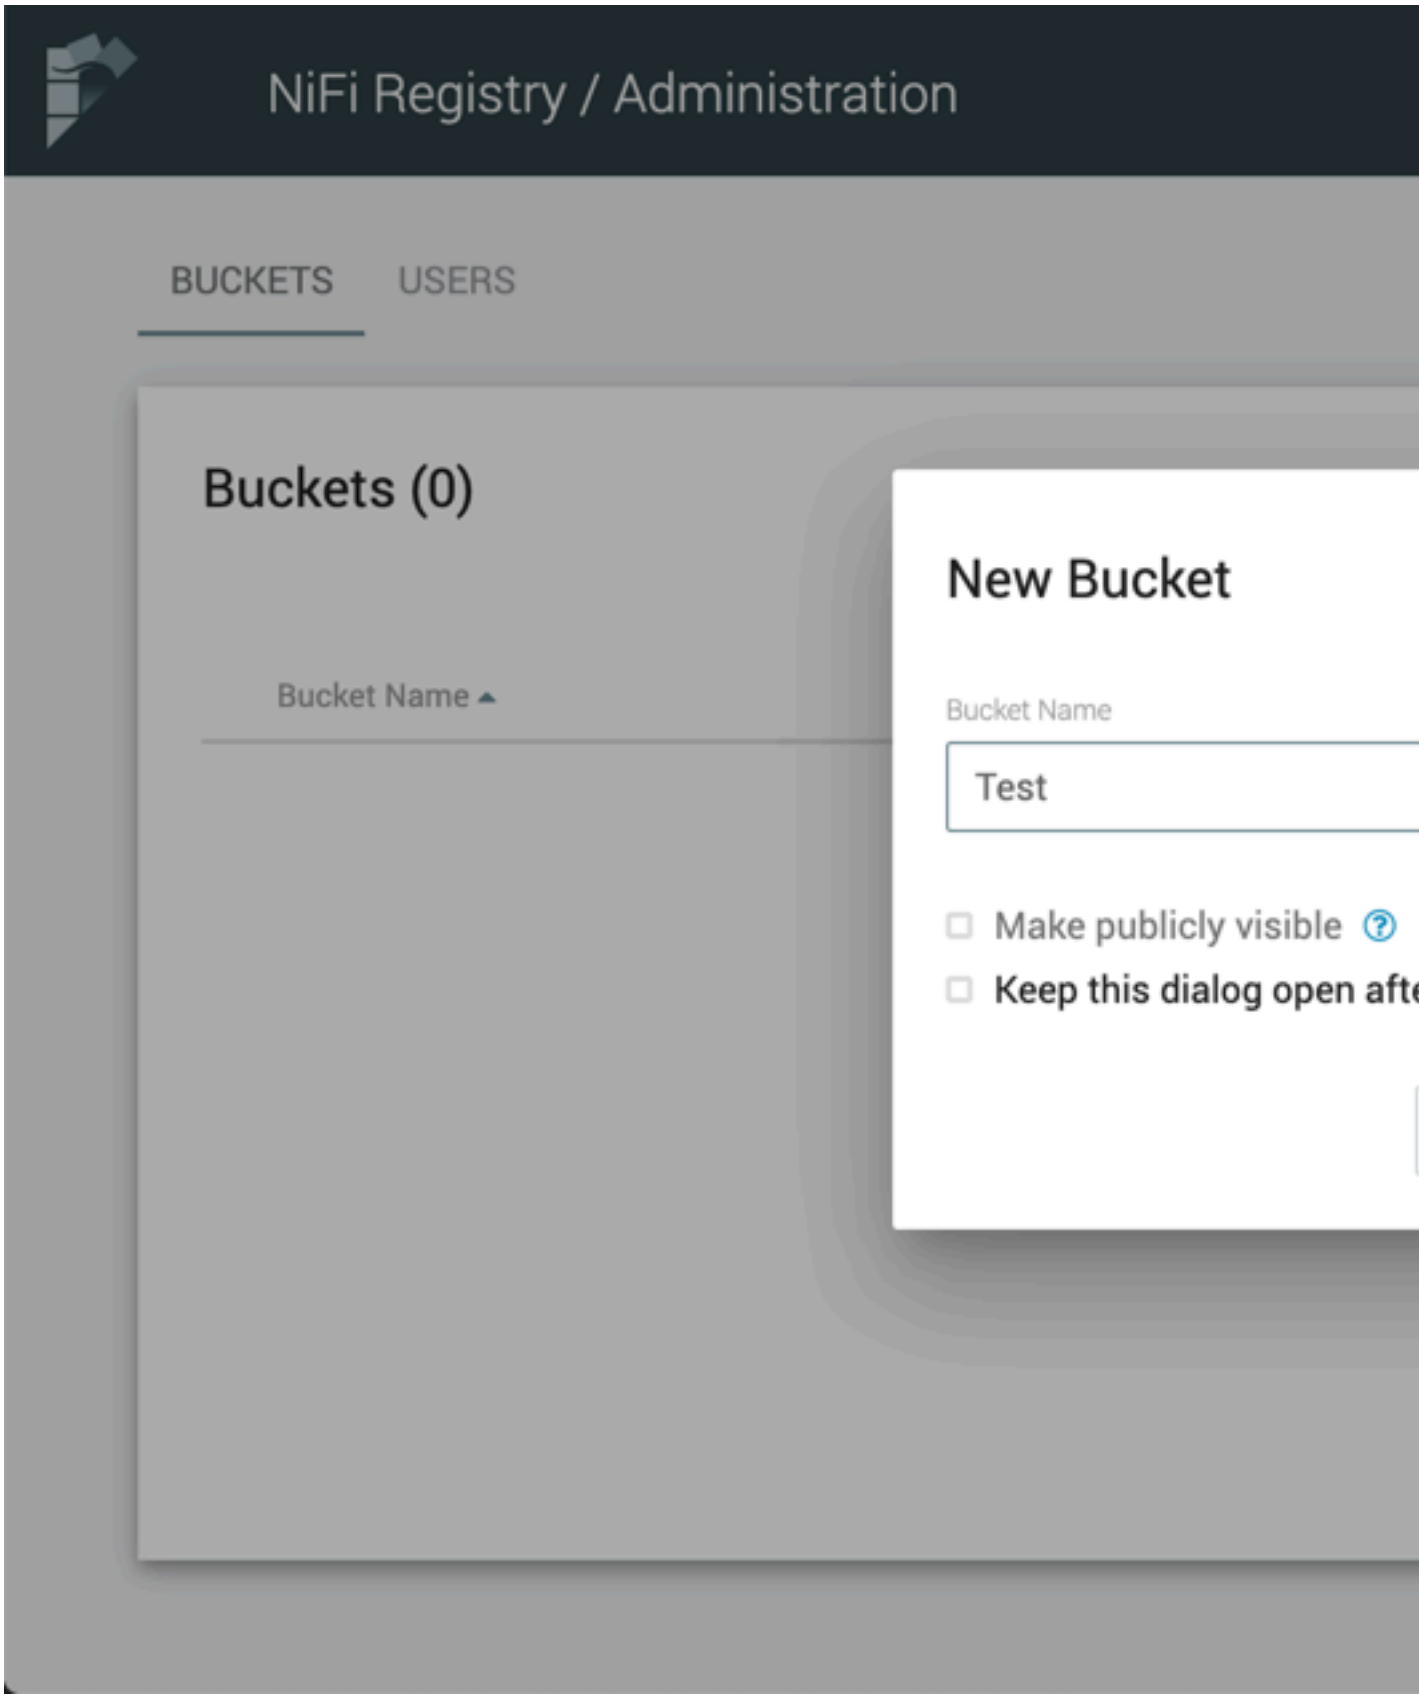

The "Test" bucket is created:

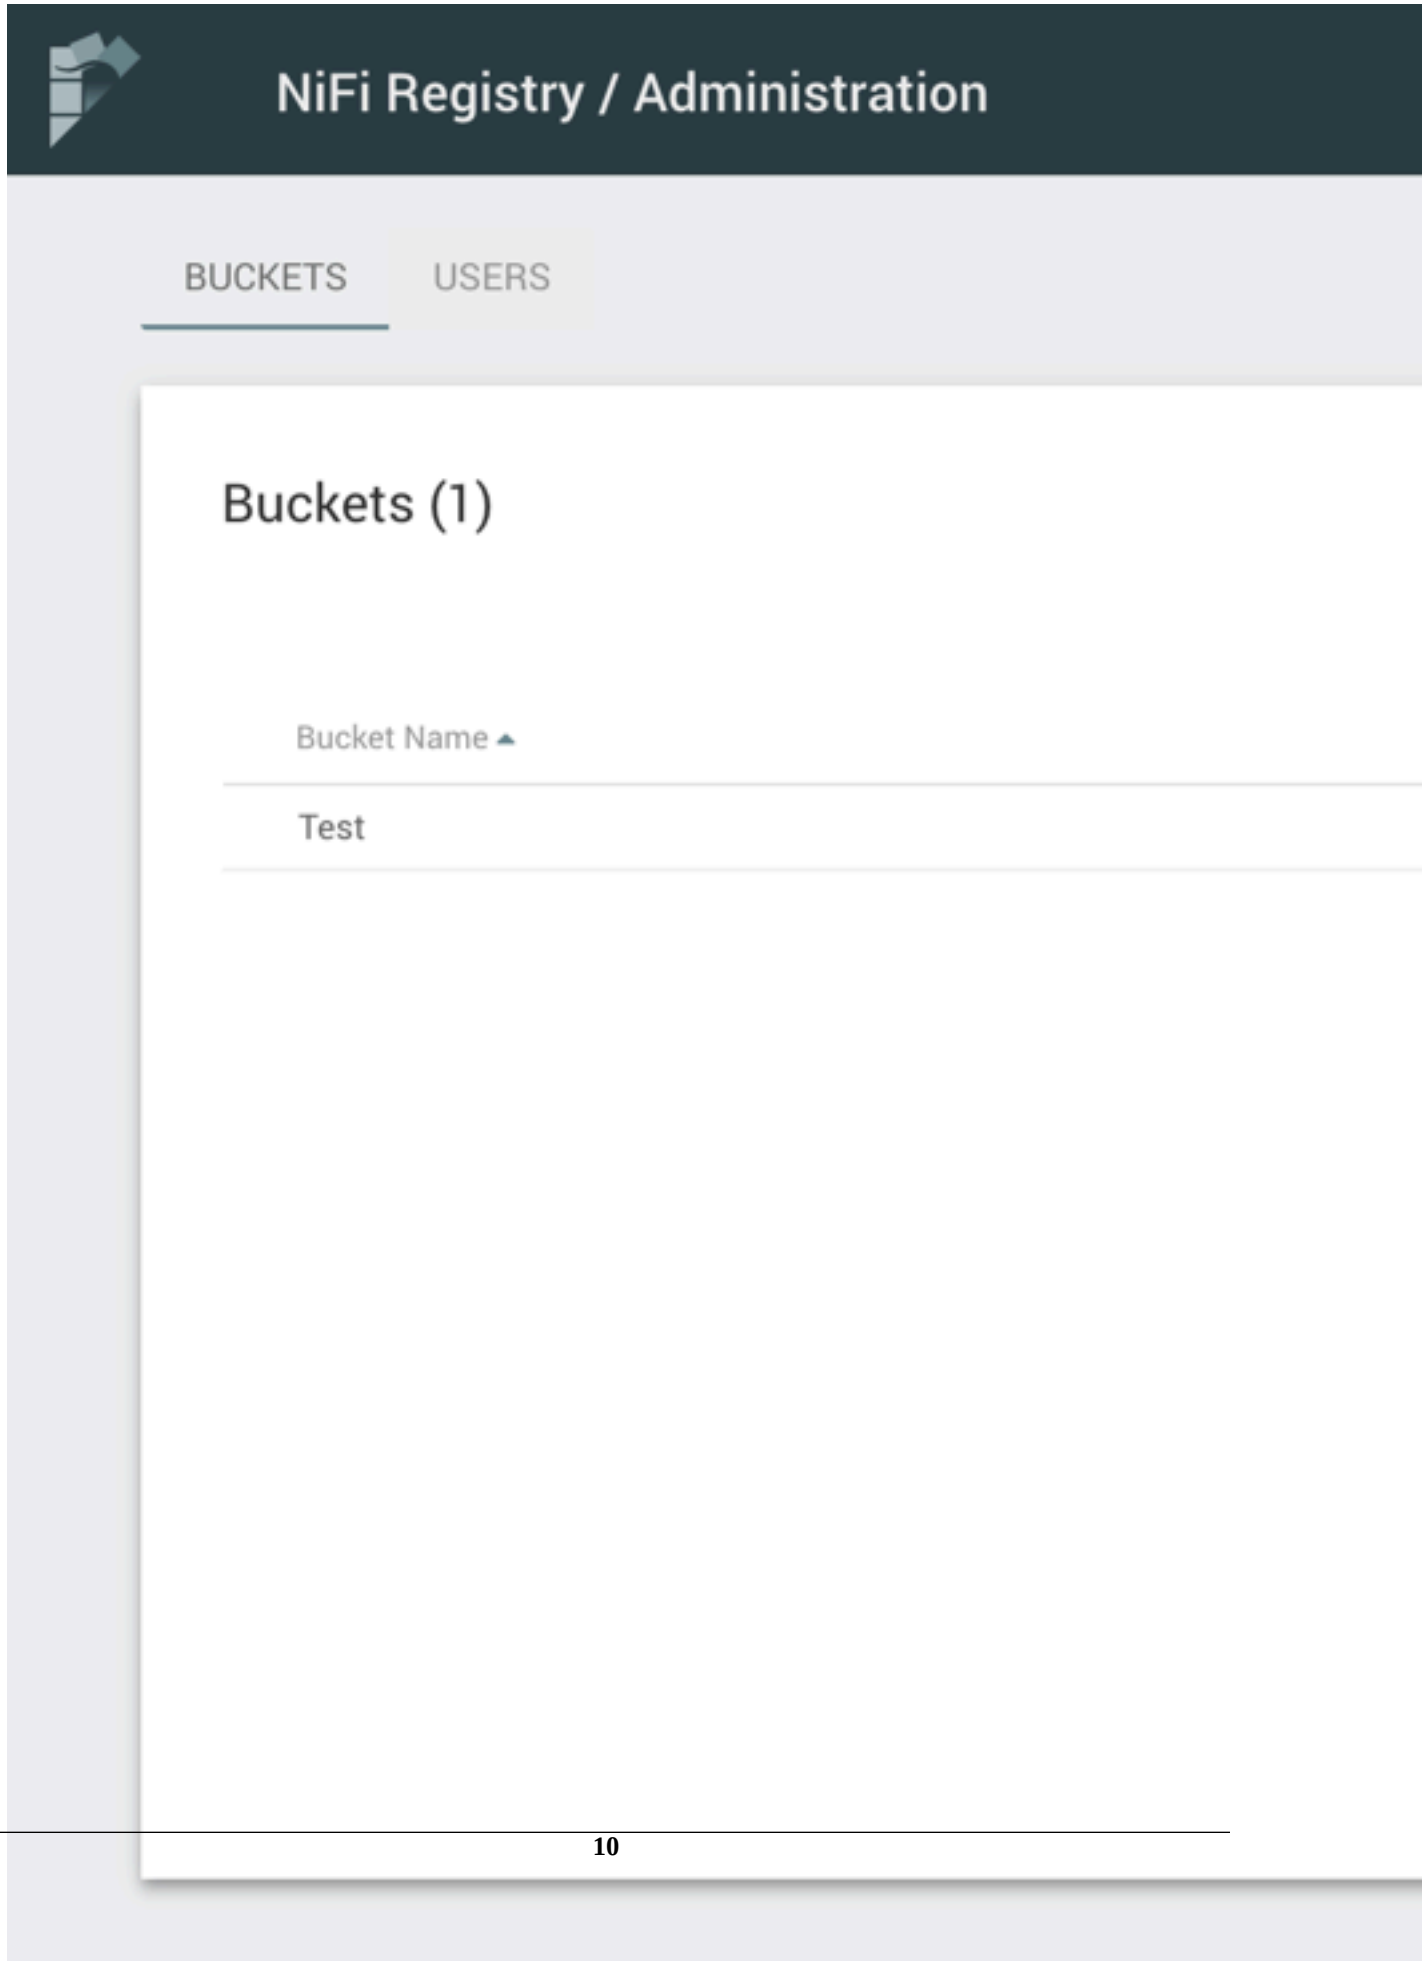
NiFi Registry / Administration
BUCKETS
USERS
Buckets (1)
Bucket Name
Test

There are no permissions configured by default, so anyone is able to view, create and modify buckets in this instance. For information on securing the system, see the System Administrator's Guide.

Connect NiFi to the Registry

Now it is time to tell NiFi about the local registry instance.

