

Cloudera Flow Management 4.0.0

Cloudera Flow Management Migration Tool

Date published: 2019-06-26

Date modified: 2025-02-27

CLOUDERA

<https://docs.cloudera.com/>

Legal Notice

© Cloudera Inc. 2025. 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

Cloudera Flow Management Migration Tool Release Notes.....	4
What's new.....	4
Support Matrix.....	4
Download location.....	4
Cloudera Flow Management Migration Tool overview.....	5
Key features.....	5
Migration stages.....	6
Recommended migration workflow.....	7
Migration best practices.....	8
Migrating a data flow using the Cloudera Flow Management Migration	
 Tool.....	10
Preparing for migration.....	10
Setting up your environment.....	10
Configuring the Migration Tool.....	12
Migrating a flow.....	14
Migrating templates.....	14
Migrating variables.....	19
Migrating components.....	25
Using migrate-all.....	34
Cloudera Flow Management Migration Tool Reference.....	36
Commands.....	37
migrate-templates.....	37
migrate-variables.....	38
migrate-components.....	39
migrate-all.....	41
help.....	41
Activity Log.....	41
Troubleshooting flow migration issues.....	44

Cloudera Flow Management Migration Tool Release Notes

Learn about the new features, known and fixed issues, limitations, and unsupported features in this Technical Preview release of the Cloudera Flow Management Migration Tool.

What's new

The Cloudera Flow Management Migration Tool helps you transition to NiFi 2 more efficiently. It simplifies and accelerates the migration of flows from Cloudera Flow Management 2.x powered by NiFi 1 to Cloudera Flow Management 4.x powered by NiFi 2. The tool automates complex and repetitive tasks in updating flow configurations, reducing manual effort and ensuring compatibility with NiFi 2 features.

You can use this command-line tool to transform templates, variables, and components to align with NiFi 2 features.

This initial release supports on-premises environments and facilitates migrations from Cloudera Flow Management 2.1.7.1000 to Cloudera Flow Management 4.0.0, in compliance with the NiFi 1.26-to-2.0 General Availability (GA) ruleset.

Support Matrix

Before using the Cloudera Flow Management Migration Tool, review the following compatibility and system requirements to ensure your system aligns with these specifications and to avoid migration issues.

Cloudera Flow Management Migration Tool 1.0 supports on-premises usage only. A given Migration Tool version always supports exactly one pair of Cloudera Flow Management versions. Using different versions than those specified may lead to incorrect migration results or failure of the process.

Requirement		Details
Operating system		Linux
Java version		Java 21
Source version	Cloudera Flow Management	2.1.7.1000
	Apache NiFi	1.26.0
Target version	Cloudera Flow Management	4.0.0
	Apache NiFi	2.0.0
Disk space		At least 25 GB for NiFi distributions and working directories



Important:

Ensure that you use the specified Cloudera Flow Management Migration Tool versions and meet the other system requirements as well for a successful migration.

Download location

The Cloudera Flow Management Migration Tool software artifacts can be downloaded from the Cloudera Archive.

To access the artifacts, you must have:

- An active Cloudera subscription agreement
- A license key file
- The required authentication credentials (username and password).

Download the Migration Tool from the following location: [Cloudera Flow Management Migration Tool Repository](#).

Cloudera Flow Management Migration Tool overview

The Cloudera Flow Management Migration Tool enables the semi-automatic migration of flows from Cloudera Flow Management 2.1.7.1000 powered by NiFi 1 to Cloudera Flow Management 4.0.0 powered by NiFi 2, using predefined transformation rules and logic. It supports a clear, step-by-step approach for migrating your data flows.

When using the Cloudera Flow Management Migration Tool, you can issue various commands that enable partial or complete migration of the incoming data flow. These commands are customizable through specific arguments to run either individual steps or the full migration process. The Migration Tool stops running after the command is completed and can be run again with the same or different commands, as needed. This modular design allows you to customize the migration workflow according to your needs.

Running Migration Tool commands does not modify the source data flow or the associated NiFi instances. Repeating the same command provides identical results, ensuring predictability. Exceptions are the generation of unique IDs, which may vary across runs, and potential differences in component order within the serialized flow.json file. While file-based comparisons may show variations, the functional outcome remains identical.

Different Migration Tool commands apply specific transformation logic to the input data flow. Each command must be configured appropriately to achieve the desired outcome. For comprehensive details on each command's functionality, as well as the arguments and parameters available for configuration, see the *Cloudera Flow Management Migration Tool Command Reference*.

The Migration Tool works with flow.json as the input source. From this source, the tool can perform the following actions:

- **Transforming variables to parameter contexts:** It translates variables into parameter contexts for improved organization and compatibility with NiFi 2.
- **Converting templates:** It extracts and converts templates into separate flow_definition.json files.
- **Updating components:** It updates components to align with NiFi 2 requirements wherever possible, potentially applying broader modifications when needed.

