

Test Sessions

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

Test sessions.....	4
Starting a test session.....	4
Connecting an external application to a test session.....	5
Stopping a test session.....	6
Restarting a test session.....	7
Verifying flow file content.....	7
Queue management.....	12
Emptying a queue.....	12
Emptying all queues.....	13

Test sessions

Starting a Test Session provisions NiFi resources, acting like a development sandbox for a particular draft. It allows you to work with live data to validate your data flow logic while updating your draft. You can suspend a test session any time and change the configuration of the NiFi cluster, then resume testing with the updated configuration.

Starting a test session

To validate your draft flow, start a test session, provisioning a NiFi cluster.

Before you begin


- Make sure that you have DFDeveloper permission to perform this task. For information on account and resource roles, see Cloudera DataFlow Authorization.
- You have created a draft in Flow Designer.

Procedure

1. Select Flow Options Test Session .


- Add Custom NAR Configuration if your flow requires custom processors that you have built.
- Select Inbound Connections Allow NiFi to receive data if your draft uses listen processors with inbound connection endpoint support.

The deployment wizard auto synchronizes port information with the configuration settings you made for your listen processor(s).



Note: You can also change NiFi version but, unless you have a specific reason to select an earlier one, Cloudera recommends using the latest version.

2. Click Start Test Session.

Test Session status changes to  Initializing Test Session...

3.

Wait for the status to change to  Active Test Session.

4. Click Flow Options Services to enable Services.

5. Select a service you want to enable, then click Enable.

Repeat this step for all Services.



Tip: You can also select Enable Service and Referencing Components. It will not only enable the controller service, but also any component that references it. This way, you do not need to enable the component separately to run the test session.

6. Click Back To Flow Designer to return to the flow design Canvas.

7.

Start all processors that are not running (you see a  Stopped symbol next to the processor name) by selecting the processor on the Canvas then clicking  Start.

Results

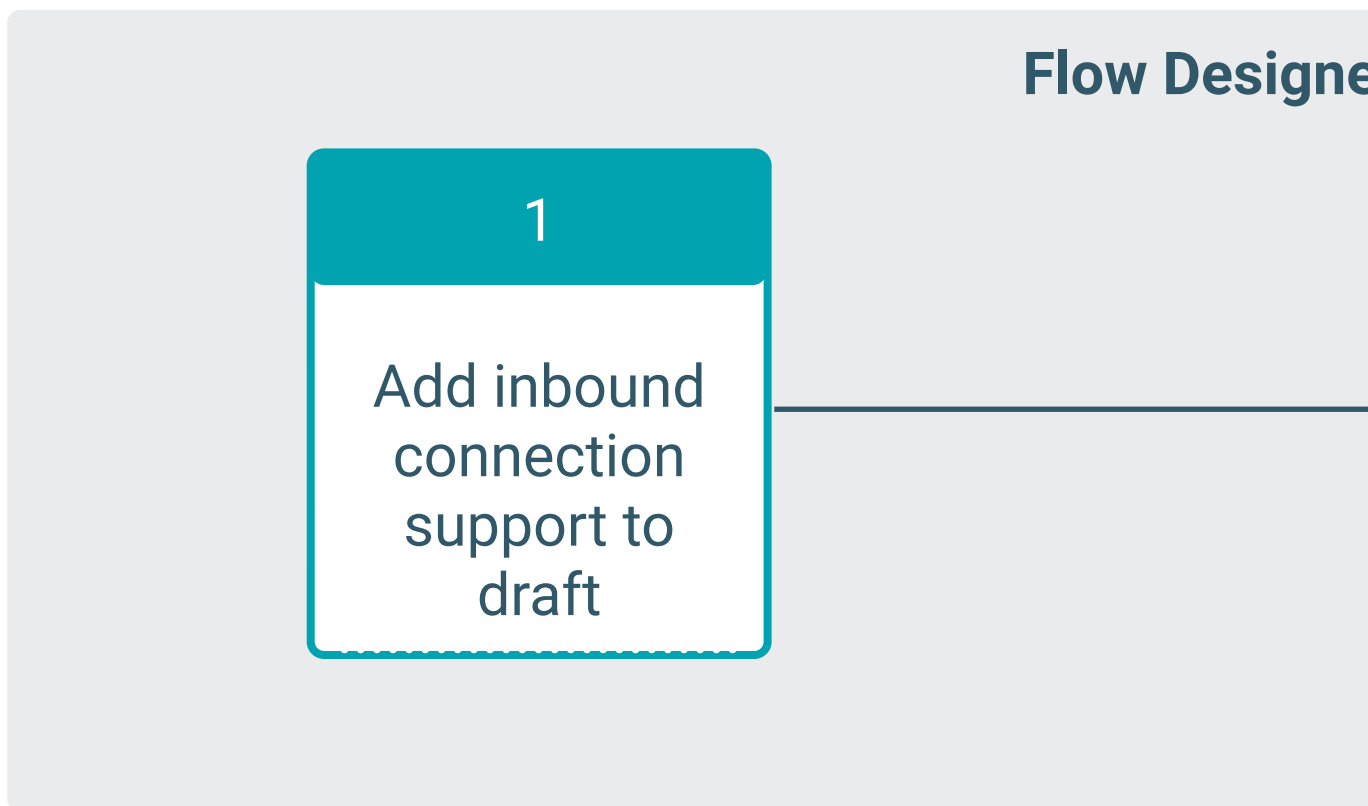
The flow starts executing. On the Flow Design Canvas you can observe statistics on your processors change as they consume data and execute their respective tasks. You can also observe one or more blue Notification Pill, providing information about the current task. Potential issues that need to be addressed are called out in an orange