Start a NiFi instance if one isn't already running and bring up the UI. Go to controller settings from the top-right menu:

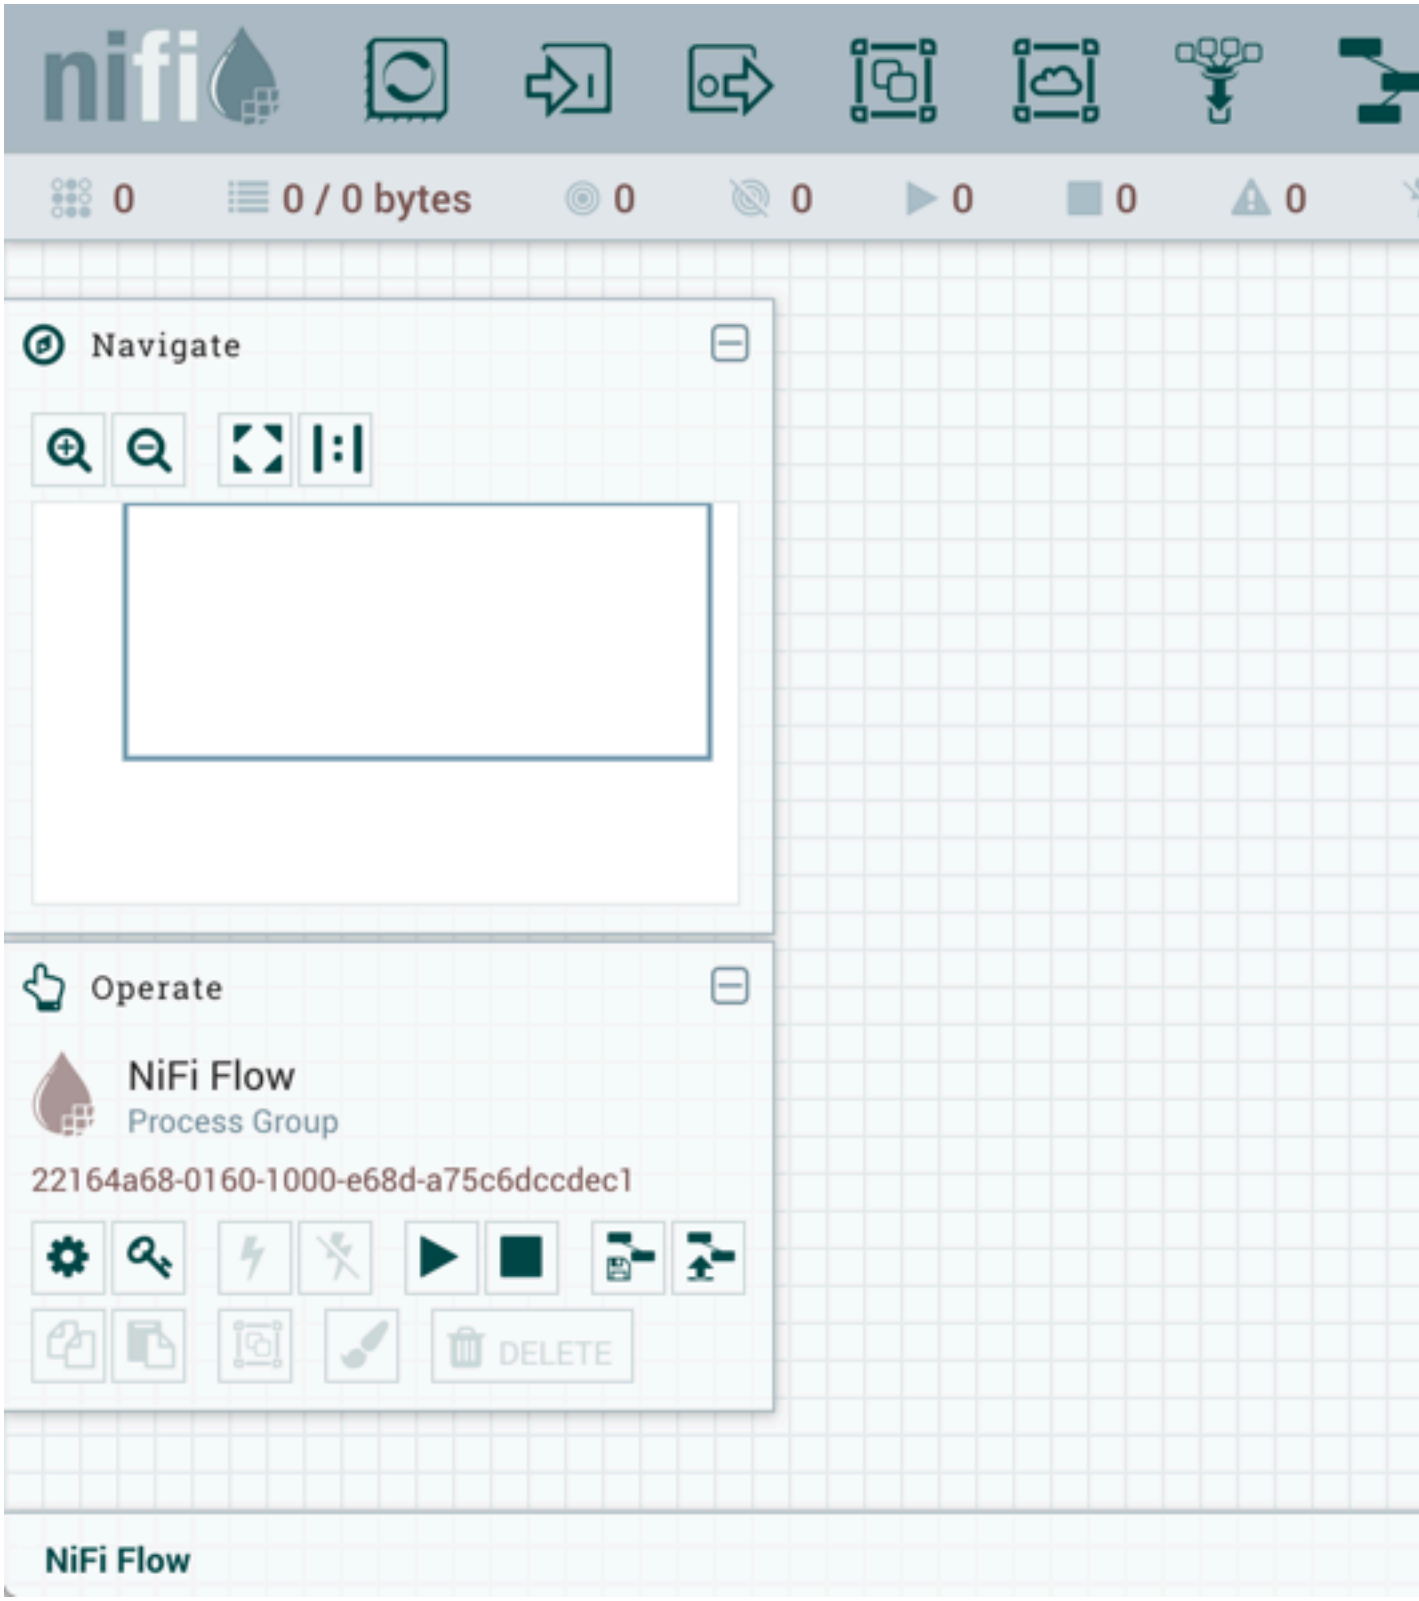

Select the Registry Clients tab and add a new Registry Client giving it a name and the URL of http://localhost:18080::

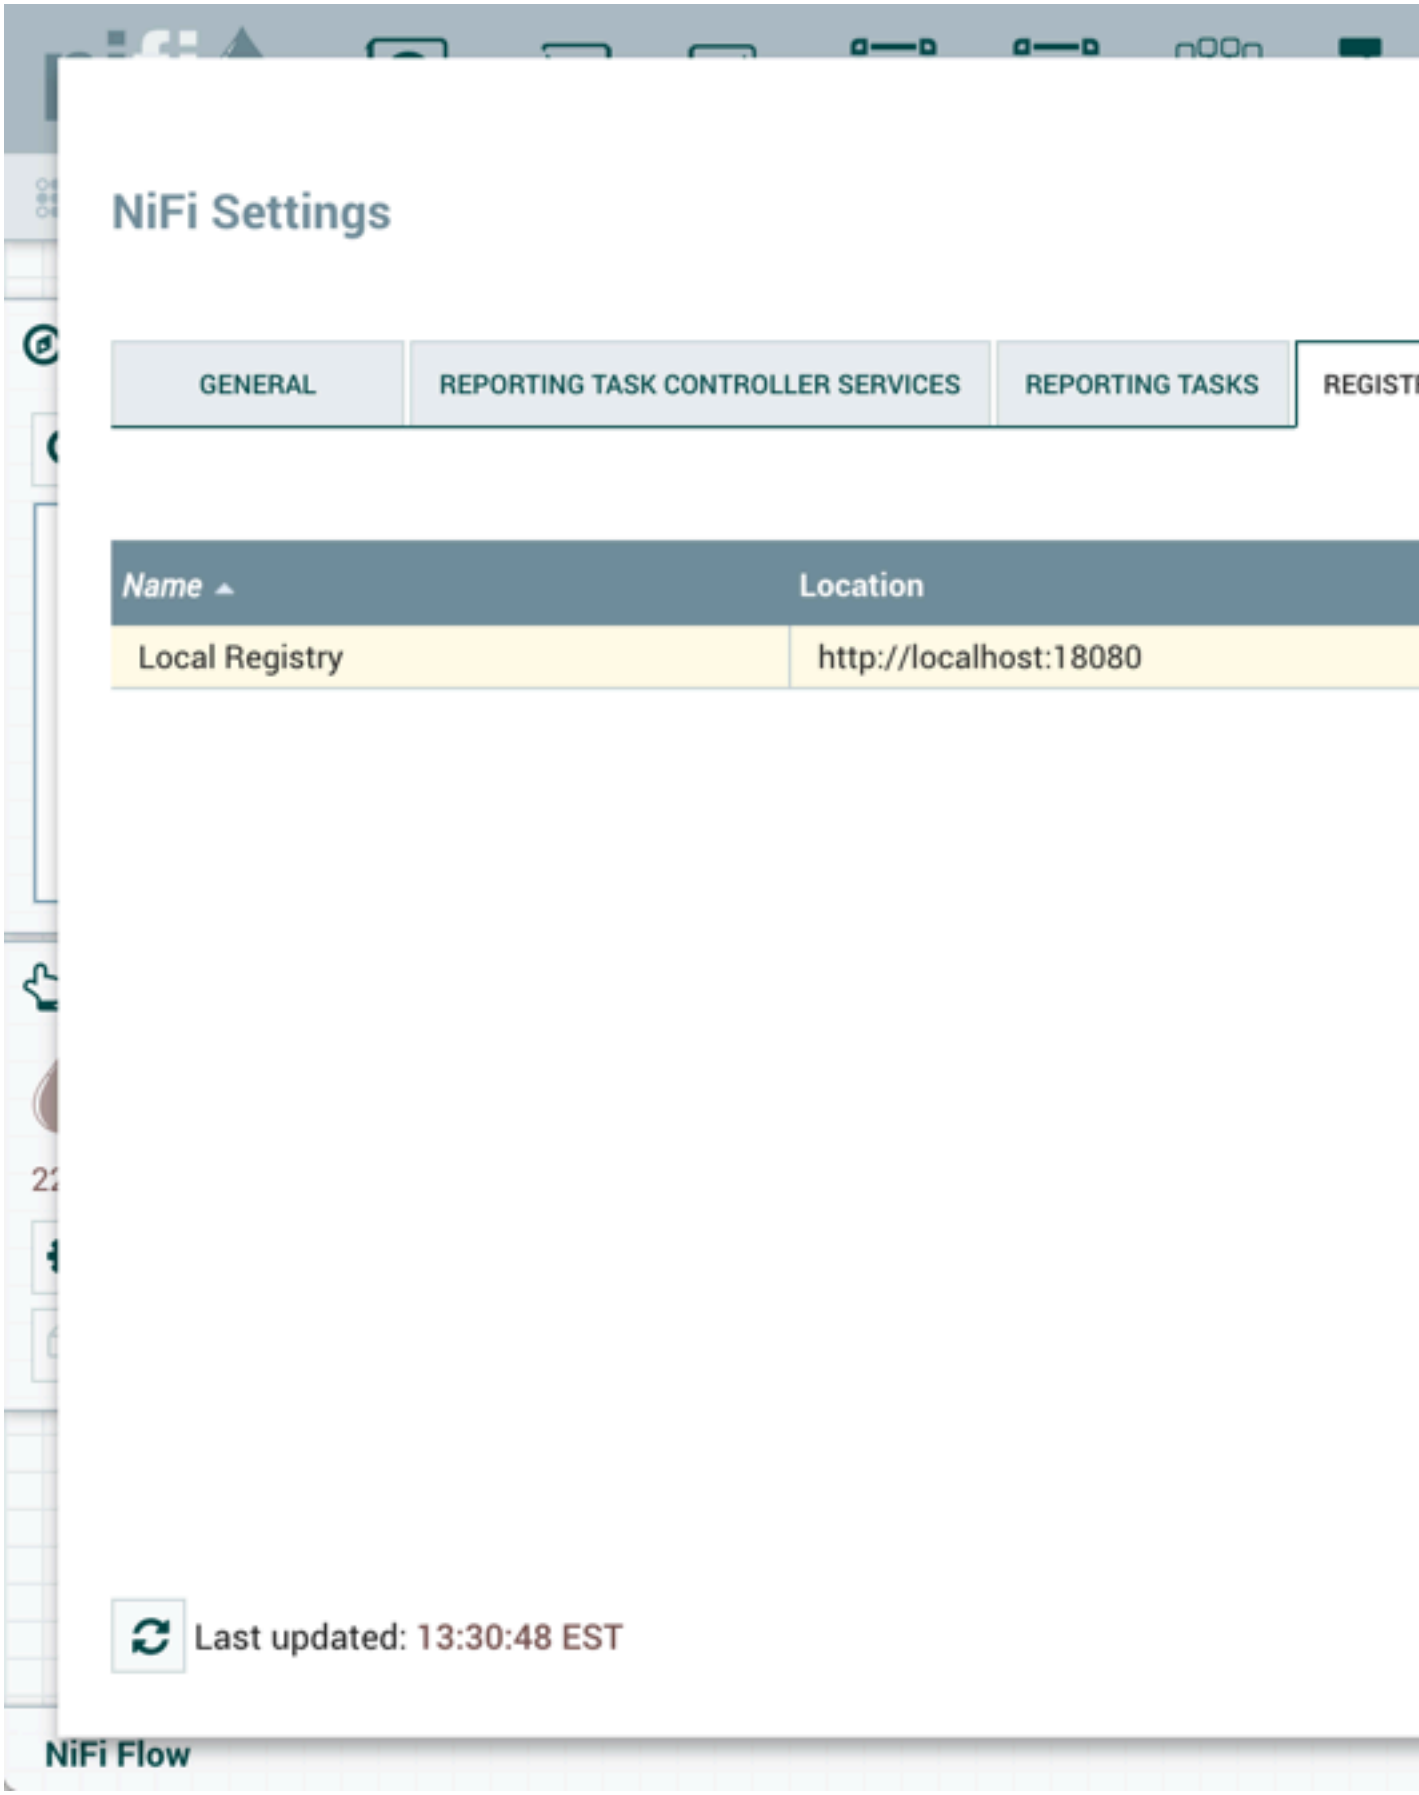
NiFi Settings
GENERAL
REPORTING TASK CONTROLLER SERVICES
REPORTING TASKS
REGISTI
Name
Location
Local Registry
http://localhost:18080
Last updated: 13:30:48 EST
NiFi Flow

Start Version Control on a Process Group

With NiFi connected to a NiFi Registry, dataflows can be version controlled on the process group level.