The Migration Tool generates various output files. These output artifacts are saved in the directory specified by the --outputDirectory argument. Results are organized into subdirectories (v1 for Stage 1 and v2 for Stage 2).

Related Information

[Cloudera Flow Management Migration Tool Command Reference](#)

Key features

The Cloudera Flow Management Migration Tool provides several features to support the migration of flows in smaller, manageable chunks, enhancing control and validation.

Reusability

Since the Migration Tool does not modify the input files and produces deterministic results (except for unique identifiers of newly generated components), you can rerun the migration on the same input multiple times. This can be useful if adjustments are needed before or after migration and the results need to be compared.

Activity log

The Activity Log is an important tool for supervising changes in the flow. It lists every modification and provides the reasoning (change-info) behind them.

Stages

The Migration Tool allows you to handle manual migration steps while continuing to use the source NiFi version for most of the migration process. This approach allows you to apply a high number of expected changes without encountering version-related differences, making validation and tracking modifications easier.

Separable commands

Flow migration steps such as template migration can be run separately, allowing you to work on smaller, more manageable parts of the flow.

Process groups

Migration can be scoped to specific process groups by setting the Process Group ID argument. This limits transformations to the specified group and its children, enabling targeted migrations.

Iterative migration (Loopback)

The main output of the Migration Tool, typically a flow.json file, can be used as input for subsequent migrations. This iterative approach helps identify and address additional issues after manual adjustments.

Important considerations

Business logic validation

The Cloudera Flow Management Migration Tool ensures that the migration follows the predefined ruleset, but since the NiFi 2 feature set differs from NiFi 1, some business behaviors may change and require human validation. This post-migration validation is essential to ensure the final flow meets your business requirements.

Generating an initial report

Although the Migration Tool lacks an explicit "preview" feature, performing a full migration generates artifacts that can provide an initial overview of compatibility issues. Review the Activity Log entries for an early assessment. The number of change entries in the Activity Logs, in relation to the flow size, as well as the count of manual-change-requests, serves as a reliable indicator of the expected number of changes.

Migration stages

Migration is performed in two sequential stages (referred to as “Stage 1” and “Stage 2”) ensuring a structured transition from source to target compatibility. A complete migration workflow consists of a Stage 1 migration of the incoming flow and a consecutive Stage 2 migration of the result of the Stage 1 migration. The final output of Stage 2 is the completed migration result: a fully migrated and NiFi 2 compatible flow.

Stage 1: Source compatibility

- Processes the incoming NiFi flow (input flow.json) to create a partially updated flow as the output, compatible with NiFi 1.



Note: This intermediate state is meant only for manual adjustments in preparation for Stage 2. Do not use a partially migrated flow for production or business operations.

- Running Stage 1 commands for a template, variable, component, or using the aggregate command for all steps results in transformations compatible with the original version.
- Generates a list of manual changes needed for further adjustments.
- Requires iterative execution after applying manual changes until the flow reaches stability (no further modifications are required).

Stage 2: Target compatibility

- Takes the output of Stage 1 and processes it as input to produce a flow compatible with NiFi 2.
- Applies transformations that go beyond the source NiFi version's capabilities to meet NiFi 2 standards.
- Ensures that the final output is a fully migrated and compliant data flow.

Both migration stages consist of multiple steps that can be executed individually or as a whole. For best results, Cloudera recommends completing all steps within a stage before proceeding to the next one to ensure consistency and avoid potential issues.

Migration steps are present in both stages but involve different transformations and serve different purposes:

- In **Stage 1**, component migration applies changes compatible with the source version.
- In **Stage 2**, component migration applies changes required by the target version.



Important: Running only a Stage 1 migration does not complete the full migration process and may produce incorrect or incomplete results.

Limitations and manual steps

Some aspects of the migration are not handled automatically by the Migration Tool and require manual intervention.

Deprecated components

Certain deprecated component types lack replacements in the target NiFi version or cannot be algorithmically replaced. These components are not migrated automatically when using the Migration Tool and need to be updated manually. In most cases, the Migration Tool provides information on the expected manual changes.

Custom components

Custom components, such as processors not included in the Cloudera Flow Management version in use, are not supported. The definitions of such components, including attributes and bundle versions, are preserved without modification. You must manually update these components.



Note:

As NiFi 1 and NiFi 2 component APIs are incompatible, custom implementations also need to be replaced.

Ghost components in templates

Components saved as ghost components within templates are not migrated during the template migration process.

Variables, dynamic properties, and flow file attributes

Overlapping names among variables, dynamic properties, and flow file attributes are not migrated automatically. You must manually resolve these conflicts and update the elements to ensure compatibility.

Recommended migration workflow

Cloudera recommends dividing the process into manageable phases and validating the results at each step. Below is the suggested migration workflow to ensure a smooth and successful transition.

Stage 1 migration

1. Template migration
 - a. Migrate templates for Stage 1.
 - b. Validate the results.
 - c. Address any issues and manual change requests.

