

Deploy a Flow

Date published: 2021-04-06

Date modified: 2024-01-09

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.

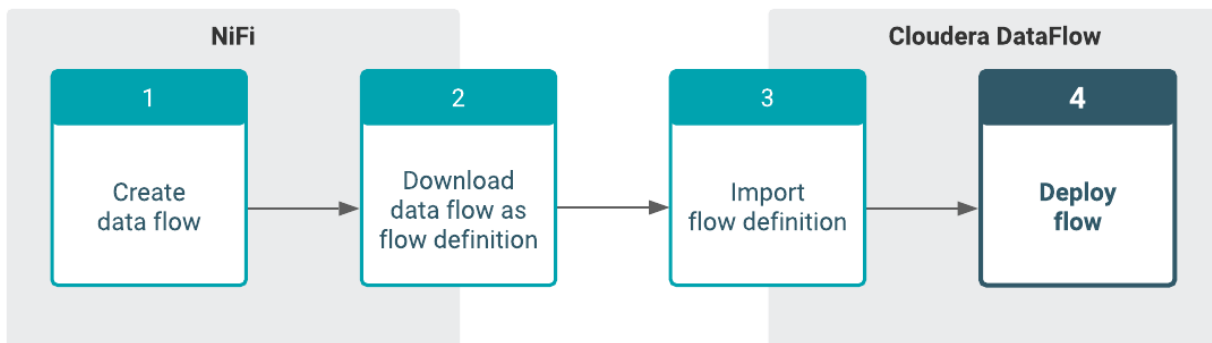
Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES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

Deploying a flow definition using the wizard.....	4
Deploying a flow definition using the CLI.....	7

Deploying a flow definition using the wizard

Deploy a flow definition to run NiFi flows as flow deployments in CDF. To do this, launch the Deployment wizard and specify your environment, parameters, sizing, and KPIs.



About this task

The CDF Catalog is where you manage the flow definition lifecycle, from initial import, to versioning, to deploying a flow definition.

Before you begin

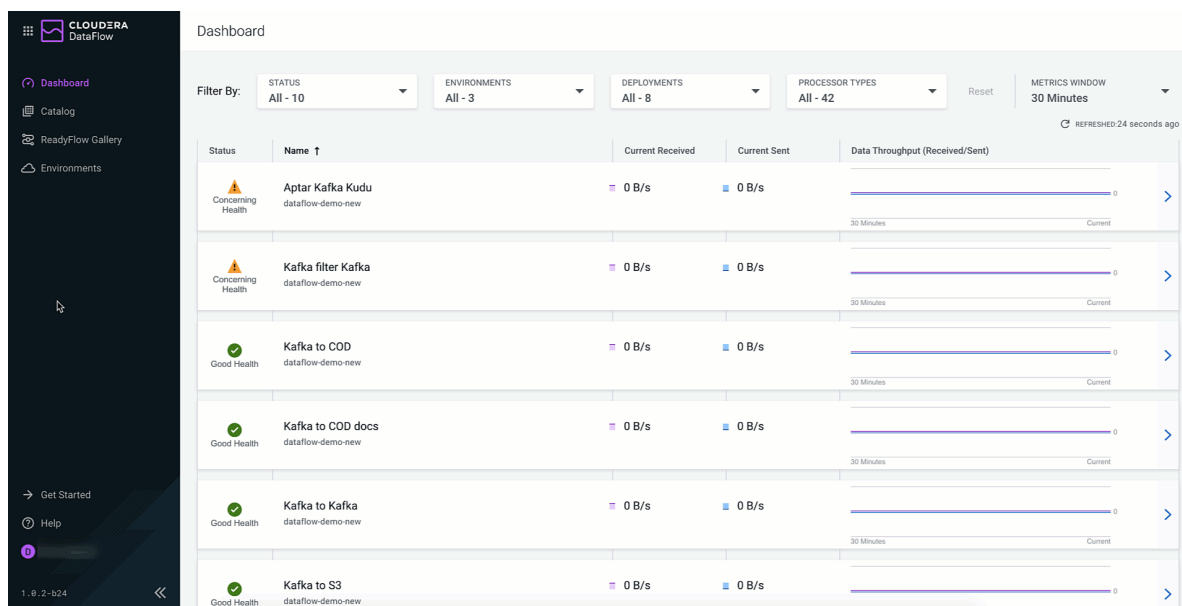
- You have an enabled and healthy CDF environment.
- You have imported at least one flow definition.
- You have been assigned the DFCatalogAdmin or DFCatalogViewer role granting you access to the CDF Catalog.
- You have been assigned the DFFlowAdmin role for the environment to which you want to deploy the flow definition.
- You have been assigned DFProjectMember role for the Project where you want to deploy the flow definition.
- If you are deploying custom Processors or Controller Services, you may need to meet additional prerequisites.