Right-click on a process group and select "Version#Start version control" from the context menu:

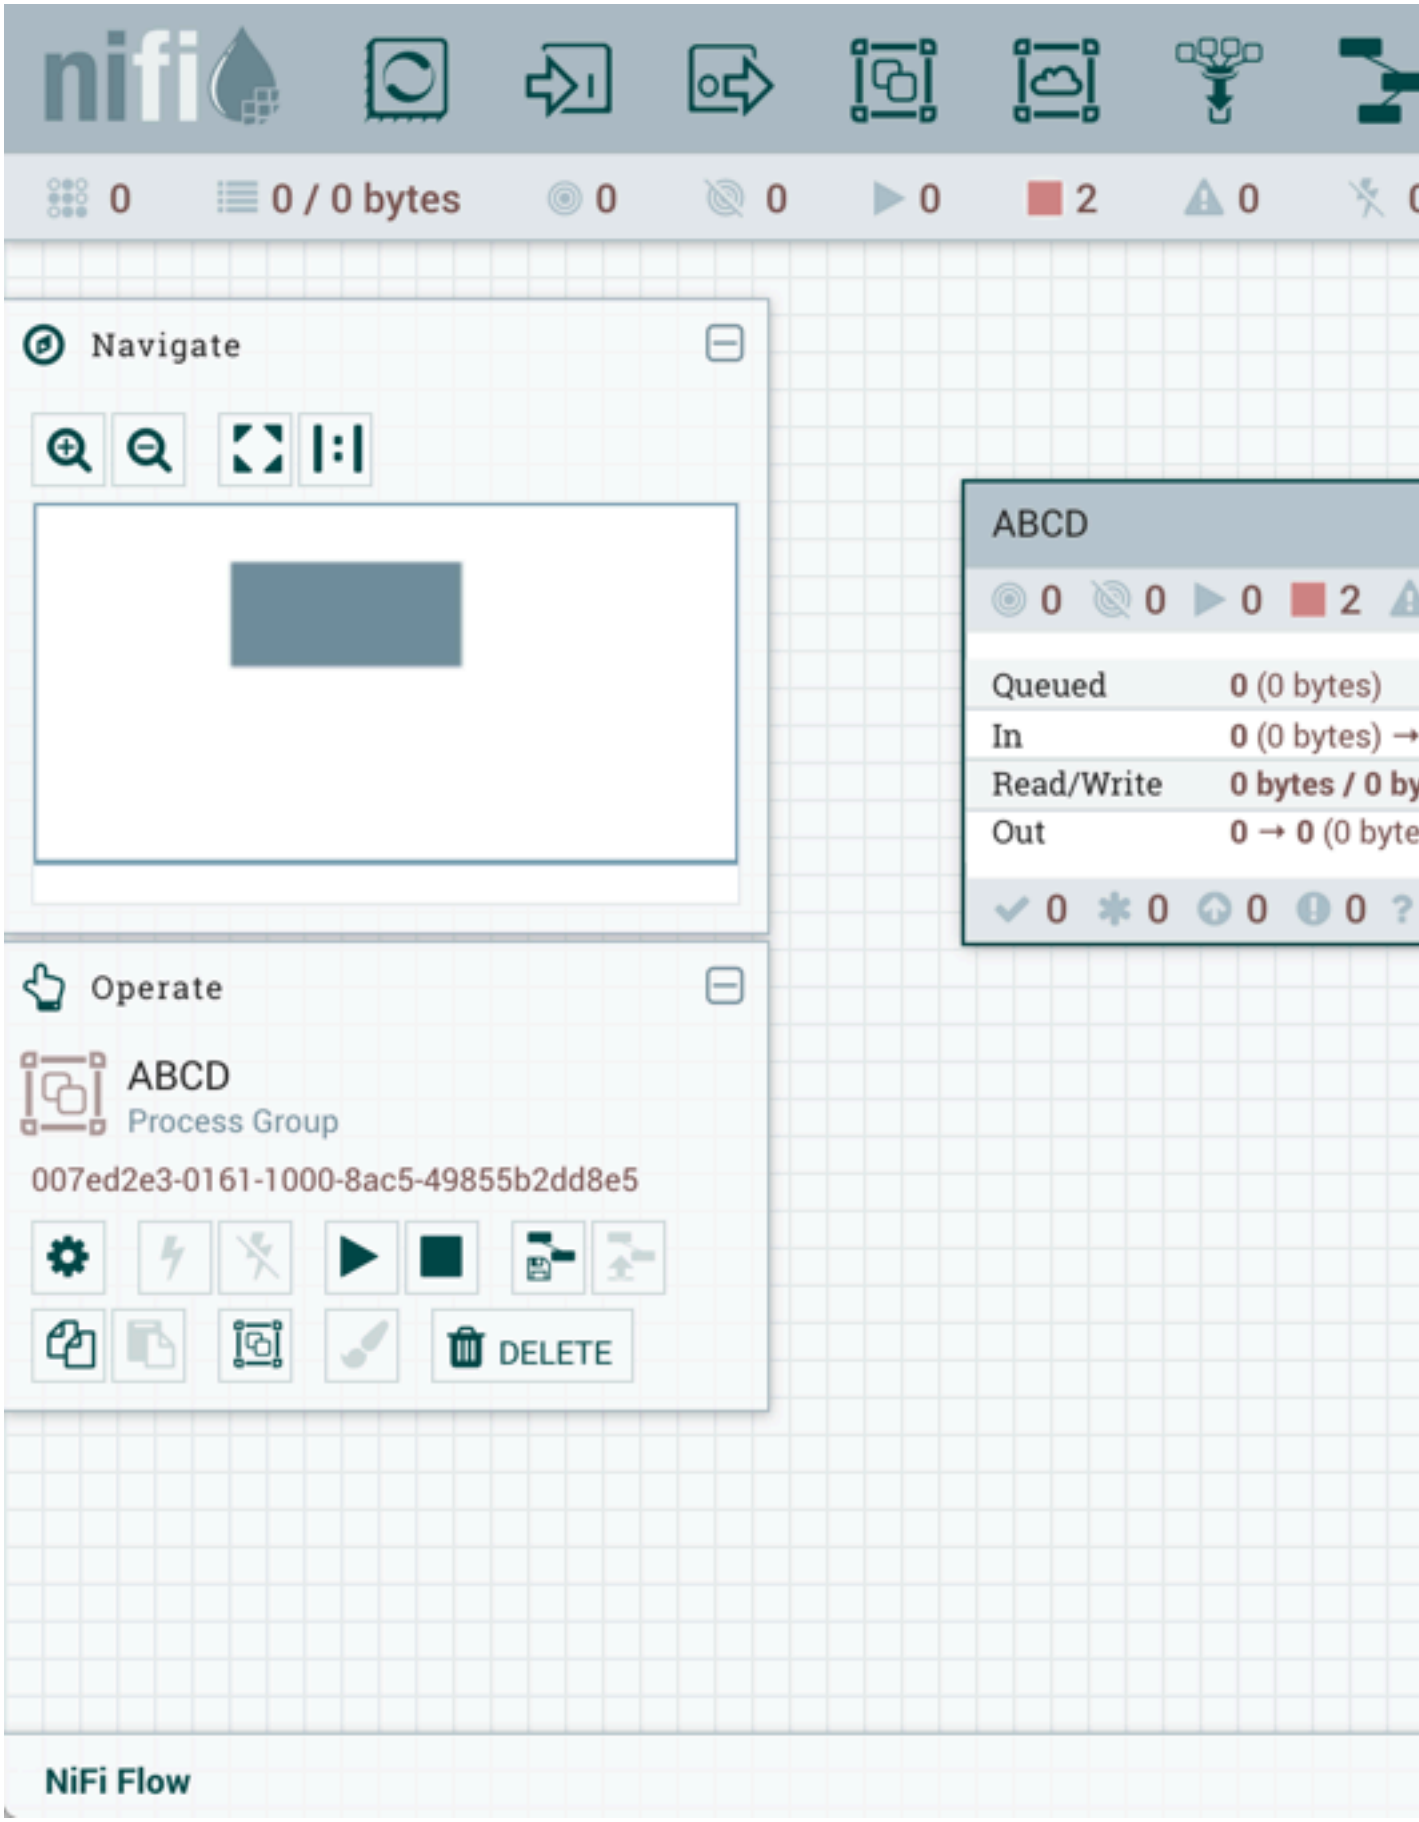
nifi
0
0 / 0 bytes
0
0
0
2
0
Navigate
ABCD
0
0
0
2
Queued
0 (0 bytes)
In
0 (0 bytes)
Read/Write
0 bytes / 0 by
Out
0 → 0 (0 byte
0
0
0
0
Operate
ABCD
Process Group
007ed2e3-0161-1000-8ac5-49855b2dd8e5
DELETE
NiFi Flow

The local registry instance and "Test" bucket are chosen by default to store your flow since they are the only registry connected and bucket available. Enter a flow name, flow description, comments and select "Save":

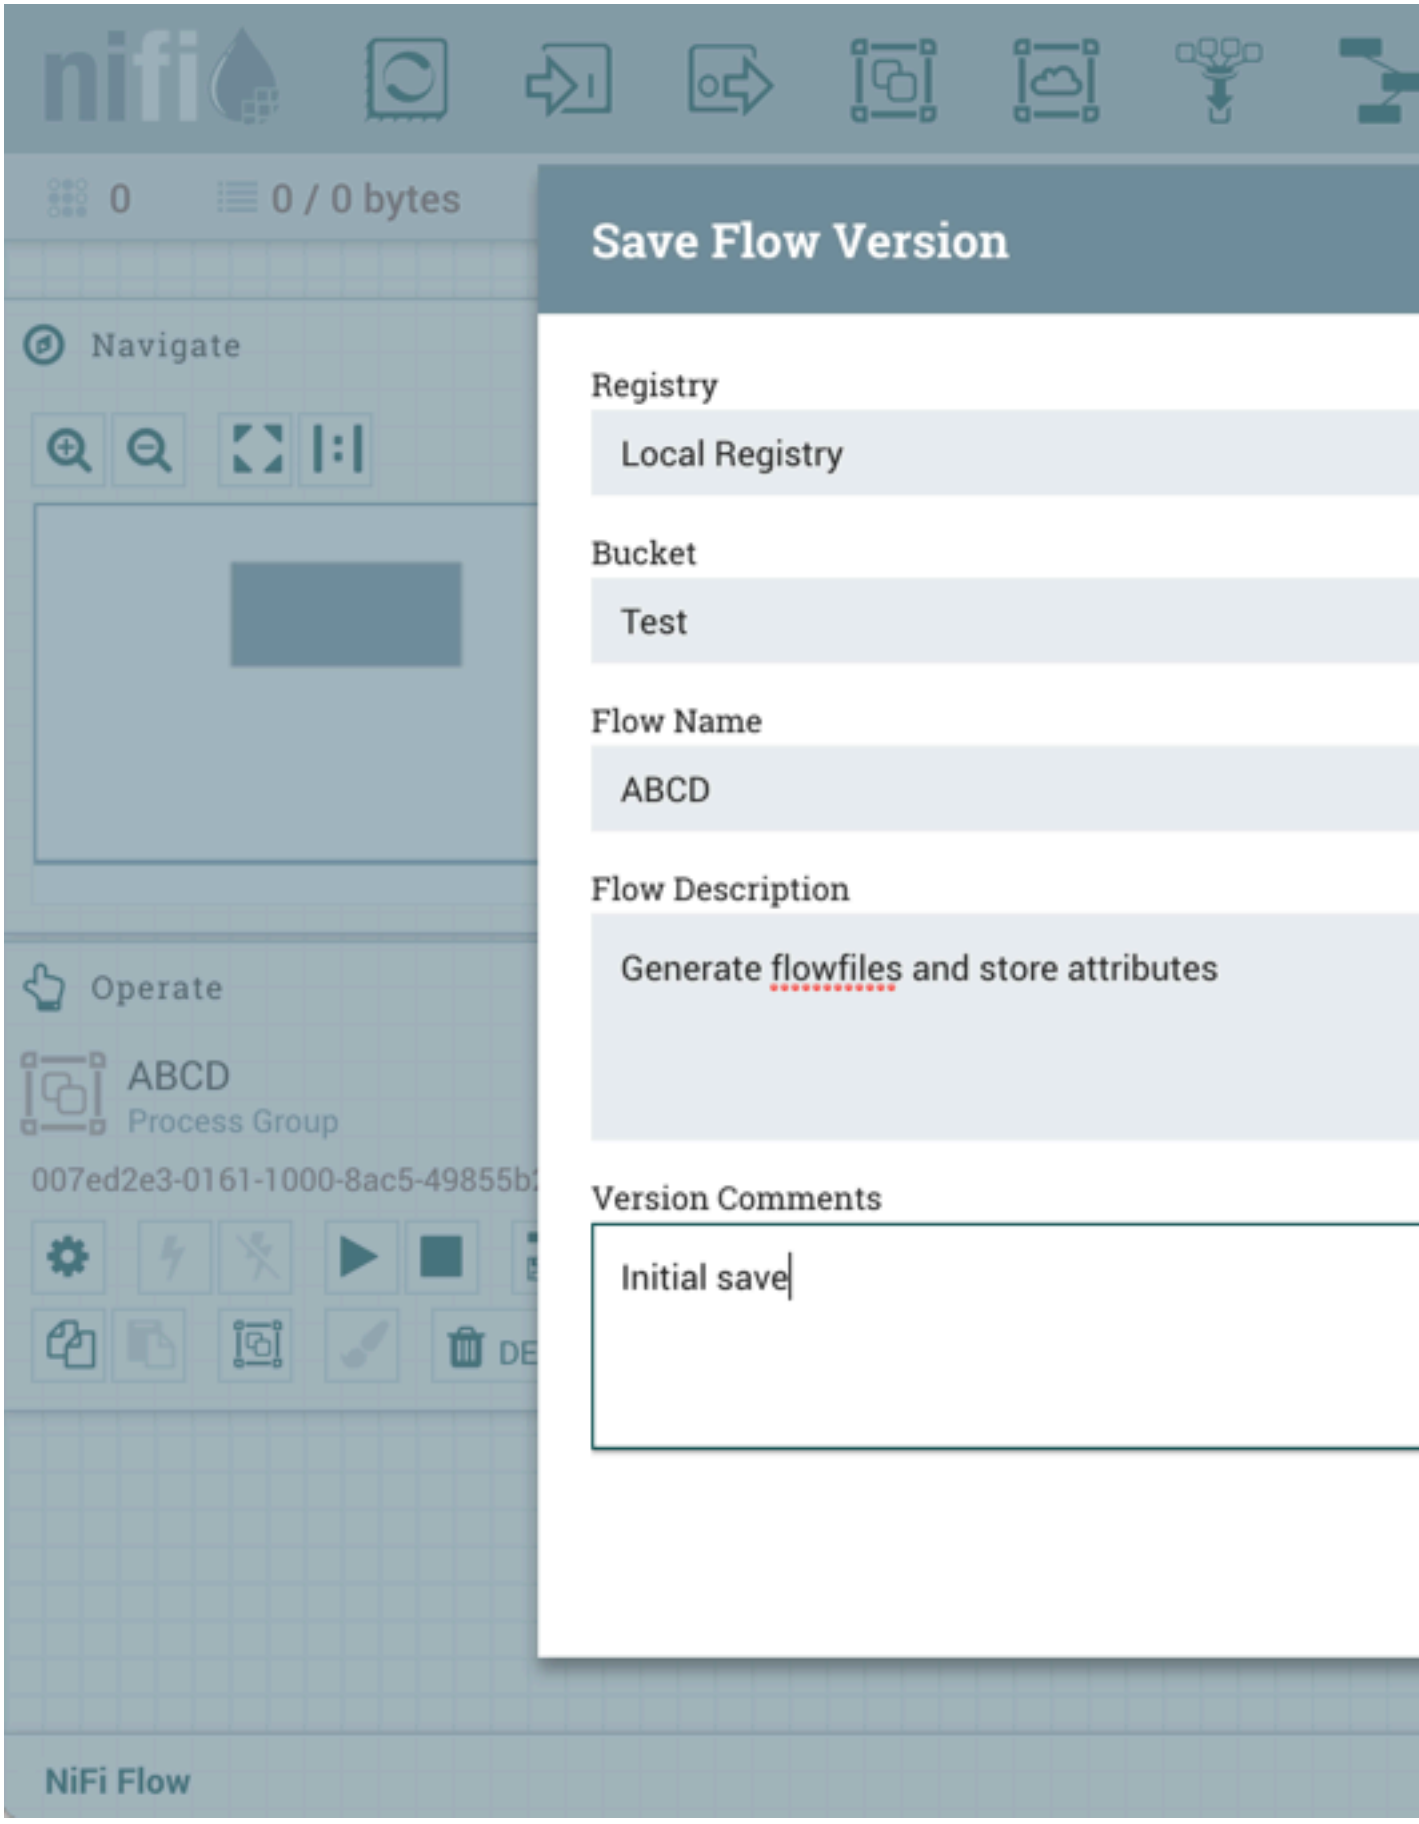
nifi
0
0 / 0 bytes
Navigate
Operate
ABCD
Process Group
007ed2e3-0161-1000-8ac5-49855b
NiFi Flow
Save Flow Version
Registry
Local Registry
Bucket
Test
Flow Name
ABCD
Flow Description
Generate flowfiles and store attributes
Version Comments
Initial save

As indicated by the Version State icon () in the top left corner of the component, the process group is now saved as a versioned flow in the registry.