2. Variable migration
 - a. Migrate variables to parameter contexts for Stage 1.
 - b. Validate the results.
 - c. Address any issues and manual change requests.
3. Component migration
 - a. Migrate components in smaller batches by specifying Process Group IDs for Stage 1.
 - b. Validate the results for each batch.
 - c. Address any issues and manual change requests.
 - d. Repeat iteratively until no further changes are required.
4. Management-level component migration
 - a. Perform component migration without specifying a Process Group ID to ensure management-level components such as Controller Services are migrated.
 - b. Validate the changes.
 - c. Address any issues and manual change requests.
5. Final validation
 - a. When all Stage 1 steps are complete, validate the entire flow in a NiFi 1 instance for final validation.
 - b. After this point, the flow will no longer be compatible with the source version.

Stage 2 migration: Target version compatibility

- Run the migration commands without specifying the Stage argument to process both stages in sequence and address any remaining compatibility issues with NiFi 2.

**Note:**

- This runs both Stage 1 and Stage 2, but if previous recommendations were followed, Stage 1 has no impact on the flow. While Stage 1 commands should no longer modify the flow, some Stage 2 commands depend on Stage 1 outputs as input (for example for template migration).
- Follow the same steps as in Stage 1: Migrate templates # Migrate variables # Migrate components, while validating the results and addressing manual adjustment needs.
 - Stage 2 migration is not iterative in a sense that the output of Stage 2 cannot be reintroduced into the migration process. To ensure a seamless workflow, fully complete Stage 1 before starting Stage 2, and run Stage 2 only once. The final output will reflect all transformations applied by the Migration Tool. Resolving manual change requests will not trigger further migration events. Any modifications beyond this point are considered business logic adjustments and fall outside the tool's scope.

For detailed instructions and examples of the end-to-end migration process, see *Migrating a data flow*.

Related Information

[Migrating a data flow](#)

Migration best practices

Migrating complex data flows requires careful planning, validation, and following a structured strategy to ensure success. While the Cloudera Flow Management Migration Tool automates many aspects of the process, manual oversight and adjustments in the migration process are essential for achieving accurate results. Additionally, some expectations from NiFi 2 cannot be fully addressed through automation alone.

To achieve a successful migration, Cloudera recommends a systematic approach combining the tool's features with manual intervention and careful validation.

Follow an iterative refinement process in Stage 1!

Stage 1 migration automatically converts the flow, applying transformations while maintaining compatibility with NiFi 1, and a list of manual changes is generated. You have to manually edit your flow to fix the requested changes. After making these manual edits, the updated flow should be used as input for the next iteration of Stage 1. Repeat this process on the modified flow iteratively until the flow reaches stability, meaning that no automatic changes are introduced during conversion (there are no change entries in the Activity Log) and there are no manual change requests in the Activity Log that would require attention in Stage 1.

**Important:**

If a manual change request can be resolved using NiFi 1 features, it is best to address it during Stage 1 migration iterations. However, if the resolution requires a NiFi 2 capability, you may defer it until Stage 2. While migration can proceed with unresolved manual change requests, Cloudera recommends resolving them as early as possible. Only postpone fixes if a NiFi 2 feature is required.

At this point, Stage 1 can be considered fully completed with all issues resolved, allowing you to proceed to Stage 2.

**Note:**

If you make any manual adjustments to your flow after finishing Stage 1 migration, rerun Stage 1 before proceeding to Stage 2.

Partition the migration process for better control!

Migrating complex data flows requires a more structured approach. Cloudera advises not to migrate the entire flow in a single Stage 1 migration. Instead, partition the migration and focus on one part of the flow at a time to simplify troubleshooting and applying manual changes. You can use the following methods:

- Run the migration in smaller steps, starting with templates, followed by variables, and then components.
- Run the migration by process group.

Both approaches allow you to focus on smaller, more manageable sections of the flow. For larger flows, you can also combine these two methods. The migration logic you choose should be tailored to your specific needs and flow structure. Keep in mind that multiple levels of partitioning may require additional coordination, so it is important to select a level of granularity that provides more benefit than added complexity.

Validate after every iteration!

Load the modified flow into a NiFi instance matching the version of the stage you are on. Confirm functionality and resolve any manual inspection requests from the Activity Log.

Always review the Activity Log!

You can use the Activity log to understand the rationale behind changes and identify what manual adjustments are needed in the flow.

Run a full migration after each stage to confirm completeness!

As part of your final validation, run a full migration after each stage. This ensures that no part of the flow was overlooked during partitioning.

- At the end of Stage 1, perform a full migration only using the Stage 1 restriction to verify completeness.
- At the end of Stage 2, run a final full migration without restrictions to ensure the entire flow migration is complete.

Migrating a data flow using the Cloudera Flow Management Migration Tool

The migration workflow consists of two main parts: preparation and migration. The *Preparing for migration* section walks you through the preparation steps you should take before starting any migrations, and the *Migrating a flow* section provides examples demonstrating how you can run different types of migrations.

Preparing for migration

While the Cloudera Flow Management Migration Tool does not require NiFi to be running, certain NiFi-related steps are needed to prepare for a successful migration.