Procedure

1. In DataFlow, from the left navigation pane, select Catalog.
Flow definitions available for you to deploy are displayed, one definition per row.

2. Launch the Deployment wizard.

- Click the row to display the flow definition details and versions.
- Click a row representing a flow definition version to display flow definition version details and the Deploy button.
- Click Deploy to launch the Deployment wizard.



3. From the Deployment wizard, select the environment to which you want to deploy this version of your flow definition.



Note:

Only environments which have been enabled for DataFlow, are in a healthy state, and to which you have access display in the dropdown. Once you have selected the target environment, you cannot change it later in the flow deployment process without starting over.

4. From Overview, perform the following tasks:

- Give your flow a unique Deployment Name.

You can use this name to distinguish between different versions of a flow definition, flow definitions deployed to different environments, and similar.




Note:

Flow Deployment names need to be unique within an Environment. The Deployment wizard indicates whether a name is valid by displaying a green check below the Deployment Name text box.

- Select a Target Project for your flow deployment from the list of Projects available to you.

-

If you do not want to assign the deployment to any of the available Projects, select  Unassigned. Unassigned deployments are accessible to every user with DFFlowUser role for the environment.

- This field is automatically populated if you import a configuration and the Project referenced there exists in your environment and you have access to it.

- If you have previously exported a deployment configuration that closely aligns with the one you are about to deploy, you can import it under Import Configuration to auto-fill as much of the wizard as possible.

You can still modify auto-filled configuration values manually during deployment.

5. Specify NiFi Configuration information.

- a) Pick the NiFi Runtime Version for your flow deployment.

Cloudera recommends that you always use the latest available version, if possible.

- b) Specify whether you want to use Inbound Connections that allow your flow deployment receiving data from an external data source.

If yes, specify the endpoint host name and listening port(s) where your flow deployment listens to incoming data.

See *Creating an inbound connection endpoint* for complete information on endpoint configuration options.

- c) Specify whether you want to use NiFi Archives (NAR) to deploy custom NiFi processors or controller services.

If yes, specify the CDP Workload Username and password, and cloud storage location you used when preparing to deploy custom processors.



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.



Tip:

Make sure that you click the Apply button specific to Custom NAR Configuration before proceeding.

- d) Specify whether you want the flow deployment to auto-start once deployed.

6. In Parameters, specify parameter values like connection strings, usernames and similar, and upload files like truststores, jars, and similar.

7. Specify your Sizing & Scaling configurations.

NiFi node sizing

- You can adjust the size of your cluster between Extra Small, Small, Medium, and Large.

Number of NiFi nodes

- You can set whether you want to automatically scale your cluster according to flow deployment capacity requirements. When you enable auto-scaling, the minimum NiFi nodes are used for initial size and the workload scales up or down depending on resource demands.
- You can set the number of nodes from 1 to 32.
- You can set whether you want to enable Flow Metrics Scaling.

Storage Selection

You can select whether you want your deployment to use storage optimized for cost or performance.

8. From KPIs, you may choose to identify key performance indicators (KPIs), the metrics to track those KPIs, and when and how to receive alerts about the KPI metrics tracking.