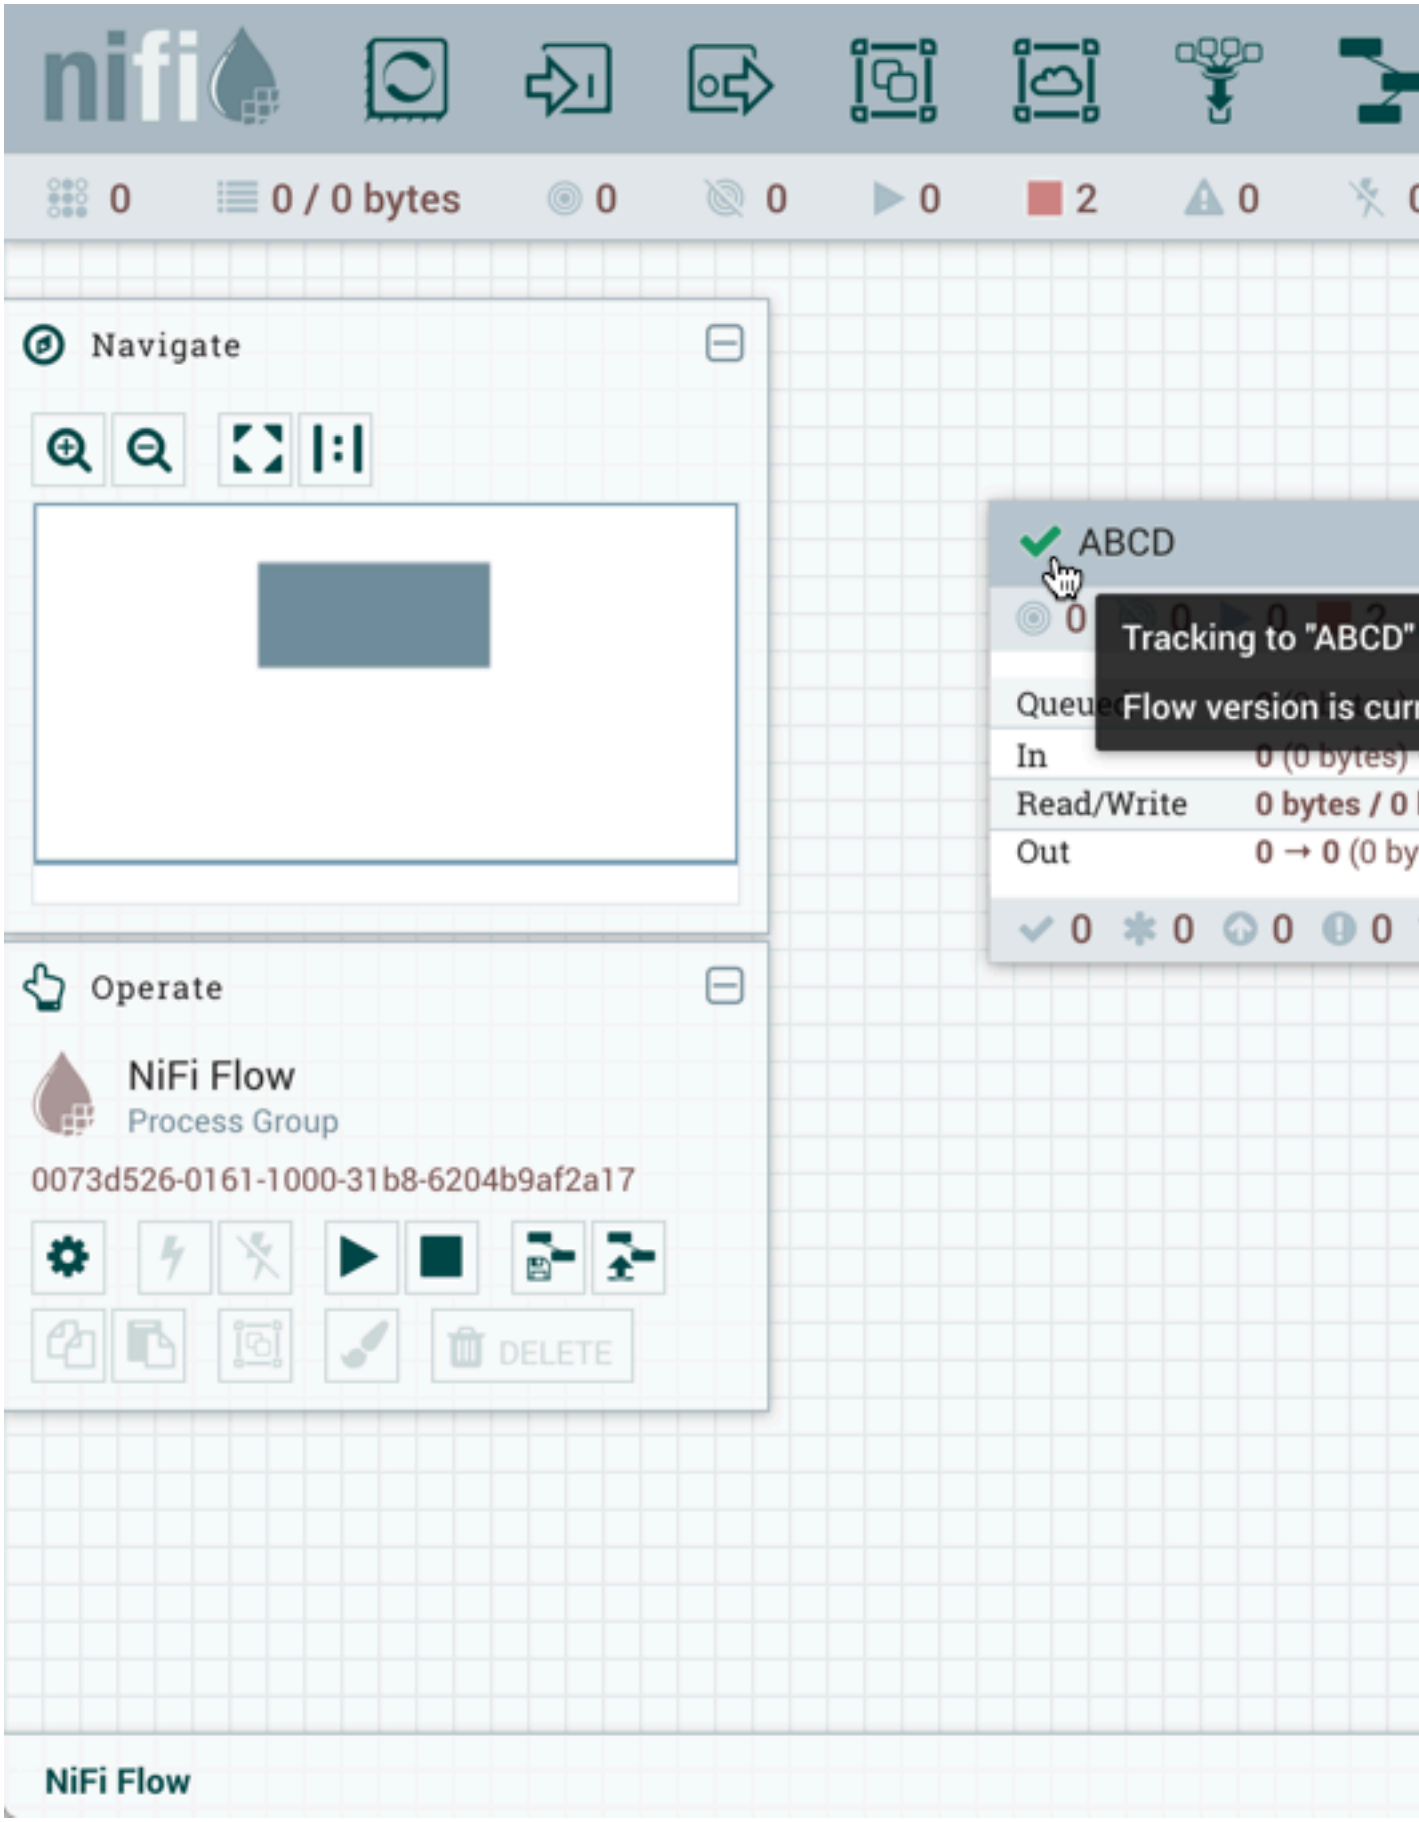
nifi
0
0 / 0 bytes
0
0
0
2
0
Navigate
Operate
NiFi Flow
Process Group
0073d526-0161-1000-31b8-6204b9af2a17
DELETE
ABCD
Tracking to "ABCD"
Flow version is cur
In
0 (0 bytes)
Read/Write
0 bytes / 0
Out
0 → 0 (0 by
0
0
0
0
NiFi Flow

Go back to the Registry UI and return to the main page to see the versioned flow you just saved (a refresh may be required):

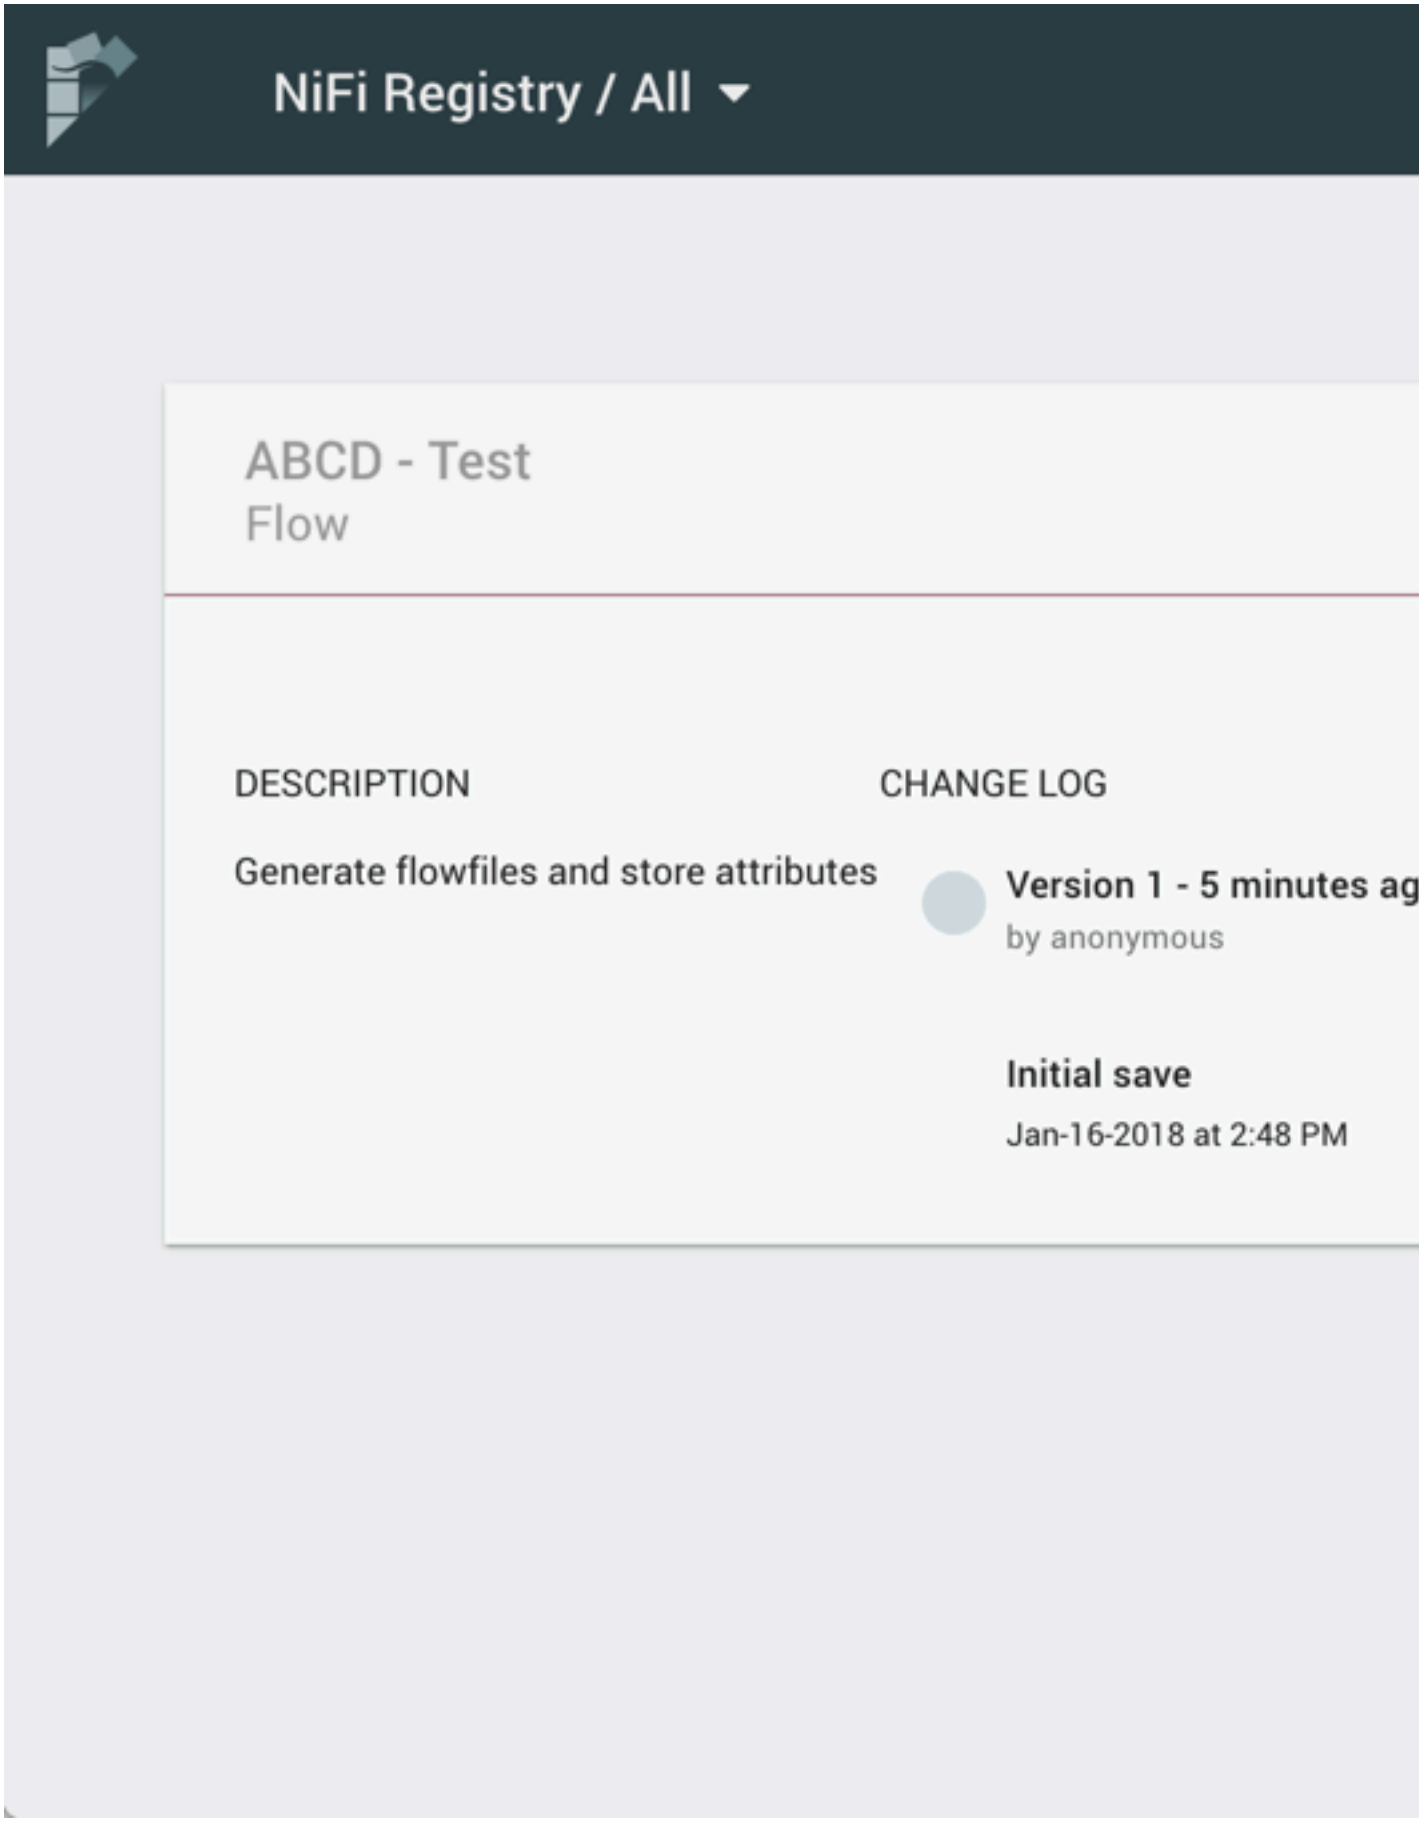
NiFi Registry / All
ABCD - Test
Flow
DESCRIPTION
CHANGE LOG
Generate flowfiles and store attributes
Version 1 - 5 minutes ag
by anonymous
Initial save
Jan-16-2018 at 2:48 PM

Save Changes to a Versioned Flow

Changes made to the versioned process group can be reviewed, reverted or saved.