Setting up your environment

Learn about the necessary setup steps, including installing Java, configuring directories, preparing the source NiFi instance, and acquiring sensitive properties.

Before you begin

- Migrating flows to Cloudera Flow Management 4.0.0 requires upgrading to Cloudera Flow Management 2.1.7.1000 before moving to 4.0.0.

For detailed guidance on preparing for the upgrade to NiFi 2 in Cloudera Flow Management 4.0.0, see [Preparing for upgrade to NiFi 2](#).

- Download the Cloudera Flow Management Migration Tool from the [Cloudera Flow Management Migration Tool Repository](#).

Procedure

1. Install and set up Java.
 - a) Install Java 21 on the machine where you are performing the migration.
 - b) Add Java 21 to the PATH environment variable.
2. Download and unpack the Cloudera Flow Management Migration Tool on the machine where the migration will run.

For example: `/etc/nifi-migration-tool-bin`

3. Configure the following directories to support the Migration Tool and its dependencies.

a) NiFi 1 and 2 library folders:

- Provide read-only access to the NiFi 1 lib folder or equivalent directory containing all necessary NAR artifacts.

For example: /etc/nifi1/lib

- Provide read-only access to the NiFi 2 lib folder or equivalent directory containing all necessary NAR artifacts.

For example: /etc/nifi2/lib

b) Work directories:

- Create two, initially empty directories for the Migration Tool to unpack NiFi 1 and NiFi 2's dependency NAR files.

For example: /etc/nifi-1-work and /etc/nifi-2-work

- Ensure the Migration Tool has write access to these directories that will serve as NiFi 1 and 2 working areas.
- Avoid using these directories for other purposes.

c) Output directory:

- Create an empty directory to store migration artifacts generated by the Migration Tool.

For example: /etc/migration-tool-output

- Multiple runs of the Migration Tool using the same output directory will overwrite existing files. For example, if a template migration is run before a component migration, template migration files may remain intact, while other files, such as the Activity Log, are overwritten during the second run.

4. Prepare the source NiFi Instance.

- a) Ensure all components are in a valid state, as invalid components may not be migrated correctly.
- b) Stop and offload all processors.
- c) Stop the NiFi instance and ensure all FlowFiles have been processed.

5. Export the flow.json.gz file from the source NiFi 1 instance to the machine performing the migration.

For example: /etc/flow.json.gz

**Important:**

Do not forget to disable all components in the source NiFi instance before exporting the flow.json file for migration purposes.

- a) Unzip flow.json.gz file to product flow.json.
- b) Ensure that the Migration Tool has access to this file.

6. Acquire the Sensitive Properties Key and Algorithm.

- a) Obtain the `nifi.sensitive.props.key` and `nifi.sensitive.props.algorithm` from the source NiFi 1 instance.

You will need these when updating the `migration.tool.properties` configuration file.

**Note:**

You will need to use the same key and algorithm later to load the migrated flow into the NiFi 2 instance.

- b) For deployments in Cloudera on premises or Cloudera on cloud, follow these steps to acquire the sensitive key:

1. Configure Cloudera Manager to allow API calls to return sensitive values. For more information, see *Disabling Redaction of sensitive information when using the Cloudera Manager API*.
2. Run the following API call:

```
https://[***HOST***].root.comops.site:7183/api/v54/clusters/[***CLUSTER_NAME***]/services/[***SERVICE_NAME***]?view=EXPORT
```

Replace `[***HOST***]`, `[***CLUSTER_NAME***]`, and `[***SERVICE_NAME***]` with the appropriate values for your cluster.

3. Locate the `random.nifi.sensitive.props.key` value, which serves as the `nifi.sensitive.props.key` property.
 4. Revert Cloudera Manager to its original state to disable sensitive value return.
7. If the source and target instances use different sensitive properties, update the `nifi.sensitive.props.key` and `nifi.sensitive.props.algorithm` properties using a copy of the original `nifi.properties` file.

For more information, see the *NiFi System Administrator's Guide*.

Related Information

[Disabling redaction of sensitive information when using the Cloudera Manager API](#)

[NiFi System Administrator's Guide](#)

Configuring the Migration Tool

Learn how to set up and customize the Migration Tool's properties file to ensure the proper operation of the Cloudera Flow Management Migration Tool.