Tip:

You can reorder the KPIs by dragging them up or down the list. The order you configure here will be reflected in the detailed monitoring view of the resulting deployment.

See *Working with KPIs* for complete information about the KPIs available to you and how to monitor them.

9. Review a summary of the information provided and make any necessary edits by clicking Previous. When you are finished, complete your flow deployment by clicking Deploy.



Tip:

Click View CLI Command to see the equivalent CDP CLI syntax in a help pane.

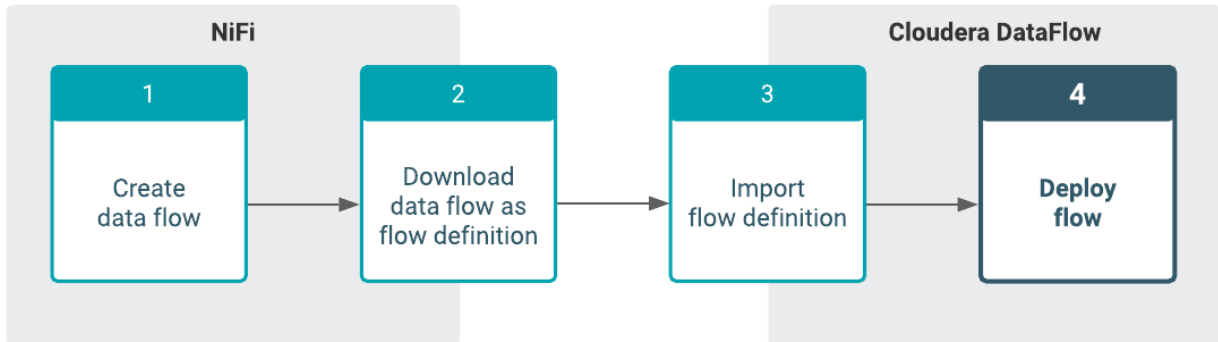
Results

Once you click Deploy, you are redirected to the Alerts tab for the deployment where you can track its progress.

Related Information[Working with KPIs](#)[Creating an inbound connection endpoint](#)[Auto-scaling](#)

Deploying a flow definition using the CLI

Deploy a flow definition to run NiFi flows as flow deployments in CDF. To do this, launch the CDP CLI and specify your environment, parameters, sizing, and KPIs.

**Before you begin**

- You have installed CDP CLI.
- Run `cdp df list-services` to get the service-crn value.
- Run `cdp df list-flows` to get the flow-version-crn value.

Procedure

To deploy a flow, enter:

```

cdp df create-deployment
  --service-crn [***service-crn-value***]
  --flow-version-crn [***flow-version-crn-value***]
  --deployment-name [***flow-deployment-name***]
  [--cluster-size-name [EXTRA_SMALL|SMALL|MEDIUM|LARGE]]
  [--static-node-count [***value***]]
  [--auto-scaling-enabled | --no-auto-scaling-enabled]
  [--auto-scale-min-nodes [***number-min-nodes***]]
  [--auto-scale-max-nodes [***number-max-nodes***]]
  [--cfm-nifi-version [***nifi-version***]]
  [--auto-start-flow | --no-auto-start-flow]
  [--parameter-groups [***json-file-location***]]
  [--kpis [***json-file-location***]]
  [---node-storage-profile-name [STANDARD_AWS|STANDARD_AZURE|PERFORMANCE_AWS|PERFORMANCE_AZURE]]
  
```

Where:

--service-crn

Specifies the service-crn value you obtained when completing the prerequisites.

--flow-version-crn

Specifies the flow-version-crn value you obtained when completing the prerequisites.

--deployment-name

Specifies a unique name for your flow deployment.

[--cluster-size-name]

Specifies the cluster size. Valid values are:

- EXTRA_SMALL
- SMALL
- MEDIUM
- LARGE

The default value is EXTRA_SMALL.

[--static-node-count]

Specifies the number of NiFi nodes when autoscaling is not enabled. You can select between 1 and 32 nodes. The default value is 1.