For example, if changes are made to the ABCD flow, the Version State changes to "Locally modified" (). The right-click menu will now show the options "Commit local changes", "Show local changes" or "Revert local changes":

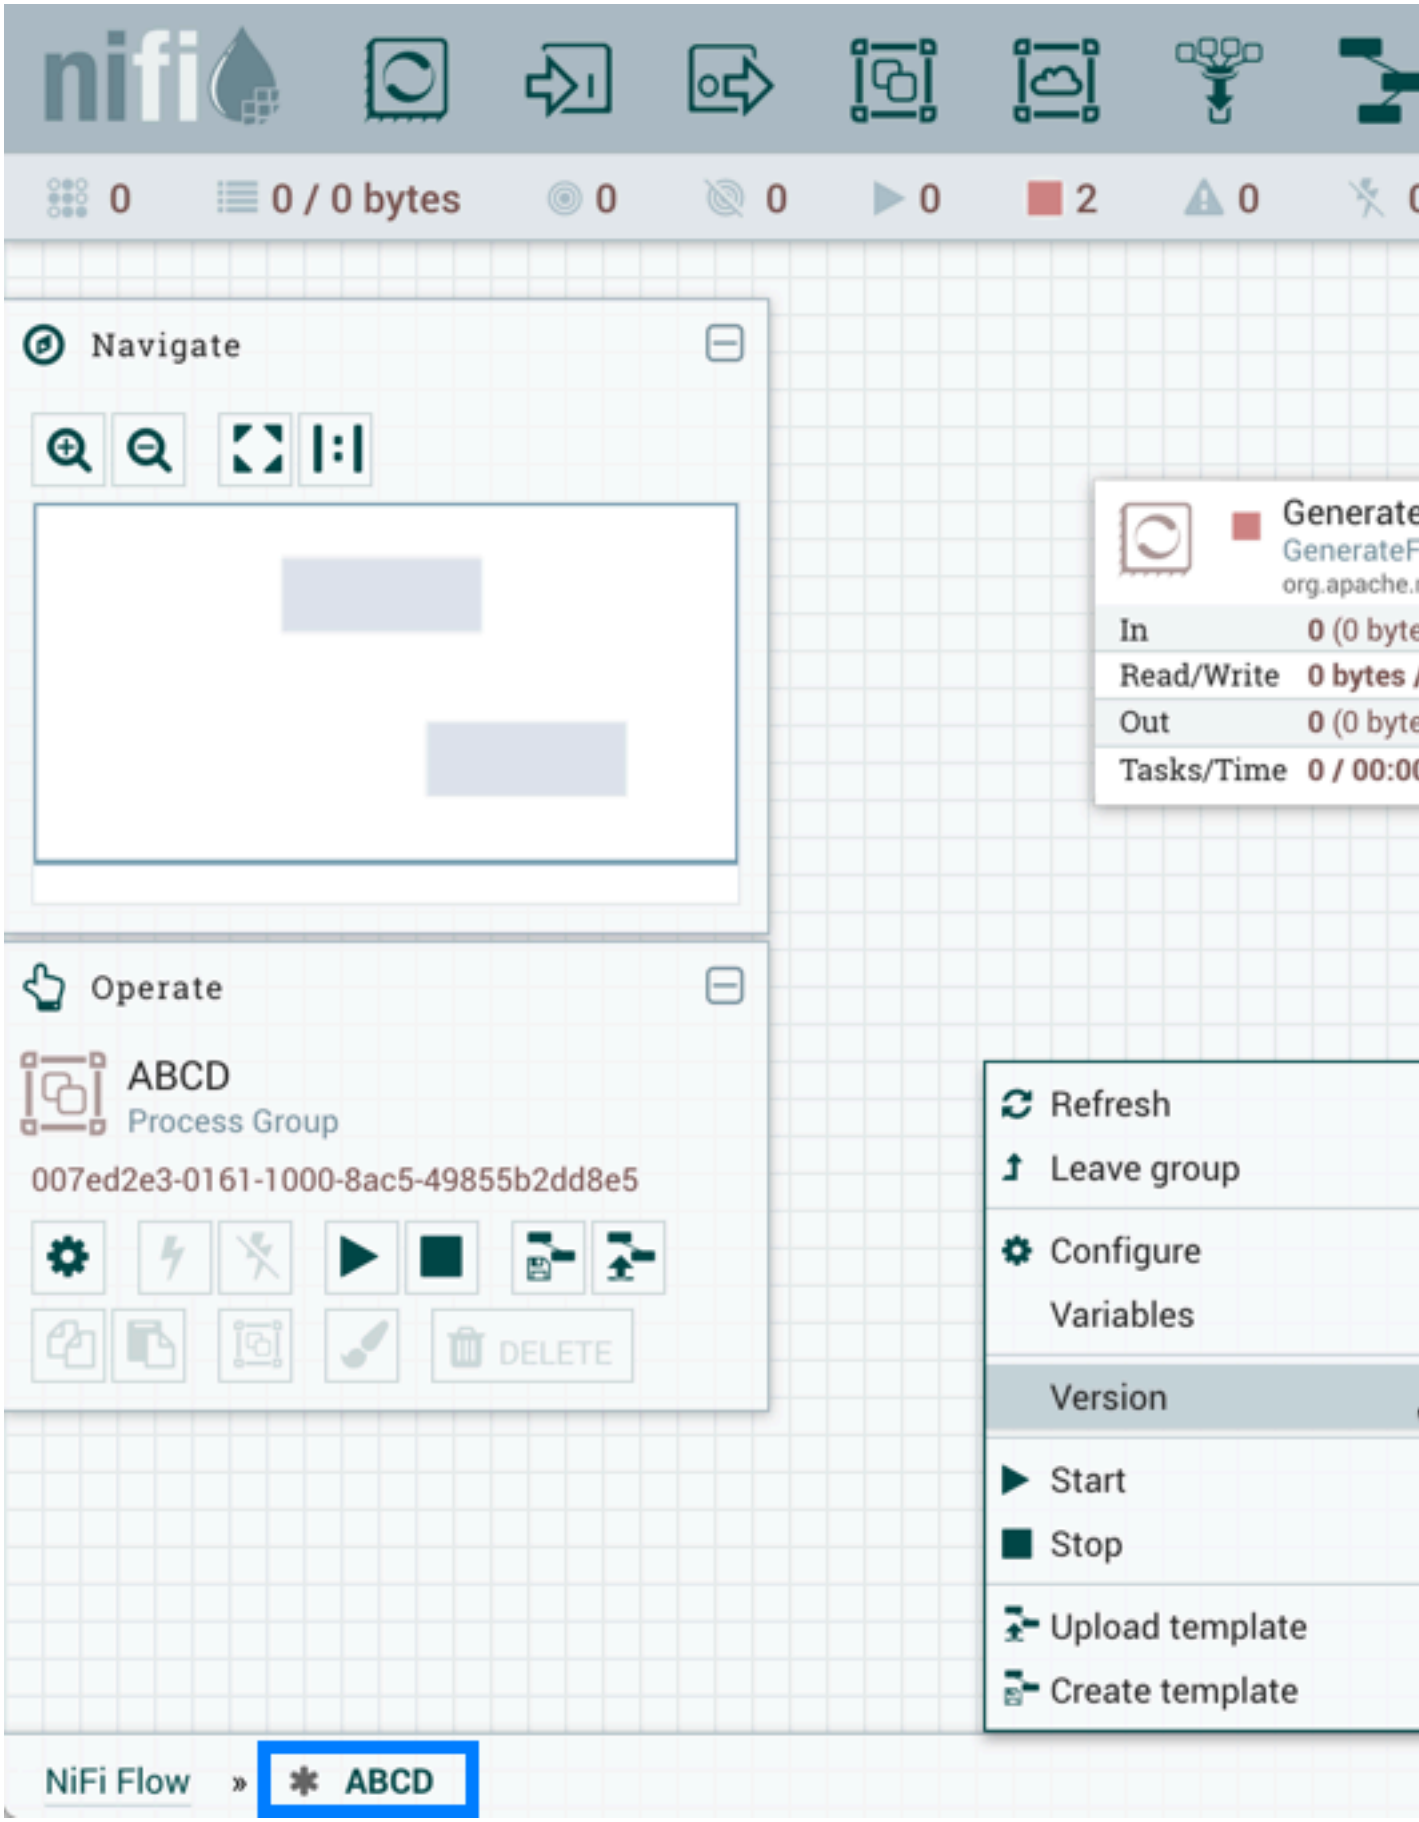
nifi
0
0 / 0 bytes
0
0
0
2
0
Navigate
Operate
ABCD
Process Group
007ed2e3-0161-1000-8ac5-49855b2dd8e5
DELETE
Generate
In
Read/Write
Out
Tasks/Time
Refresh
Leave group
Configure
Variables
Version
Start
Stop
Upload template
Create template
NiFi Flow
»
ABCD

Select "Show local changes" to see the details of the changes made:

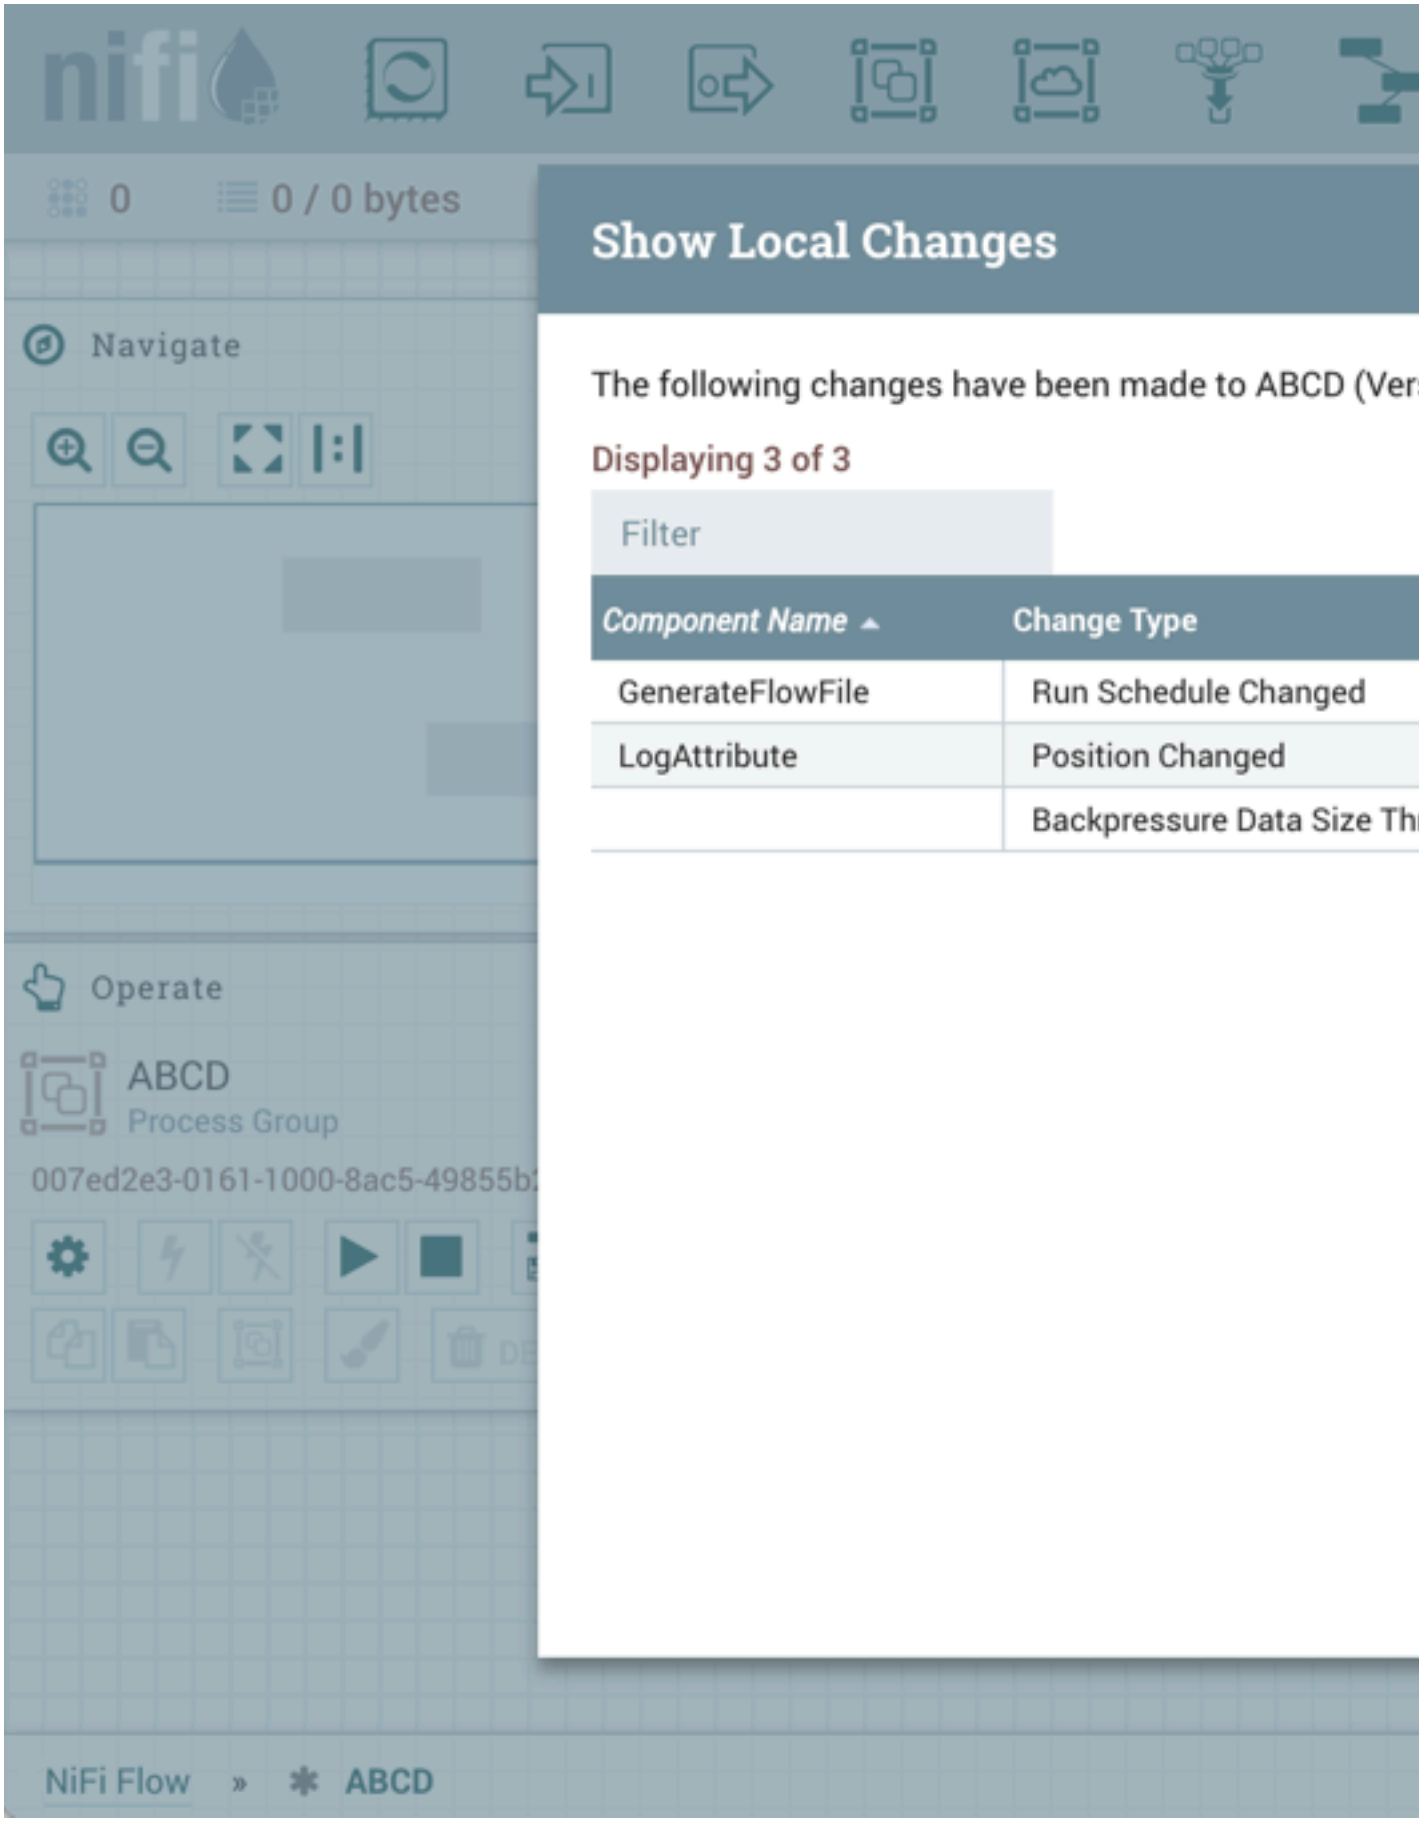
0
0 / 0 bytes
Navigate
Operate
ABCD
Process Group
007ed2e3-0161-1000-8ac5-49855b
NiFi Flow
»
ABCD
Show Local Changes
The following changes have been made to ABCD (Ver
Displaying 3 of 3
Filter
Component Name
Change Type
GenerateFlowFile
Run Schedule Changed
LogAttribute
Position Changed
Backpressure Data Size Th

Select "Commit local changes", enter comments and select "Save" to save the changes:

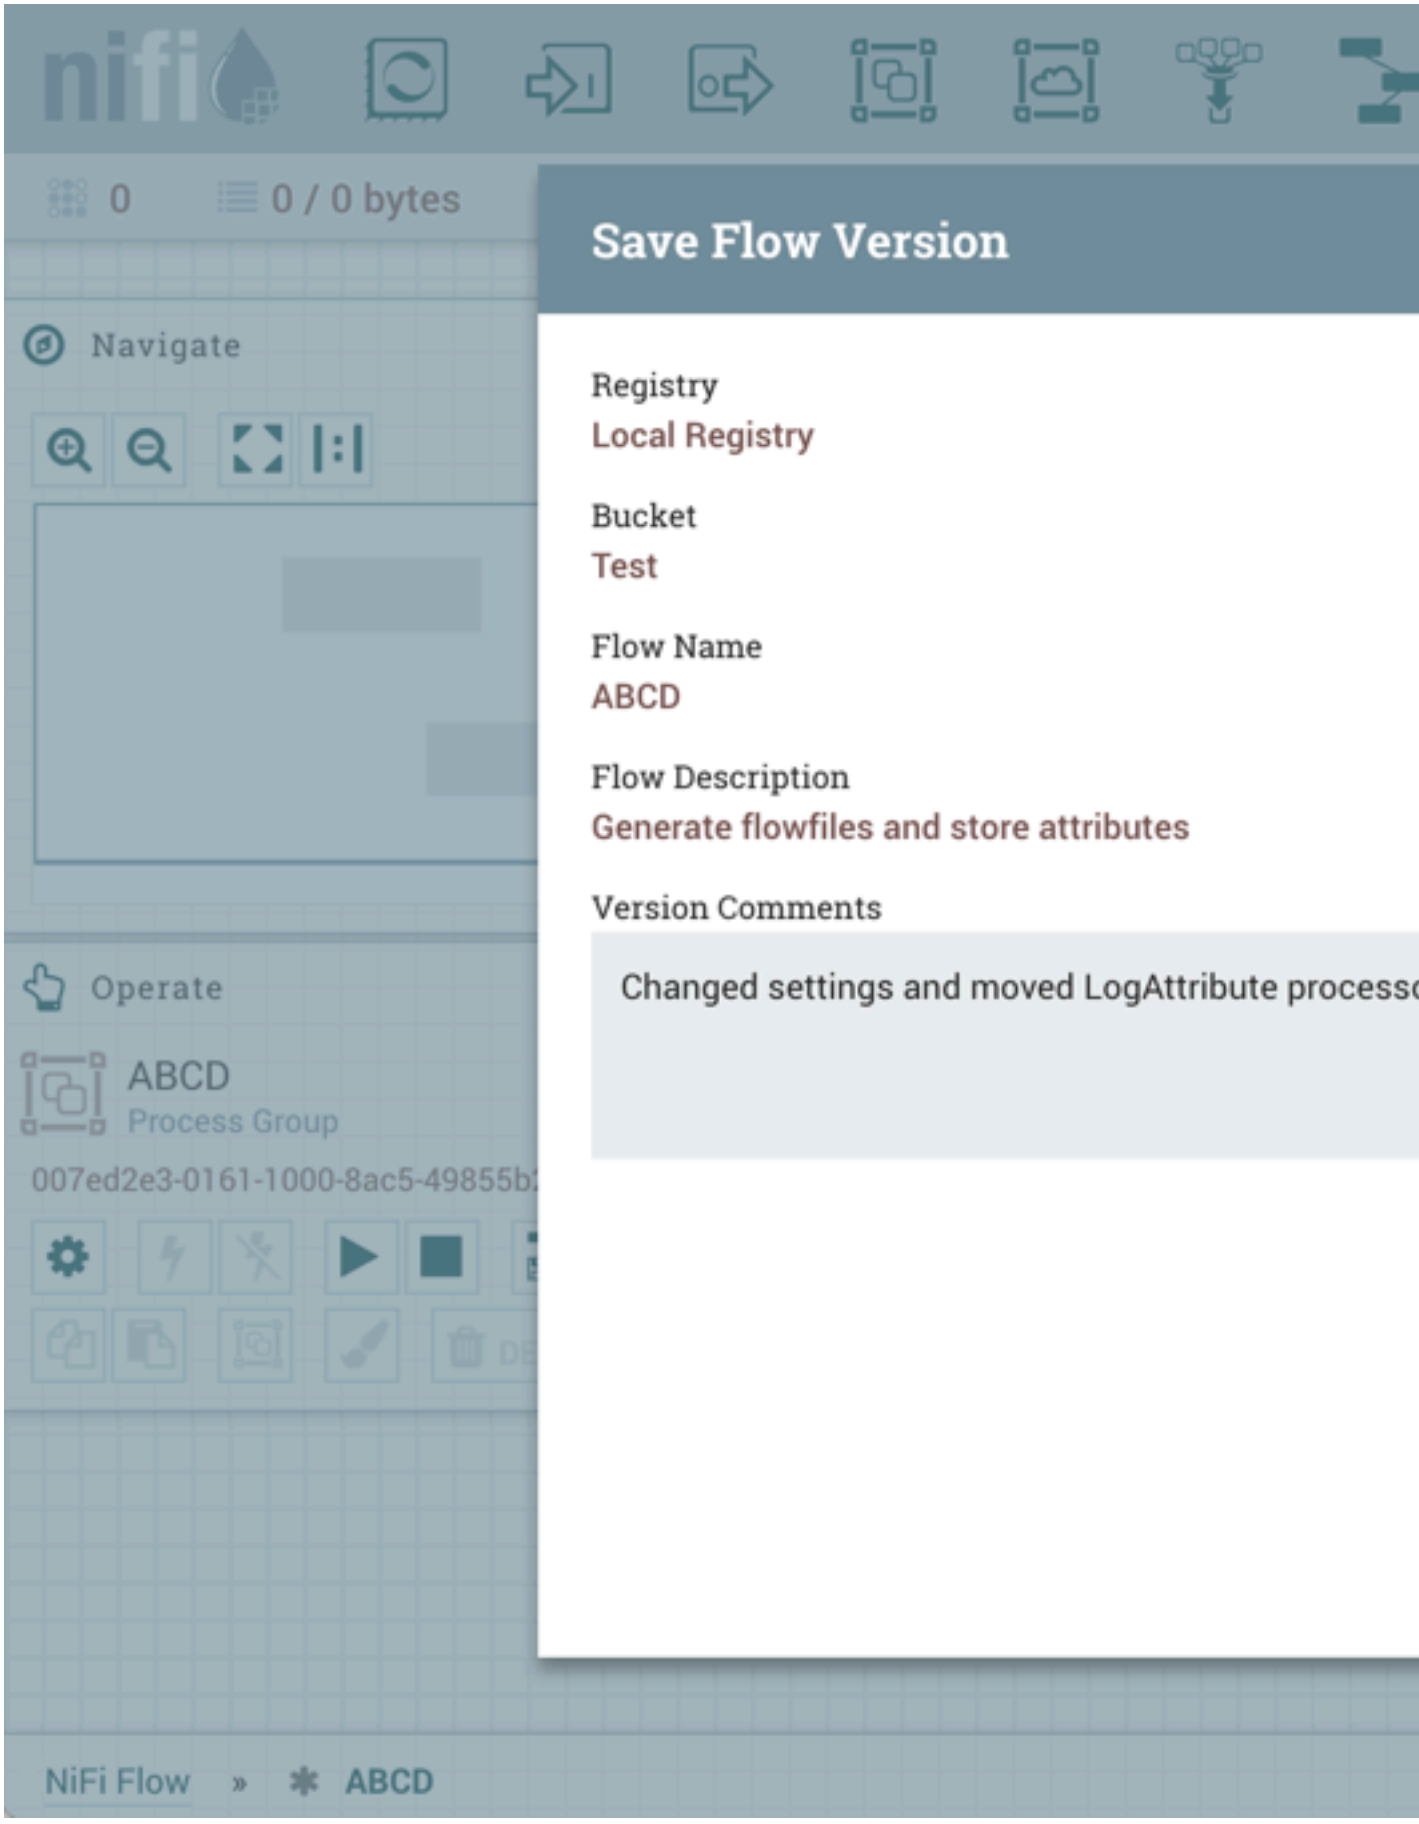
nifi
0
0 / 0 bytes
Navigate
Operate
ABCD
Process Group
007ed2e3-0161-1000-8ac5-49855b
NiFi Flow
»
ABCD
Save Flow Version
Registry
Local Registry
Bucket
Test
Flow Name
ABCD
Flow Description
Generate flowfiles and store attributes
Version Comments
Changed settings and moved LogAttribute processo

Version 2 of the flow is saved:

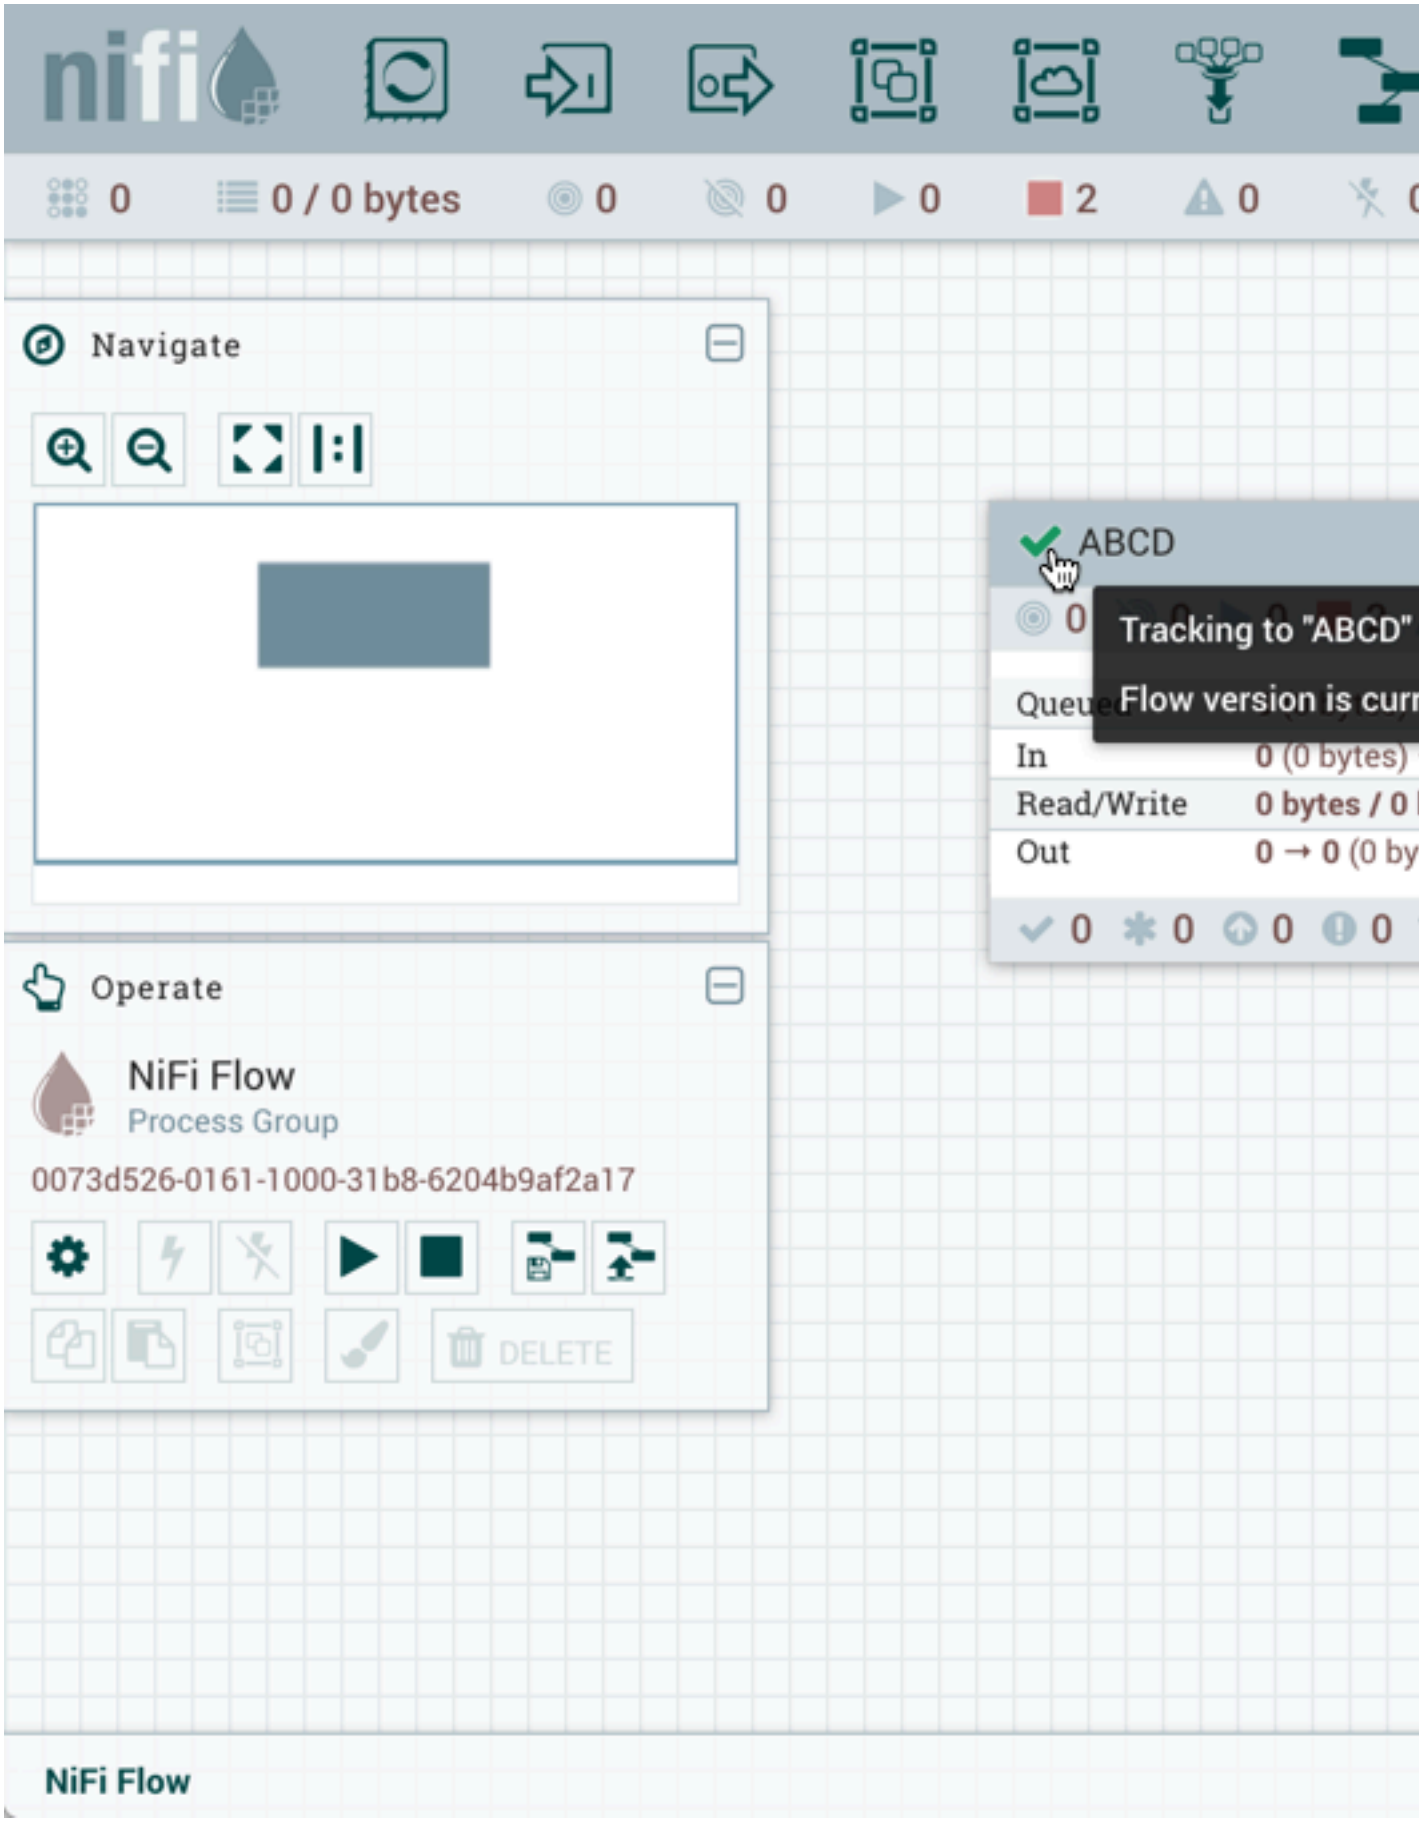
nifi
0
0 / 0 bytes
0
0
0
2
0
Navigate
Operate
NiFi Flow
Process Group
0073d526-0161-1000-31b8-6204b9af2a17
DELETE
ABCD
Tracking to "ABCD"
Flow version is curr
In
0 (0 bytes)
Read/Write
0 bytes / 0
Out
0 → 0 (0 by
NiFi Flow

Import a Versioned Flow

With a flow existing in the registry, we can use it to illustrate how to import a versioned process group.