Procedure

1. Locate `migration.tool.properties` configuration file in the `conf` directory of the Migration Tool root.

2. Configure the directory paths by updating the following mandatory properties in the configuration file.

embedded.nifi.v1.nar.directory

- Specify the read-only directory containing the NAR files associated with the source NiFi version.
- It can be a lib folder of an actual NiFi instance with a matching version or a different directory containing all NAR artifacts from a given Cloudera distribution. For example: /etc/nifi1/lib
- Although the Migration Tool does not need a running NiFi instance and does not modify the contents of this directory, it requires read access to these files.

embedded.nifi.v1.working.directory

- Designate an empty workspace for the Migration Tool when it is working with the source NiFi version.
- It should be initially empty and the Migration Tool must have write access to it, allowing the Migration Tool to perform all necessary operations.
- It must be separate from the source NiFi instance's working directory and should not be used for any other purposes to avoid potential conflicts.
- After the migration is complete, it can be cleaned up if needed.

embedded.nifi.v2.nar.directory

- Specify the read-only directory containing the NAR files associated with the target NiFi version.
- It can be a lib folder of an actual NiFi instance with a matching version or a different directory containing the required NAR artifacts.
- Although the Migration Tool does not need a running NiFi instance and does not modify the contents of this directory, it requires read access to these files.

embedded.nifi.v2.working.directory

- Designate an empty workspace for the Migration Tool when it is working with the target NiFi version.
- It should be initially empty and the Migration Tool must have write access to it, allowing the Migration Tool to perform all necessary operations.
- It must be separate from the target NiFi instance's working directory and should not be used for any other purposes to avoid potential conflicts.
- After the migration is complete, it can be cleaned up if needed.

3. To configure the sensitive properties, update the following mandatory properties in the configuration file.

embedded.nifi.sensitive.props.key

- Provide the `nifi.sensitive.props.key` from the source NiFi instance for handling sensitive properties.
- This key is used to decrypt sensitive values during migration and re-encrypt them for the target NiFi instance.
- During migration, the target NiFi must use the same `nifi.sensitive.props.key` as the source, as configured in the Migration Tool properties.
- If the target NiFi requires a different sensitive properties key than the source NiFi, update the value only after completing the migration.

embedded.nifi.sensitive.props.algorithm

- Provide the `nifi.sensitive.props.algorithm` from the source NiFi instance required for handling sensitive properties.
- This algorithm is used to decrypt sensitive values during migration and re-encrypt them for the target NiFi instance.
- During migration, the target NiFi must use the same `nifi.sensitive.props.algorithm` as the source, as configured in the Migration Tool properties.
- If the target NiFi requires a different sensitive properties algorithm than the source NiFi, update the value only after completing the migration.

4. Save the modifications of the `migration.tool.properties` file.

Migrating a flow


This section provides step-by-step examples of how to run different types of migrations using the Cloudera Flow Management Migration Tool.

Migrating templates

Learn how to use the Cloudera Flow Management Migration Tool to extract and transform templates for compatibility with NiFi 2. NiFi 2 does not support templates, so this step is required to ensure compatibility with the target version.

Example flow for migrating a template

You have a flow that has a template named **elastic_template**, which was created from a process group called **elastic**. This process group contains the following simple flow:



The process group has a variable, **Elasticsearch Index**, referenced in the index property of the PutElasticsearchJson processor.

Configure Processor | PutElasticsearchJson 1.26.0.2.17.1000-46

⚠ Invalid

SETTINGS
SCHEDULING
PROPERTIES
RELATIONSHIPS
COMMENTS

Required field
⊘
+

Property	Value
Identifier Attribute	No value set
Index Operation	index
Index	#{Elasticsearch Index}
Type	No value set
Script	No value set
Scripted Upsert	false
Dynamic Templates	No value set
Batch Size	100
Character Set	UTF-8
Client Service	No value set
Log Error Responses	false
Output Error Responses	false

CANCEL
APPLY

Your goal is to migrate this template to NiFi 2 used in Cloudera Flow Management 4.0.0.

Before you begin

1. Stop all processors and disable all controller services in NiFi.
2. Stop NiFi.
3. Copy the flow.json.gz file from NiFi's conf directory to the Migration Tool's input folder (/etc/migration-tool-input).
4. Unzip the file to obtain flow.json.

Procedure

1. Run Stage 1 template migration using the following command.

```
bin/migration.sh nifi migrate-templates \
-i /etc/migration-tool-input/flow.json \
-od /etc/migration-tool-output/templates \
-s STAGE_1
```

This generates a v1 folder that contains the output files of the migration.

```
templates
### v1
### NiFi_Flow_ad3a86a6-0194-1000-78d9-6374298d9a0c
#   ### elastic_template.json
#   ### elastic_template.xml
```



```
### activity_log.json
```

The folder name NiFi_Flow_ad3a86a6-0194-1000-78d9-6374298d9a0c indicates that the exported template belongs to the root process group, which is named NiFi Flow in NiFi. The unique ID for this group is ad3a86a6-0194-1000-78d9-6374298d9a0c.

elastic_template.json

- Flow definition containing the contents of the original template
- Equivalent to manually converting a template to a flow definition in NiFi 1 by:
 - a. Instantiating the template on the canvas.
 - b. Right-clicking the process group.
 - c. Selecting Download flow definition without external services.
- Modified by the Migration Tool to ensure compatibility:
 - Variables (not supported in NiFi 2) converted into parameters
 - Parameter context created to hold the new parameters
 - Processors updated to reference parameters instead of variables
- This is still a NiFi 1 flow definition, so do not load this file directly into NiFi 2. Use an empty NiFi 1 instance to review the modification made by the tool. The flow will appear unchanged, but variables will be replaced with parameters of the same name, which are now referenced by the processor instead.
- Changes in syntax:

Variables were previously referenced using `${}`, whereas parameters are now referenced using `#{}.`

 - **NiFi 1 variables:** `${variable_name}`
 - **NiFi 2 parameters:** `#{parameter_name}`

elastic_template.xml

- Original template exported from NiFi 1
- Identical to downloading it from the Templates menu in NiFi 1
- Compatible only with NiFi 1 as NiFi 2 does not support templates and cannot parse XML template files

activity_log.json

- Log of all actions performed during this stage of the migration.
- The following were changes made during the template migration:
 - A new parameter context was created with a new parameter in it.
 - The parameter context was assigned to the process group.
 - The PutElasticsearchJson processor was updated to reference the new parameter.
- Components are referenced by a unique ID, not by name.

2. Validate the Stage 1 template migration output.
 - a) Check the new flow definition in a NiFi 1 instance to verify that the flow matches your expectations.
 1. Start your NiFi 1 instance.
 2. Create a new process group.
 3. Load elastic_template.json to NiFi 1.
 4. Confirm that variables were correctly converted to parameters.
 - b) Review activity_log.json and address any manual-change-requests or manual-inspection-requests.
 - c) If there are manual-change-requests or manual-inspection-requests to handle, follow these steps:
 1. Start a NiFi 1 instance.
 2. Load the v1/<template_name>.json flow definition by creating a new process group and uploading the JSON file.
 3. Make the modifications on the NiFi canvas, indicated by the manual-change-requests or manual-inspection-requests in the v1/activity_log.json.
 4. Stop the NiFi instance.
 5. Fetch the flow.json.gz file from the NiFi conf directory.
 6. Unzip it to produce a flow.json file.
 7. Perform variable migration and then component migration on this flow.json file. For instructions, see the example scenario for variable and component migration.
 - d) If no additional changes are required, proceed to Stage 2.
3. Run full template migration (Stage 1 and 2), using the following command.

```
bin/migration.sh nifi migrate-templates \
-i /etc/migration-tool-input/flow.json \
-od /etc/migration-tool-output/templates
```

The contents of the previously generated v1 folder will be overwritten, but this is not a concern as the tool generates the same v1 output as before. Additionally, a v2 directory is created, containing the output files of the Stage 2 part of the migration.

```
v2
### NiFi_Flow_ad3a86a6-0194-1000-78d9-6374298d9a0c
#   ### elastic_template.json
### activity_log.json
```

elastic_template.json

- NiFi 2-compatible flow definition
- Contains the contents of the original template converted into an exported process group.
- This version is compatible with NiFi 2, but no longer supports NiFi 1.



Note:

An elastic_template.xml is not generated as NiFi 2 does not support XML templates.

activity_log.json

- List of all actions performed during this stage of the migration.

4. Validate the Stage 2 template migration output.
 - a) Check the new flow definition in a NiFi 2 instance to verify that the flow matches your expectations.

1. Start your NiFi 2 instance.
2. Create a new process group.
3. Load `elastic_template.json` into a NiFi 2 instance.

The new process group will be called **elastic_template** and will contain another process group named **elastic**, matching the name of the process group that was converted into a template in NiFi 1 before you started the migration.

4. Verify parameter replacements and processor updates.
- b) Review `activity_log.json` and address any `manual-change-requests` or `manual-inspection-requests`.

In this case, you can see the following `manual-change-request`:

```
{
  "sequence" : 4,
  "type" : "manual-change-request",
  "subject" : "a1468d69-3f30-3f83-a7c1-91dad09890f7",
  "message" : "New relationship [original] has been added, it has to be
connected to a downstream processor or terminated.",
  "context" : {
    "rule" : "6501a693-b524-4d1a-b11e-69747c2651e8"
  }
}
```

1. Open the `v2/elastic_template.json` file and search for the "subject" ID, `a1468d69-3f30-3f83-a7c1-91dad09890f7`. This ID refers to the `PutElasticsearchJson` processor.
2. Go to the NiFi 2 canvas and check the processor. You will find that it has an unbound "original" relationship that needs to be connected to a downstream component.
3. Make the required change manually on the canvas.
4. Once done, export the process group. This exported process group is now a fully NiFi 2-compatible version of the original template.
5. Save the file.

Results

You have finished the template migration process.

Migrating variables

Learn how to use the Cloudera Flow Management Migration Tool to convert variables to parameters and parameter contexts.


Example flow for migrating variables

The following NiFi flow is used to demonstrate both variable and component migration. It illustrates how variables are used within process groups and how they are referenced by individual components in the NiFi flow.