[--auto-scaling-enabled | --no-auto-scaling-enabled]

Specifies whether you want to enable autoscaling. The default is to disable autoscaling.

[--auto-scale-min-nodes]

Specifies the minimum nodes when you have autoscaling enabled. If you have autoscaling enabled, this parameter is required.

[--auto-scale-max-nodes]

Specifies the maximum nodes when autoscaling is enabled. If you have autoscaling enabled, this parameter is required.

[--cfm-nifi-version]

Specifies the NiFi runtime version. The default is the latest version.

[--auto-start-flow | --no-auto-start-flow]

Specifies whether you want to automatically start your flow once it has been deployed. The default is to enable the automatic start.

[--parameter-groups]

Specifies the location of the parameter group JSON file, if you are using one for this flow deployment.

[--kpis]

Specifies the location of the KPIs JSON file, if you are providing KPIs for this flow.

--node-storage-profile-name

Specifies the storage profile. Valid values are:

- STANDARD_AWS
- STANDARD_AZURE
- PERFORMANCE_AWS
- PERFORMANCE_AZURE

The default values are STANDARD_AWS and STANDARD_AZURE, depending on the cloud provider.

Example parameter group file:



Note: The JSON file you develop for parameter group will be different depending on your flow objectives and requirements. This is an example of the parameter group file format.

```
[
  {
    "name": "kafka-filter-to-kafka",
    "parameters": [
```



```

    {
      "name": "CDP Workload User",
      "assetReferences": [],
      "value": "srv_nifi-machine-ingest"
    },
    {
      "name": "CDP Workload User Password",
      "assetReferences": [],
      "value": "<<CDP_MISSING_SENSITIVE_VALUE>>"
    },
    {
      "name": "CSV Delimiter",
      "assetReferences": [],
      "value": ","
    },
    {
      "name": "Data Input Format",
      "assetReferences": [],
      "value": "CSV"
    },
    {
      "name": "Data Output Format",
      "assetReferences": [],
      "value": "JSON"
    },
    {
      "name": "Filter Rule",
      "assetReferences": [],
      "value": "SELECT * FROM FLOWFILE"
    },
    {
      "name": "Kafka Broker Endpoint",
      "assetReferences": [],
      "value": "streams-messaging-broker0.pm-sandb.a465-9q4k.cloudera.sit
e:9093"
    },
    {
      "name": "Kafka Consumer Group ID",
      "assetReferences": [],
      "value": "cdf"
    },
    {
      "name": "Kafka Destination Topic",
      "assetReferences": [],
      "value": "MachineDataJSON"
    },
    {
      "name": "Kafka Producer ID",
      "assetReferences": [],
      "value": "cdf"
    },
    {
      "name": "Kafka Source Topic",
      "assetReferences": [],
      "value": "MachineDataCSV"
    },
    {
      "name": "Schema Name",
      "assetReferences": [],
      "value": "SensorReading"
    },
    {
      "name": "Schema Registry Hostname",
      "assetReferences": [],

```

```

    "value": "streams-messaging-master0.pm-sandb.a465-9q4k.cloudera.site
  }
]
}
]

```

Example KPI file:



Note: The JSON file you develop for KPIs will be different depending on your flow objectives and requirements. This is an example of the KPI file format.

```

[
  {
    "metricId": "rateBytesReceived",
    "alert": {
      "thresholdLessThan": {
        "unitId": "kilobytesPerSecond",
        "value": 150
      },
      "frequencyTolerance": {
        "unit": {
          "id": "MINUTES"
        },
        "value": 1
      }
    }
  },
  {
    "metricId": "processorAmountBytesSent",
    "alert": {},
    "componentId": "a7f7df1c-a32d-3c25-9b09-e1d1036dcc04;a33a1b48-005b-32dd-bb88-b63230bb8525"
  }
]

```

Results

Successfully deploying a flow results in output similar to:

```

{
  "crn": "deployment-crn"
}

```