In NiFi, select Process Group from the Components toolbar and drag it onto the canvas:

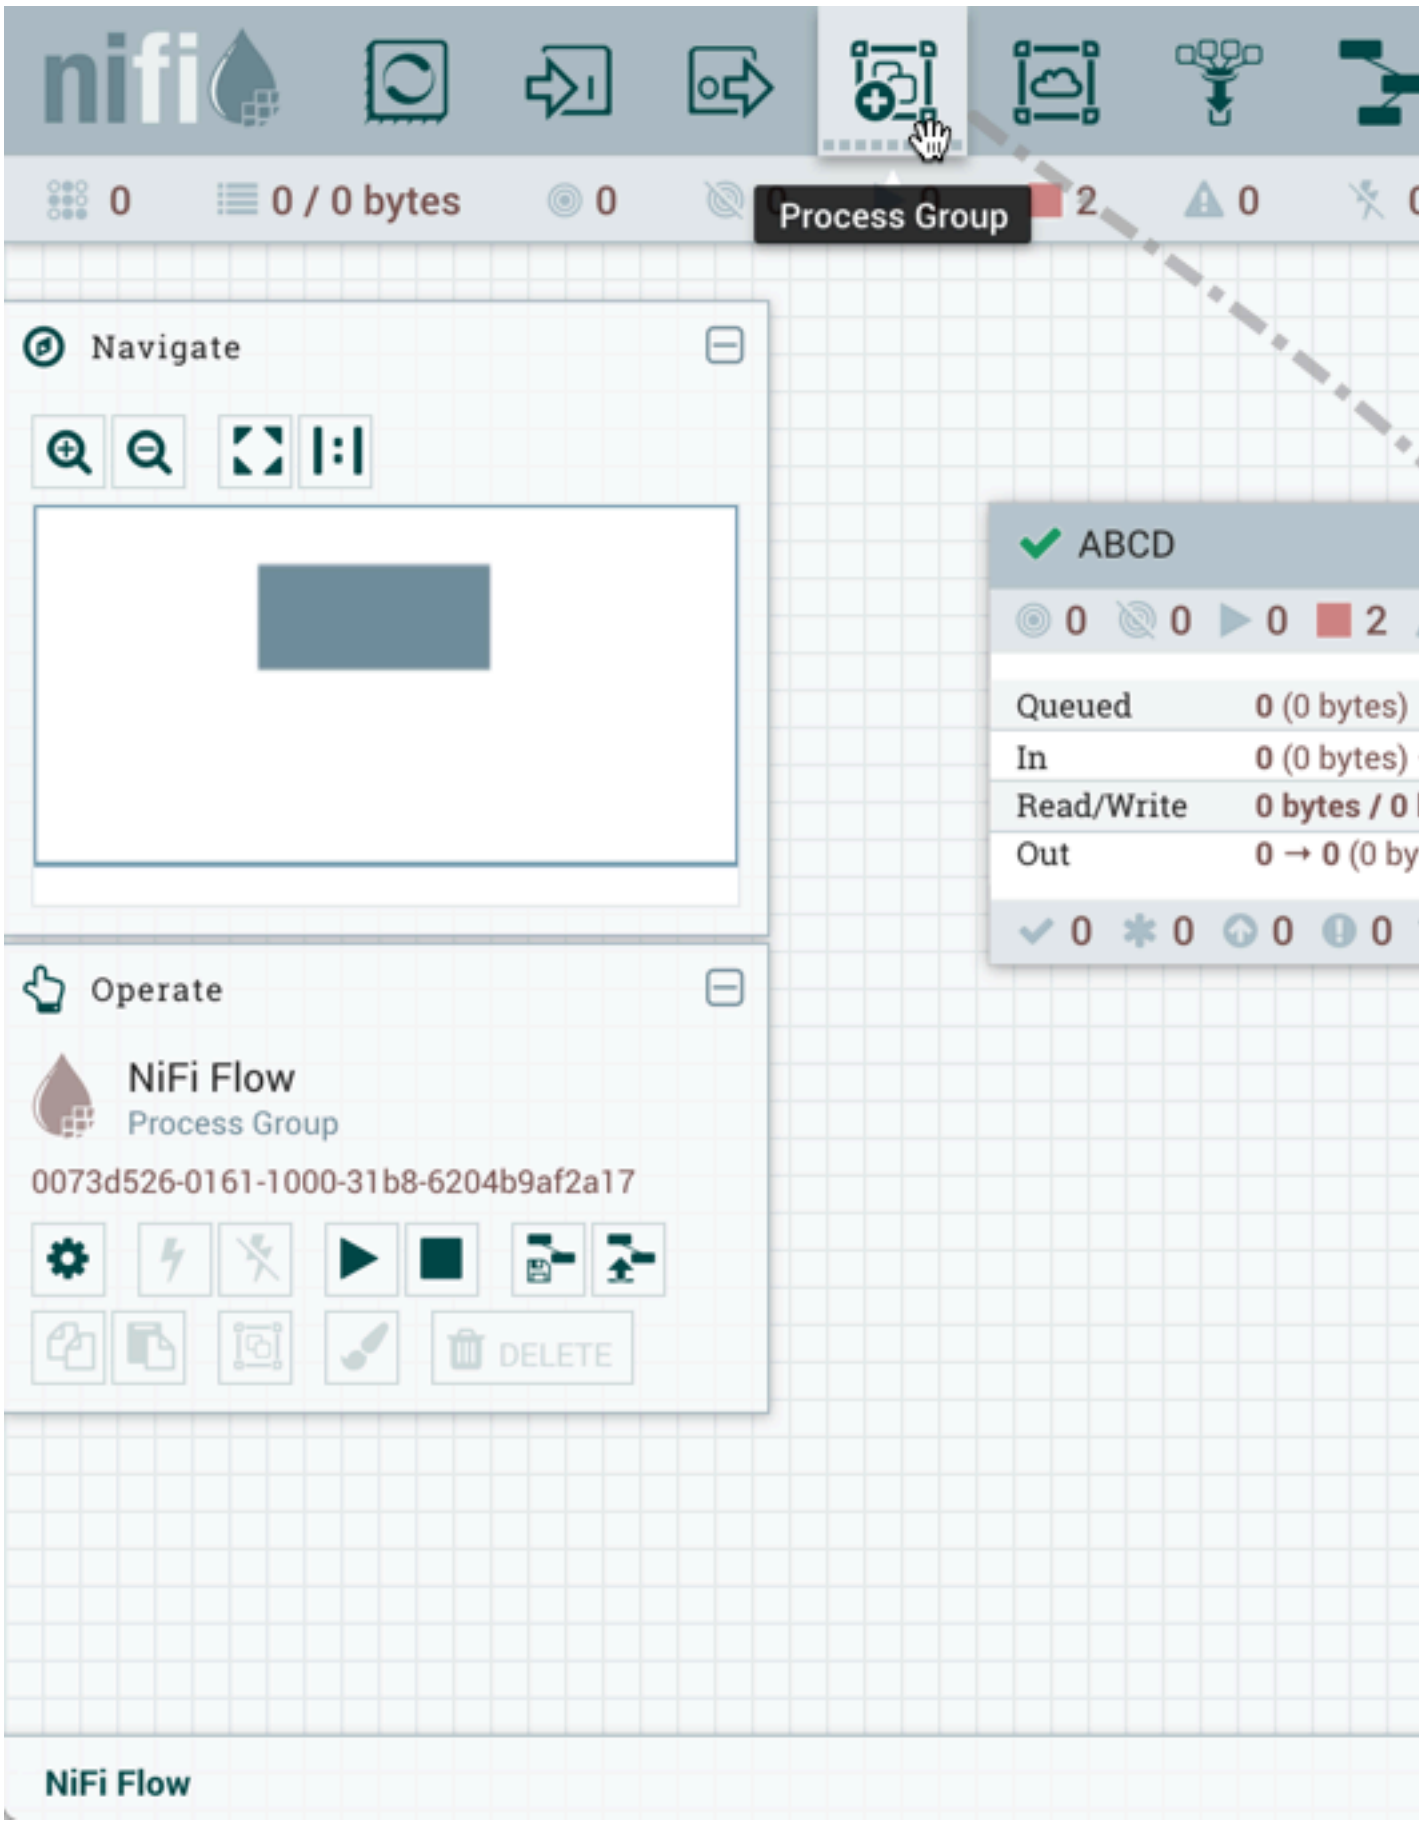
nifi
0
0 / 0 bytes
0
Process Group
2
0
Navigate
ABCD
0
0
0
2
Queued
0 (0 bytes)
In
0 (0 bytes)
Read/Write
0 bytes / 0
Out
0 → 0 (0 by
0
0
0
0
Operate
NiFi Flow
Process Group
0073d526-0161-1000-31b8-6204b9af2a17
DELETE
NiFi Flow

Instead of entering a name, click the Import link:

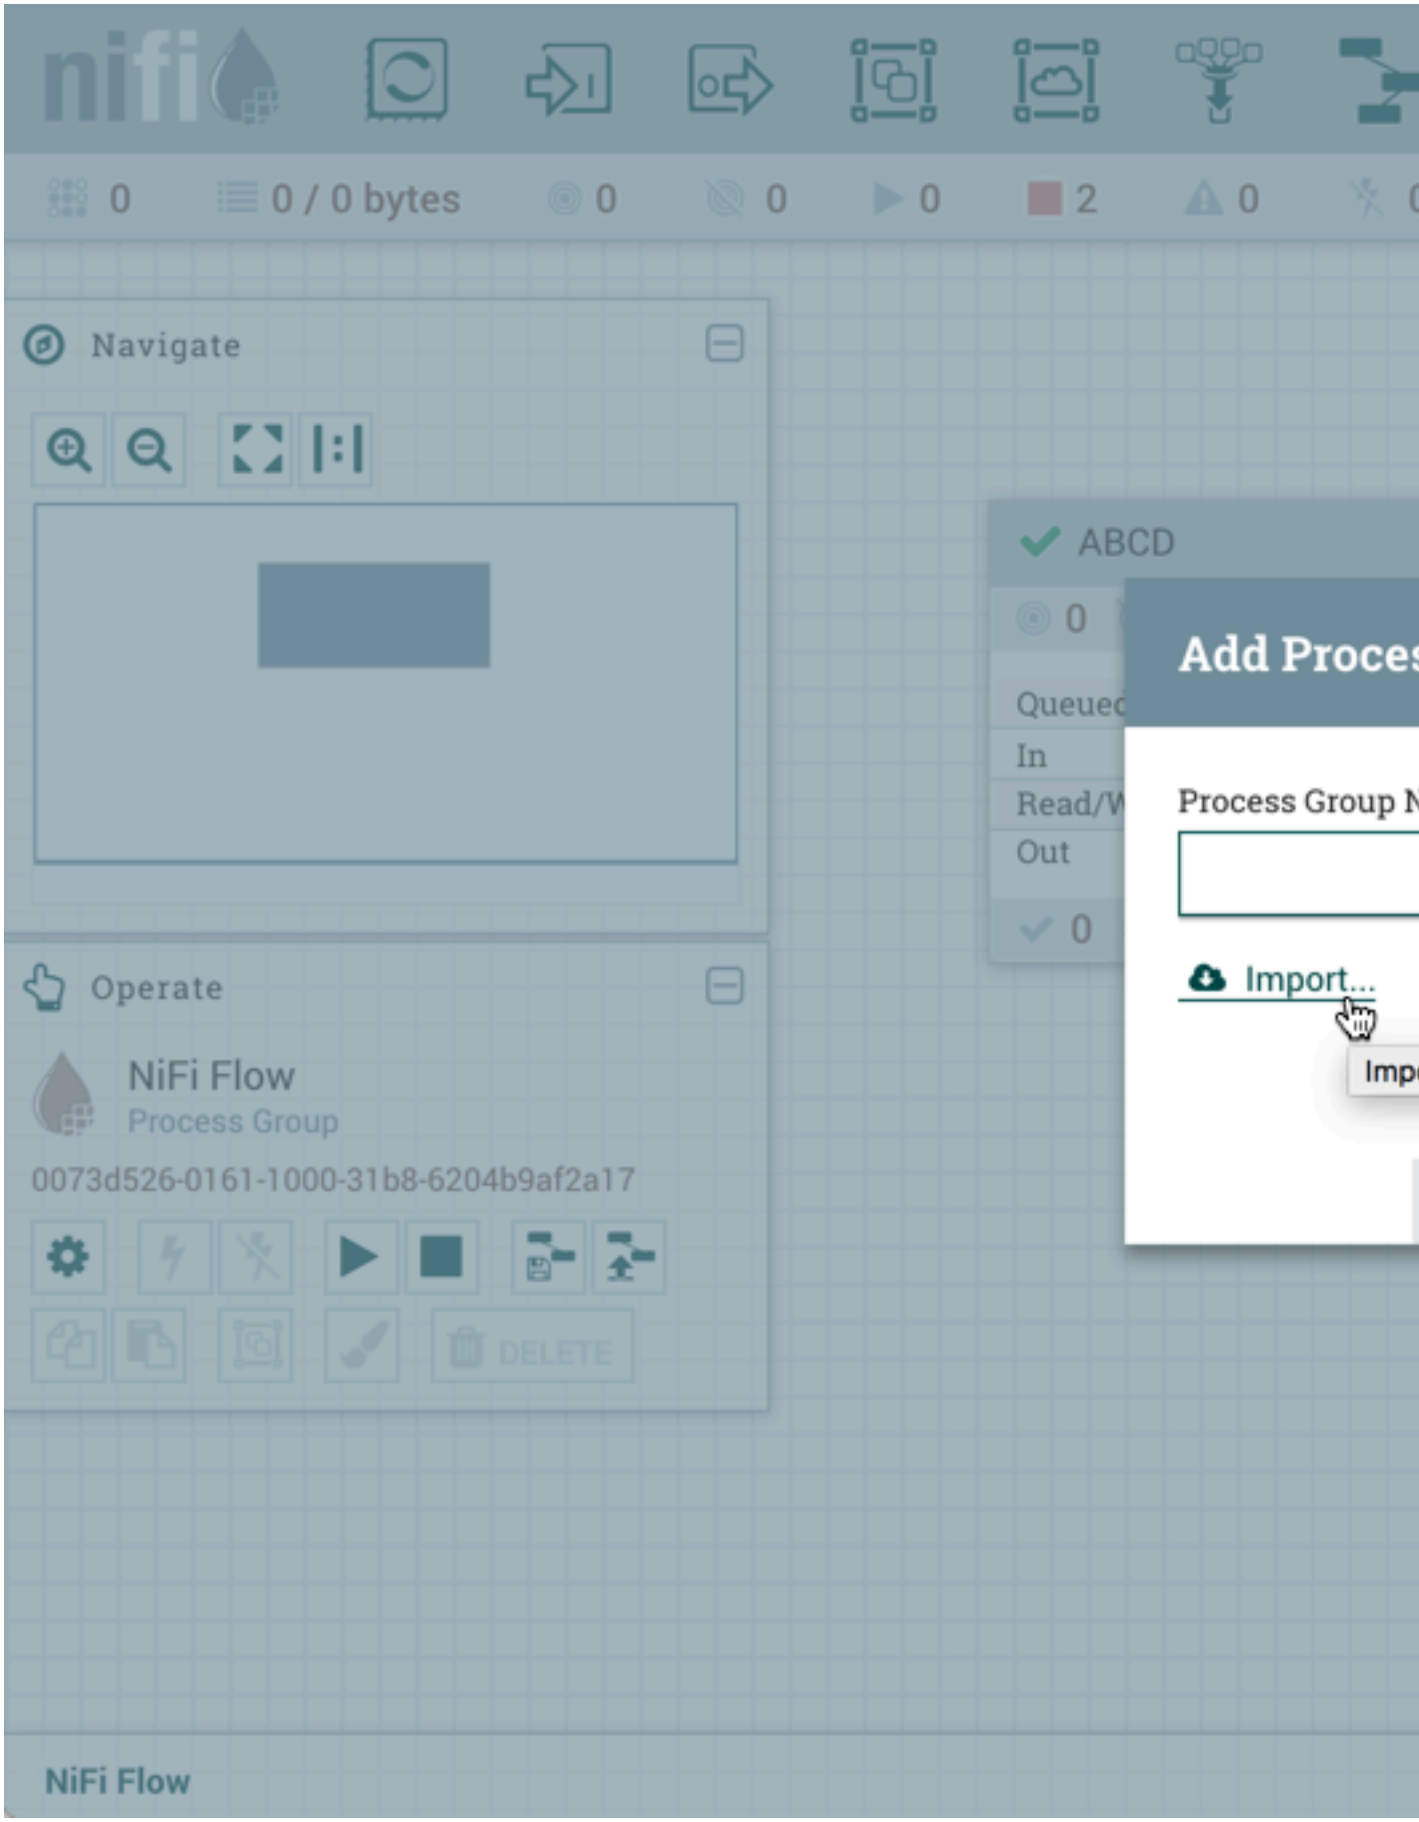
nifi
0
0 / 0 bytes
0
0
0
2
0
Navigate
Operate
NiFi Flow
Process Group
0073d526-0161-1000-31b8-6204b9af2a17
DELETE
ABCD
0
Queued
In
Read/W
Out
0
Add Proces
Process Group N
Import...
Impo
NiFi Flow

Choose the version of the flow you want imported and select "Import":