The flow consists of two process groups:

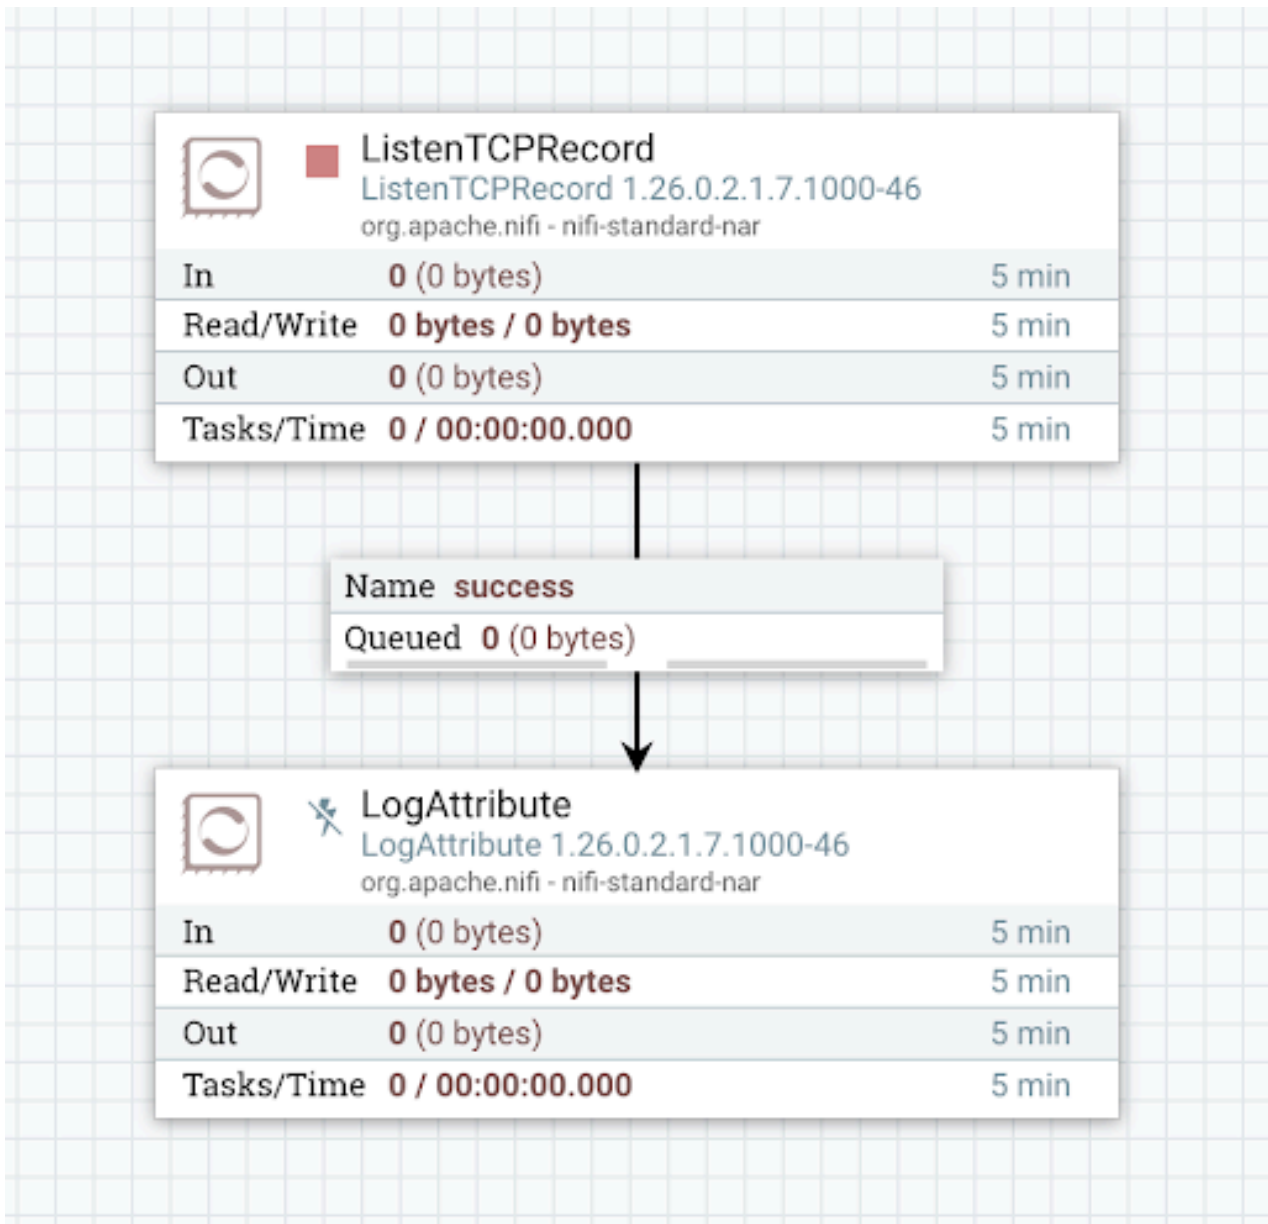
- **TCP Listener** (ID: `b41940d7-0194-1000-42fc-458834630567`)
- **Elastic** (ID: `b42881c7-0194-1000-3cdf-1bd453a0ed0f`)

At the root level, the flow is structured as follows:



TCP Listener Process Group

This process group contains the following simple flow:



The process group defines a variable called **TCP Listener Port**, which is referenced by the ListenTCPRecord processor.