Notification Pill on the individual processors.

Connecting an external application to a test session

Once a test session with inbound connection is running, you can go on and connect an external application to start sending data.



Before you begin

- A test session with inbound connection configured is running.
- A network connection through which the client can reach the NiFi endpoint is available.
- You have been assigned at least the DFFlowDeveloper role for the environment to which you want to configure the inbound connection.

Procedure

1. Select the draft where you want to send data and go to Flow Options Inbound Connection Details .
2. Make a note of the ENDPOINT HOSTNAME and LISTENING PORTS.
3. Click Download Client Certificate.

The X.509 client certificate downloads to your computer in PEM format.

4. Click Download Private Key to obtain the RSA Private Key.

The unencrypted RSA Private Key encoded with PKCS8 downloads to your computer in PEM format.

5. Depending on your client, you may have to convert the certificate and the private key to a different format.

For example, to convert PEM to PKCS12 format, use the following command:

```
openssl pkcs12 -export -in [***DOWNLOADED PEM CERT FILE***] -in
key [***DOWNLOADED PEM PRIVATE KEY***] -out certificate.p12
```

To further convert the PKCS12 file to JKS format for a Java client, run the following command:

```
keytool -importkeystore -srckeystore [***CERTIFICATE NAME***].p12 -srcsto
retype pkcs12 -destkeystore [***DESTINATION KEYSTORE***].jks
```

6. Add the certificate file and the private key files to the keystore of your application.
7. Configure your application to stream data to the Endpoint Hostname, port, and protocol of the NiFi instance behind the Test Session.


Stopping a test session


You can stop a test session to reconfigure your NiFi cluster. You can continue editing your draft flow in the meantime.

Before you begin


- Make sure that you have DFDeveloper permission to perform this task. For information on account and resource roles, see Cloudera DataFlow Authorization.
- You have an active test session.

Procedure

1. You have two options to end an active test session:
 - Click the toggle next to the Active Test Session notification.
 - Click Flow Options then click  Stop under Test Session.
2. Click End.

Test session status first changes to  Ending Test Session... then, after the NiFi cluster shuts down, it changes to  Test Session Not Active.

What to do next

Once the test session stops, you can update the configuration of your NiFi testbed cluster. To do so, select Flow Options Test Session  Edit Configuration .



Restarting a test session

Learn about restarting an inactive test session.