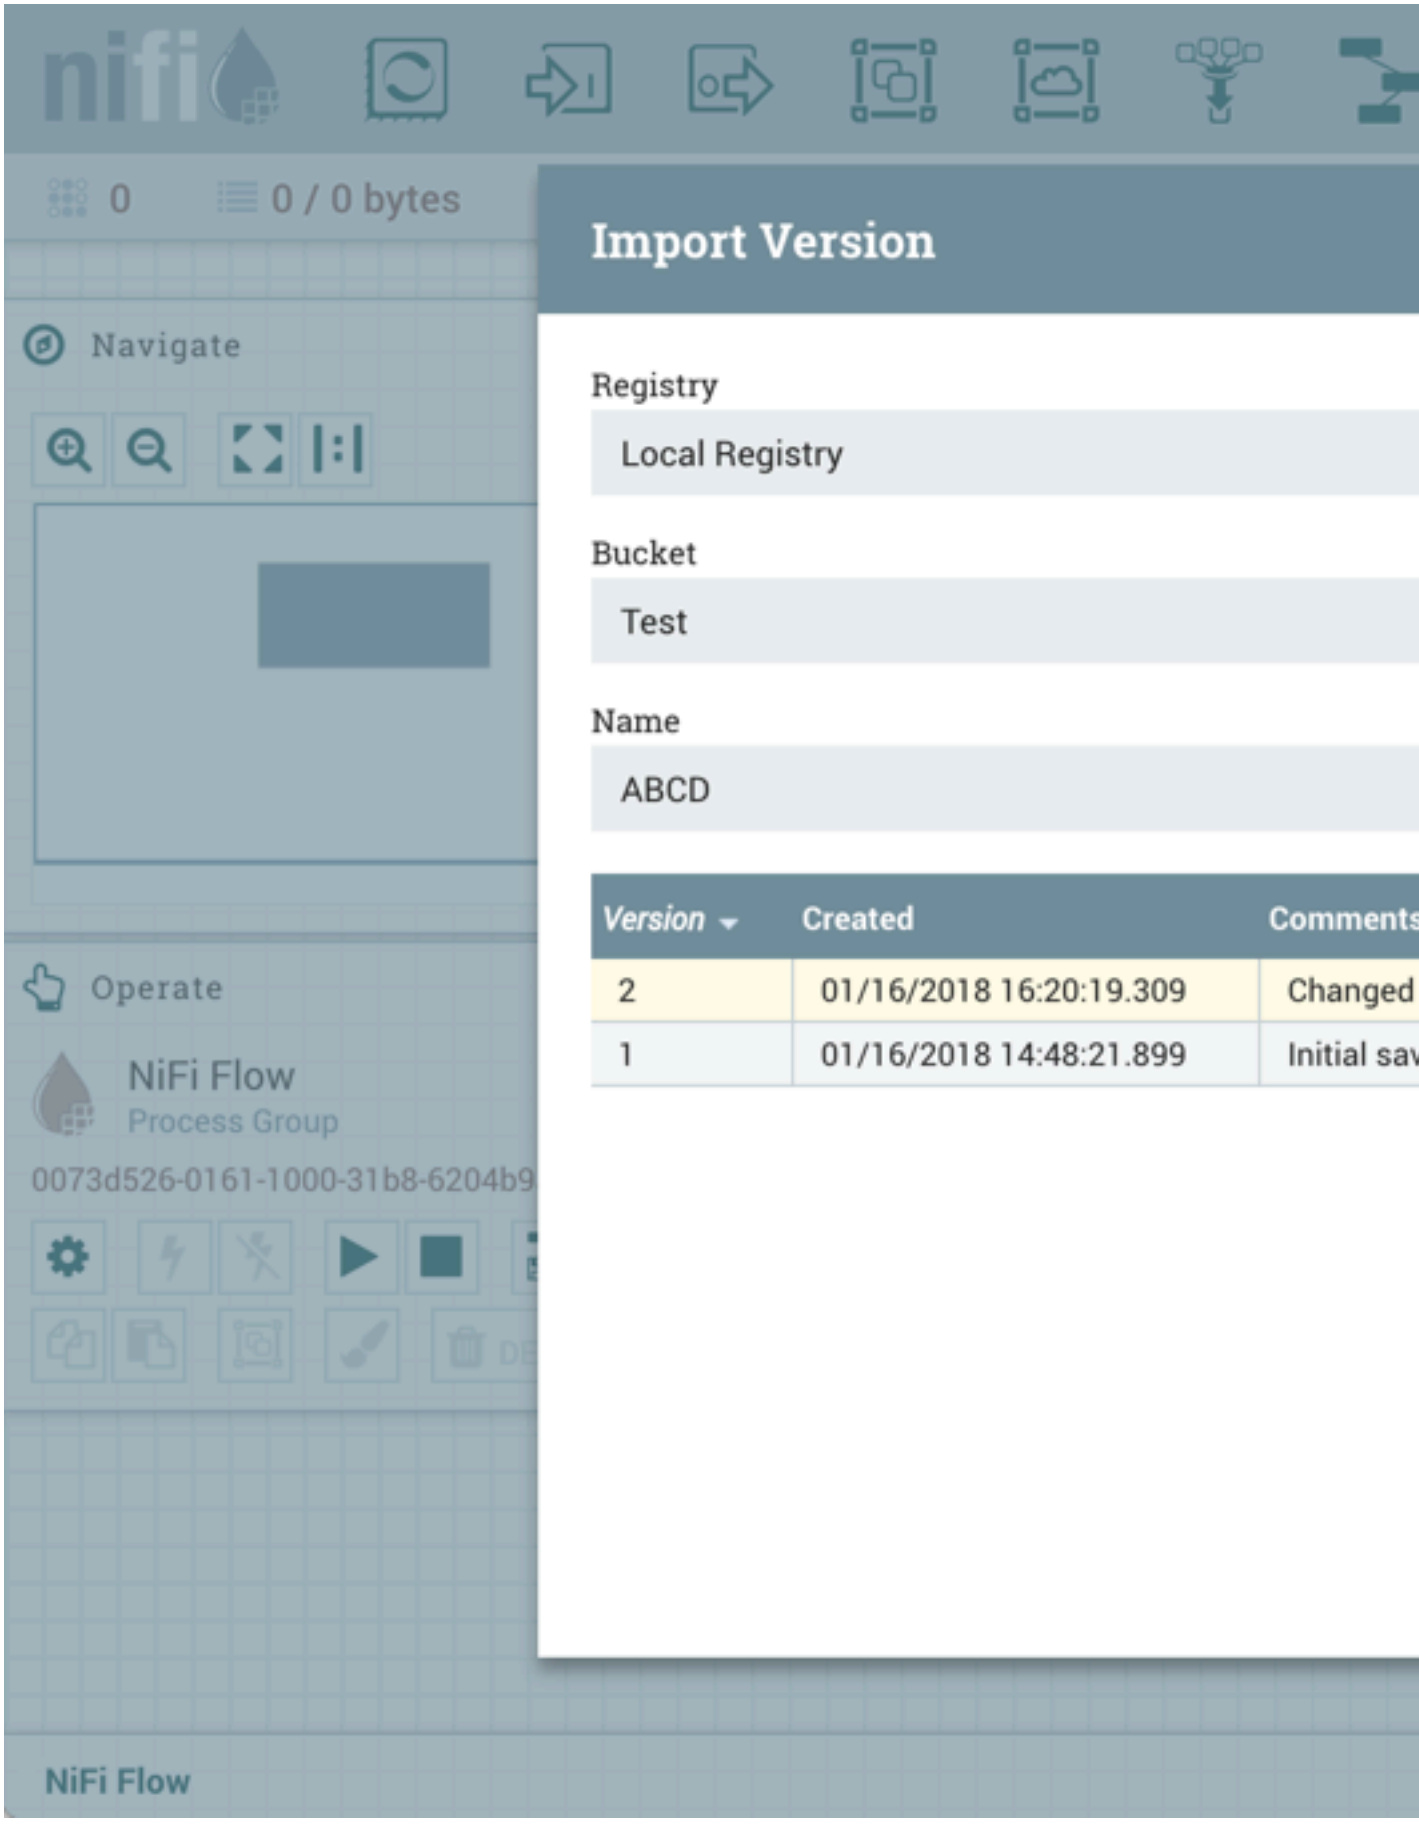
nifi
0
0 / 0 bytes
Navigate
Operate
NiFi Flow
Process Group
0073d526-0161-1000-31b8-6204b9
NiFi Flow
Import Version
Registry
Local Registry
Bucket
Test
Name
ABCD
Version
Created
Comments
2
01/16/2018 16:20:19.309
Changed
1
01/16/2018 14:48:21.899
Initial sav

A second identical PG is now added:

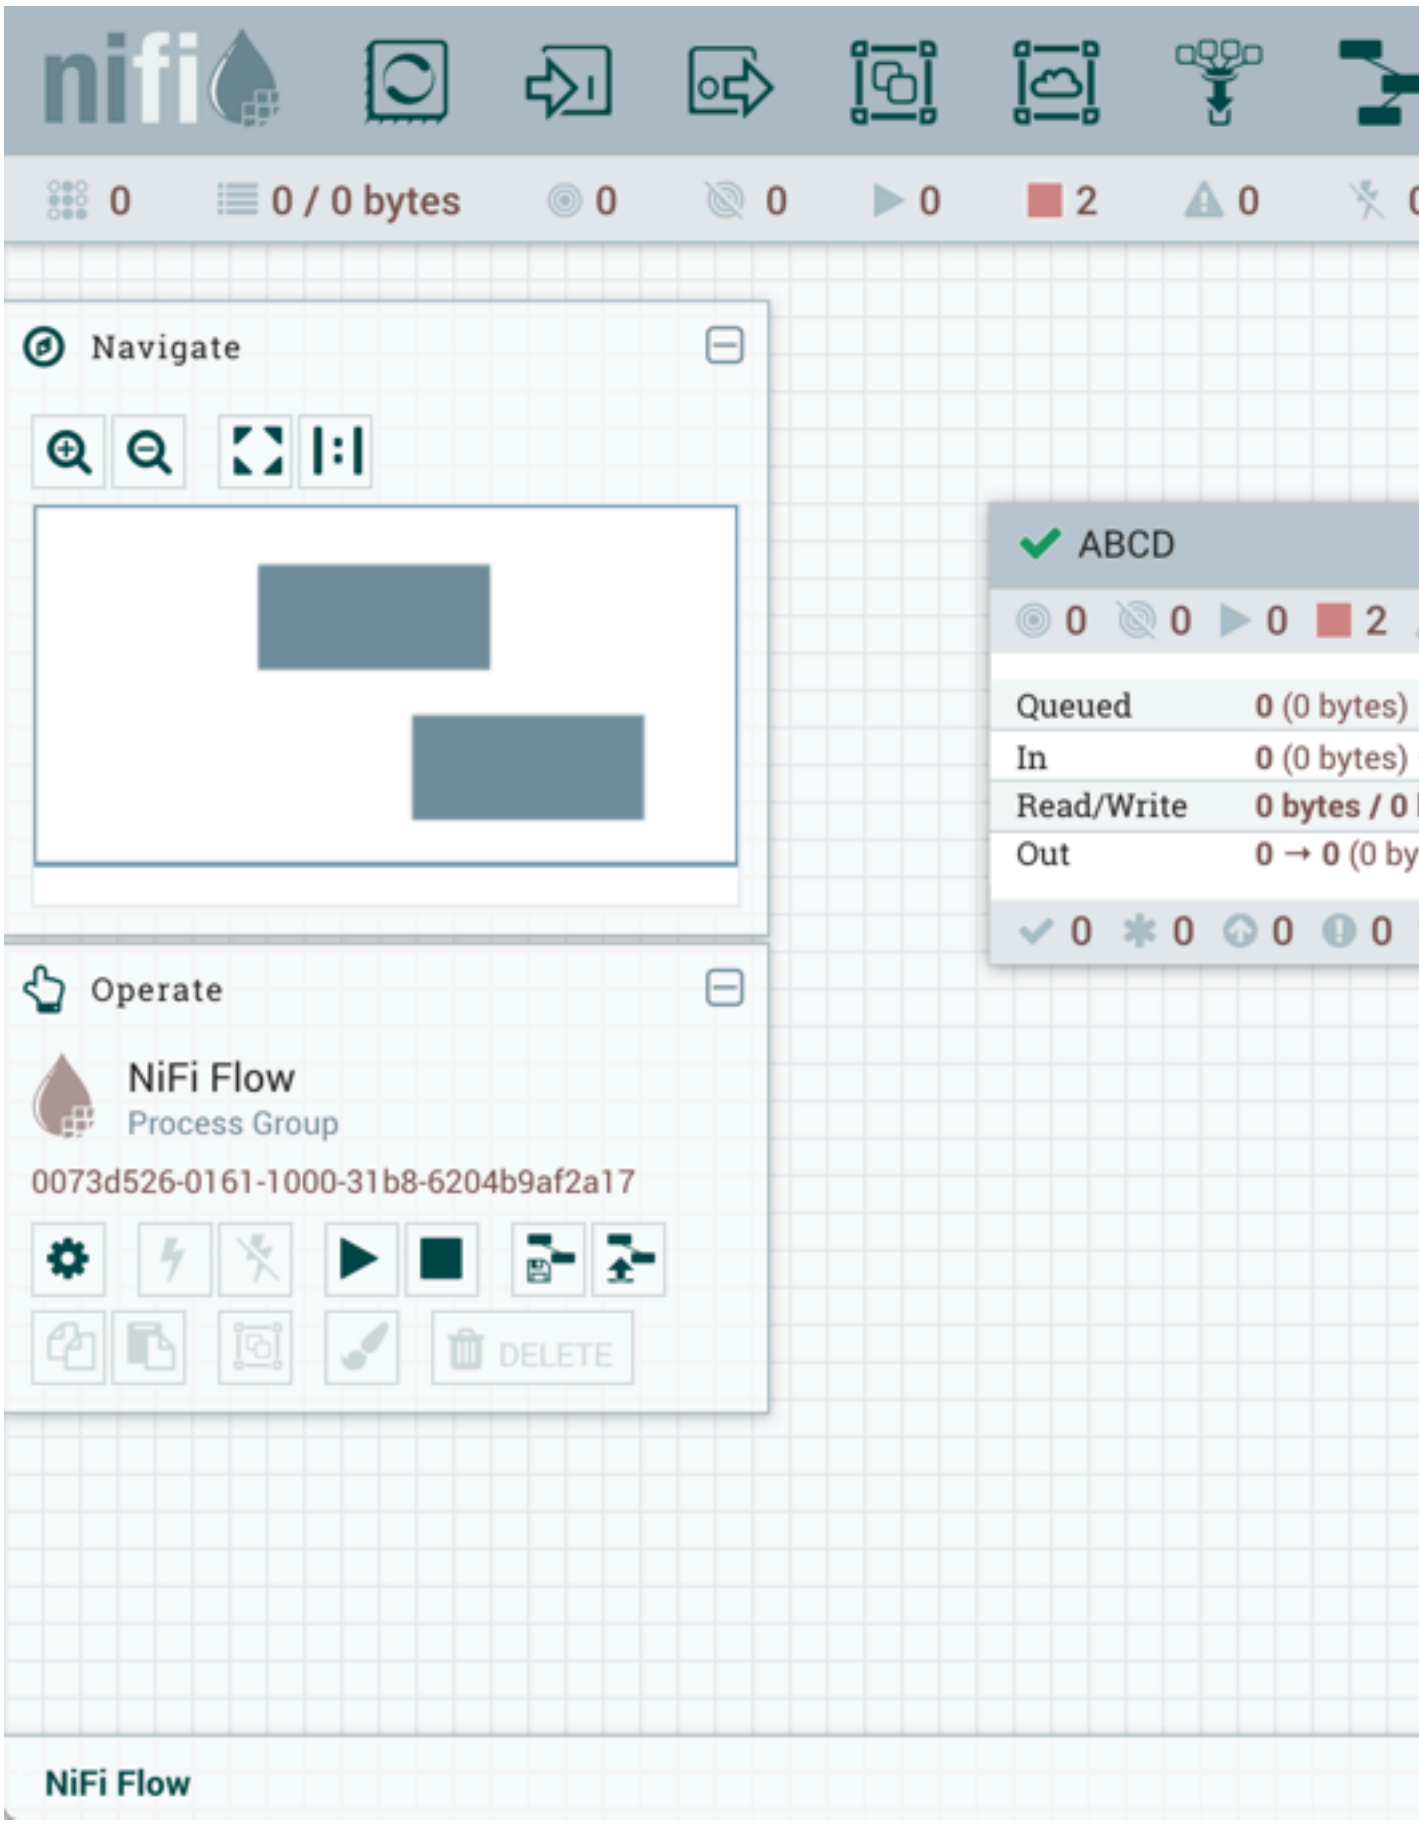
nifi
0
0 / 0 bytes
0
0
0
2
0
Navigate
Operate
NiFi Flow
Process Group
0073d526-0161-1000-31b8-6204b9af2a17
DELETE
ABCD
0
0
0
2
Queued
0 (0 bytes)
In
0 (0 bytes)
Read/Write
0 bytes / 0
Out
0 → 0 (0 by
0
0
0
0
NiFi Flow

Where To Go For More Information

In addition to this Getting Started Guide, more information about NiFi Registry and related features in NiFi can be found in the following guides:

- Apache NiFi Registry User Guide - This guide provides information on how to navigate the Registry UI and explains in detail how to manage flows/policies/special privileges and configure users/groups when the Registry is secured.
- Apache NiFi Registry System Administrator's Guide - A guide for setting up and administering Apache NiFi Registry. Topics covered include: system requirements, security configuration, user authentication, authorization, proxy configuration and details about the different system-level settings.
- Apache NiFi User Guide - A fairly extensive guide that is often used more as a Reference Guide, as it provides information on each of the different components available in NiFi and explains how to use the different features provided by the application. It includes the section Versioning a Dataflow which covers the integration of NiFi with NiFi Registry. Topics covered include: connecting to a registry, version states, importing a versioned flow and managing local changes.
- Contributor's Guide - A guide for explaining how to contribute work back to the Apache NiFi community so that others can make use of it.

In addition to the guides provided here, you can browse the different NiFi Mailing Lists or send an e-mail to one of the mailing lists at users@nifi.apache.org or dev@nifi.apache.org.

Many of the members of the NiFi community are also available on Twitter and actively monitor for tweets that mention @apachenifi.