

Creating an Inbound Connection Endpoint

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.

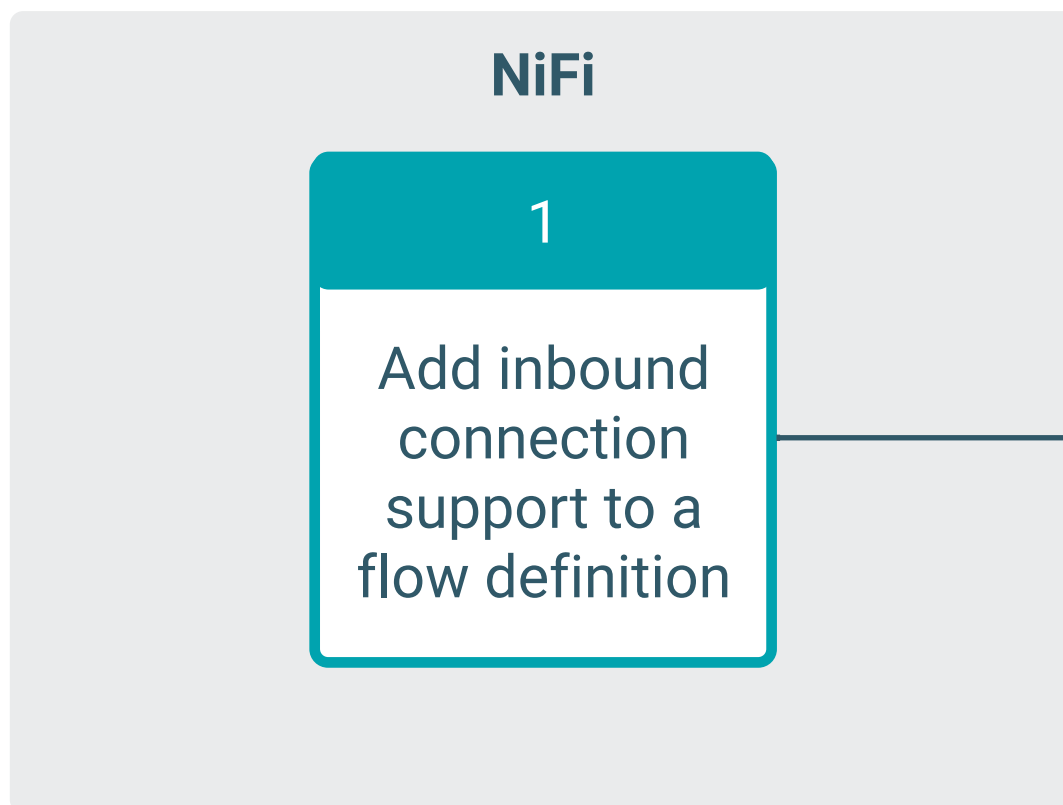
Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES 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

Create an Inbound Connection Endpoint during flow deployment.....	4
--	----------

Create an Inbound Connection Endpoint during flow deployment

During flow deployment, you can enable an endpoint that allows NiFi to listen to external data



sources.

About this task

You can create flow deployments with an Inbound Connection Endpoint during flow deployment, using the Cloudera DataFlow Deployment wizard.



Note: Currently you cannot create flow deployments with Inbound Connection Endpoints using the CLI.

Before you begin

- A flow definition supporting inbound connections has been created and added to the Catalog.
- You have DFCatalogViewer role to view flow definitions in the Catalog.
- You have DFFlowAdmin right for the environment where you want to create the endpoint.

Procedure

1. During flow deployment, at the NiFi Configuration step, select **Inbound Connections Allow NiFi to Receive Data**.

An Endpoint Hostname, generated from the Deployment Name you provided at the Overview step is offered and port configuration options become available.

dataflow-azure / New Flow Deployment

The screenshot shows the 'NiFi Configuration' step in a deployment wizard. On the left is a vertical navigation pane with steps: Overview (checked), NiFi Configuration (active), Parameters, Sizing & Scaling, Key Performance Indicators, and Review. The main content area is titled 'NiFi Configuration' and contains several sections:

- NiFi Runtime Version:** A dropdown menu showing 'CURRENT VERSION' and 'Latest Version (1.16.0.2.3.4.0-33)'. A 'Change Version' link is on the right.
- Autostart Behavior:** A checkbox labeled 'Automatically start flow upon successful deployment' is checked.
- Inbound Connections:** A checkbox labeled 'Allow NiFi to receive data' is checked. Below it, the 'Endpoint Hostname' is 'listenhttpflow'. A 'Select Preexisting Endpoint' link is on the right. A tooltip indicates 'This endpoint is valid'. The 'Listening Ports' section shows a table with one entry: TCP on port 7001. An 'Edit Ports' link is on the right.
- Custom NAR Configuration:** A checkbox labeled 'This flow deployment uses custom NARs' is unchecked.

On the right side of the wizard, an 'Overview' panel shows: FLOW DEFINITION: ListenHTTP filter to Kafka v.1; ENVIRONMENT DEPLOYING TO: dataflow-azure; DEPLOYMENT NAME: ListenHTTP Flow. At the bottom are 'Cancel', 'Previous', and 'Next' buttons.

2. You can accept the generated hostname as is, or you can change the prefix. You can also click **Select Preexisting Endpoint** to select an available endpoint within your environment.

Endpoints are unavailable if they are currently used by another flow deployment.

3. Select Listening Ports.

Provide the protocol and port or ports where your flow will listen to incoming data. Cloudera DataFlow automatically suggests ports based on listen processors detected in the flow definition. Custom Processors, if applicable, are not detected and must be added manually.

Using ports with mixed protocols (TCP and UDP) is not allowed. As long as the suggested ports use mixed protocols, you are not allowed to add new ports or to proceed to the next step in the Wizard.



Note: If you want your dataflow to use HTTP, select TCP protocol.

✕

Protocol

TCP
▼

Port ?

Specify a port
⊕ Add Port

Protocol	Port	
TCP ?	7001	⊖

Cancel

OK

4. Fill in Parameters.

If the listening port in your flow definition is parameterized instead of having a pre-defined value, you have to fill in the same listening port numbers you provided in the NiFi Configuration step. The listening ports you specified in the NiFi Configuration step are visible in the Overview pane.

5. Configure Sizing & Scaling, then click Next.

6. Add Key Performance Indicators, then click Next.

7. Review the information provided and click Deploy.

As part of the flow deployment, a Layer-4 Load Balancer (LB) is created to forward traffic to the flow deployment's NiFi cluster on the configured ports and protocols. DNS records are created for Inbound Connection Endpoint hostname to resolve to the Layer-4 LB.



Note: There is an initial delay for the cloud provider (AWS or Azure) to fully provision and start the LB resources, during which the flow deployment appears in the Cloudera DataFlow monitoring dashboard as healthy, but it is not yet able to receive data from clients. The load balancer provisioning normally completes within 10 minutes of the flow deployment being finished, but can take up to three hours.

Related Information

[Deploying a flow definition](#)