Before you begin

- Make sure that you have DFDeveloper permission to perform this task. For information on account and resource roles, see Cloudera DataFlow Authorization.
- You have an inactive test session.

Procedure

1. You can optionally update the configuration of your NiFi development sandbox before restarting it. Click Flow Options then click Edit Configuration under Test Session.
2. You have two options to restart an inactive test session:
 - Click the toggle next to the **Test Session Not Active**  Test Session Not Active notification.
 - Click Flow Options then click  Restart under Test Session.

Verifying flow file content

Before you begin

- Make sure that you have DFDeveloper permission to perform this task. For information on account and resource roles, see Cloudera DataFlow Authorization.
- You have an active test session.

Procedure

- 1. On the Flow Design Canvas, click the connection where you want to inspect the flow files.

Flow Design / cdf-priv / Hello World CSV

Active Test Session

Flow Options

READ/WRITE0.00 B / 39.01 KB

OUT10 (39.01 KB)

TASKS10 / 00:00:01.063

5 min

NAMEResponse

QUEUED200 (854.31 KB)

Convert JSON to CSV

ConvertRecord 1.20.0.2.3.8.0-31

IN0 (0.00 B)

READ/WRITE0.00 B / 0.00 B

OUT0 (0.00 B)

TASKS0 / 00:00:00.000

5 min

NAMEConverted changes

QUEUED0 (0.00 B)

SOURCEGet Recent Wikipedia Ch...→DESTINATIONConvert JSON to CSV

More Details

Settings

Connection Name

Connection Name

Flowfile Expiration

0 sec

Back Pressure Object Threshold

10000

Back Pressure Size Threshold

1 GB

Load Balance Strategy

Do not load balance

Relationships

☐Original

☐Failure

☐Retry


☐No Retry

☒Response

Apply

Discard Changes

2.

In the Configuration pane select  List Queue .



















In the List Queue view, All Queued Flow Files for [***RELATIONSHIP NAME***] Relationship opens.

Flow Design / cdf-priv / Hello World CSV / List Queue Active Test Session Flow Options

[← Back To Flow Designer](#)

All Queued Flow Files for Response Relationship(s)

Displaying 100 (917 KB) REFRESHED: 2023-03-06 13:30 CET

Position	Filename	UUID	File Size	Queue Duration	Lineage Duration	Penalized	
1	19e761c1-6a...	 19e761c...	4 KB	01:46:50.561	01:46:50.561	No	>
2	88d8d94c-39...	 88d8d94...	4 KB	01:46:20.452	01:46:20.452	No	>
3	9b352478-46...	 9b35247...	4 KB	01:46:12.580	01:46:12.581	No	>
4	ec4e9adb-75...	 ec4e9adb...	6 KB	01:45:42.470	01:45:42.470	No	>
5	2df3f2c2-e55...	 2df3f2c2...	4 KB	01:45:12.353	01:45:12.353	No	>
6	16e849f1-508...	 16e849f1...	4 KB	01:44:42.248	01:44:42.248	No	>
7	dad34e7b-56...	 dad34e7...	4 KB	01:44:12.134	01:44:12.134	No	>
8	d6e80363-71f...	 d6e8036...	4 KB	01:43:42.031	01:43:42.031	No	>
9	660c2318-f2c...	 660c231...	4 KB	01:43:11.918	01:43:11.919	No	>
10	db57f0ee-2e0...	 db57f0ee...	4 KB	01:42:41.806	01:42:41.806	No	>
11	bb950d3b-26...	 bb950d3...	5 KB	01:42:11.696	01:42:11.697	No	>
12	bd255467-61...	 bd25546...	4 KB	01:41:41.580	01:41:41.580	No	>
13	2db9a267-a9...	 2db9a26...	4 KB	01:41:11.474	01:41:11.475	No	>
14	3427f0bc-19b...	 3427f0bc...	5 KB	01:40:41.349	01:40:41.349	No	>
15	df14c3e5-8a4...	 df14c3e5...	5 KB	01:40:11.244	01:40:11.244	No	>
16	b51123ae-05...	 b51123a...	5 KB	01:39:41.135	01:39:41.136	No	>
17	1c0a81fb-598...	 1c0a81fb...	4 KB	01:39:11.028	01:39:11.028	No	>
18	bfd3d9e1-dc8...	 bfd3d9e1...	4 KB	01:38:40.918	01:38:40.918	No	>

Attributes (29)

Pinned (0) Copy

date
Mon, 06 Mar 2023 10:43:10 GMT

server
mw2284.cdfw.wmnet

invokehttp.tx.id
65ebe530-3c4b-4b28-99c7-aed19f70688b

vary
Accept-Encoding,Treat-as-Untrusted,X-Forwarded-Proto,Cookie,Authorization

x-frame-options
SAMEORIGIN

x-client-ip
34.213.89.169

invokehttp.remote.dn
CN=*.wikipedia.org

invokehttp.request.duration
226


uuid
19e761c1-6a14-4b25-a30b-77c030338541

invokehttp.request.url
https://en.wikipedia.org/w/api.php?action=query&list=recentchanges&format=jso...

path
/

3. Select a flow file by clicking it to view its Attributes in the right-hand pane.

4.

If there is an attribute that is of special interest to you, hover the cursor over it and select  Pin this attribute. The selected attribute is pinned to the top of the list of attributes. Pinning persists, if you select a different flow file, or return to the Canvas and decide to view flow files later, you will find the pinned attribute(s) at the top of the list.



19e761c1-6a14-4b25-a30b-77c030338541



19e761c1-6a14-4b25-a30b-77c030338541

[More Details](#)



Attributes (29)

Pinned (1)

Copy

invokehttp.status.code
200



date
Mon, 06 Mar 2023 10:43:10 GMT

server
mw2284.codfw.wmnet

invokehttp.tx.id
65ebe530-3c4b-4b28-99c7-aed19f70688b

vary
Accept-Encoding,Treat-as-Untrusted,X-Forwarded-Proto,Cookie,Authorization

x-frame-options
SAMEORIGIN


x-client-ip
34.213.89.169

invokehttp.remote.dn
CN=*.wikipedia.org

invokehttp.request.duration
226


uuid
19e761c1-6a14-4b25-a30b-77c030338541

invokehttp.request.url
<https://en.wikipedia.org/w/api.php?action=query&list=recentchanges&format=json>...


5. To view flow file content, click  Open in Data Viewer on top of the Attributes pane.



Tip:

- Click More details  to view file metadata.
- To format the file printout, select an available grammar from the View as drop-down.

The contents of the flow file is displayed in a new tab in your browser.



The screenshot shows the Cloudera Data Viewer interface. On the left, the 'Content' pane displays a JSON flow file. The 'View as' dropdown is set to 'Original'. On the right, the 'Attributes (29)' pane shows file metadata:

Attribute	Value
FILE SIZE	4 KB
MIME TYPE	application/json; charset=utf-8
LINEAGE DURATION	02:37:37.933
QUEUE POSITION	No value set
PENALIZED	No
QUEUE DURATION	02:37:37.932

Below the attributes, a list of flow files is shown, including 'invokehttp.status.code' with a value of 200.

Queue management

You can remove all flow files from a selected connection or all flow files on process group level.

Before you begin

- Make sure that you have DFDeveloper permission to perform this task. For information on account and resource roles, see Cloudera DataFlow Authorization.
- You have an active test session.



Procedure

- To view the queue of a connection, right click on the connection and select List Queue. The top 100 flow files are listed on the **List Queue** page.

Emptying a queue

Learn how to delete enqueued flow files from a connection.

Procedure



- Right click on the connection where you want to delete the enqueued flow files.
- Select  Empty Queue.
- Click  Empty.

All FlowFiles waiting on the selected connection at the time of the request are removed.

Emptying all queues

Learn how to delete enqueued flow files from all connections on Process Group level.

Procedure

1. Click on an empty area of the Canvas.
The **Configuration** pane on the right displays configuration options for the process group you are currently in.
2. Click  Empty All Queues.
3. Click  Empty.

All FlowFiles waiting on all connections within the process group at the time of the request are removed.