Configure Processor | ListenTCPRecord 1.26.0.2.1.7.1000-46

■ Stopped

SETTINGS
SCHEDULING
PROPERTIES
RELATIONSHIPS
COMMENTS

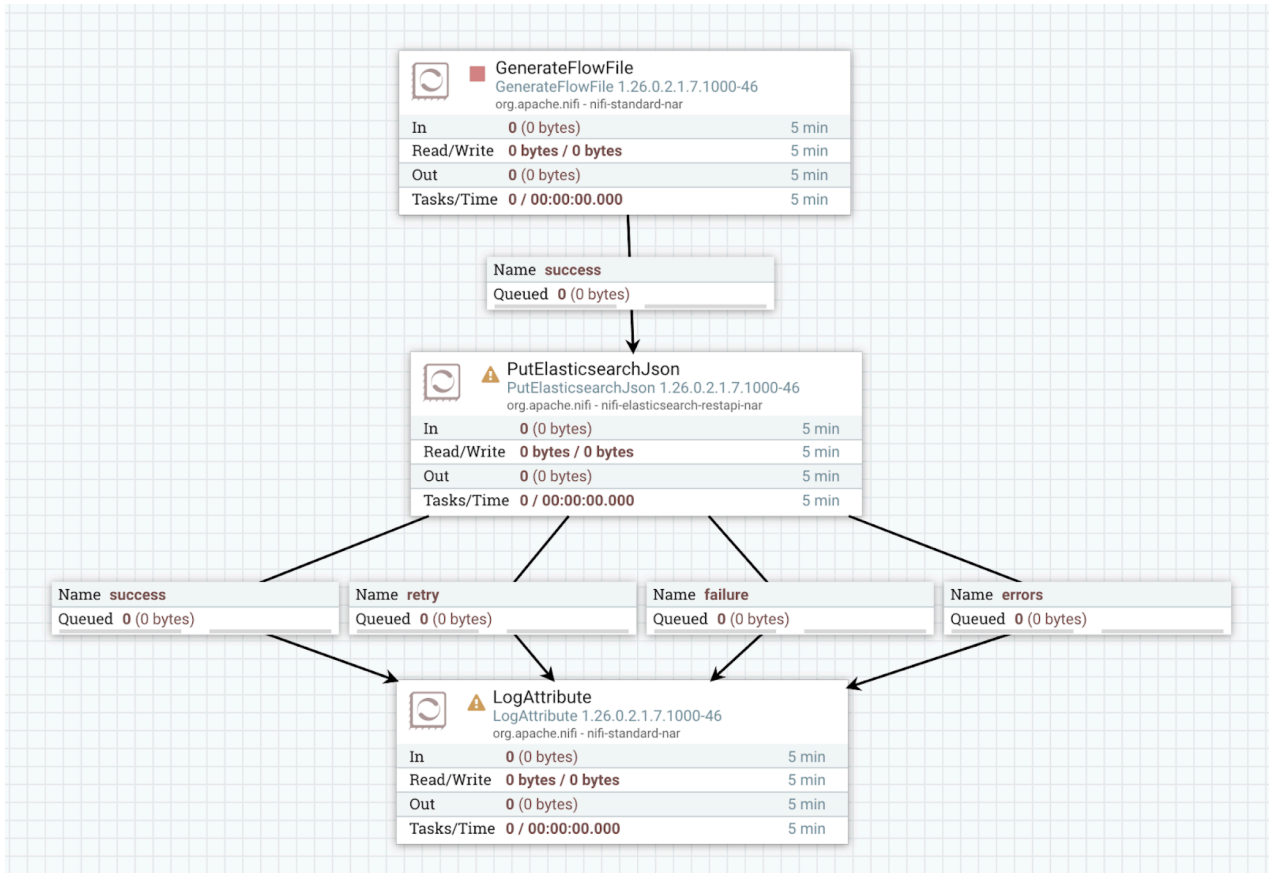
Required field
☑
+

Property	Value	
Local Network Interface	🔍 No value set	
Port	🔍 \${TCP Listener Port}	
Max Size of Socket Buffer	🔍 1 MB	
Max Number of TCP Connections	🔍 2	
Read Timeout	🔍 10 seconds	
Record Reader	🔍 JsonTreeReader	→
Record Writer	🔍 AvroRecordSetWriter	→
Read Error Strategy	🔍 Transfer	
Record Batch Size	🔍 1000	
SSL Context Service	🔍 No value set	
Client Auth	🔍 NONE	

CANCEL
APPLY

Elastic Process Group

This process group also contains a simple flow:



The process group defines a variable called **Elasticsearch Index**. This variable is referenced in the PutElasticsearchJson processor's "Index" property.

Configure Processor | PutElasticsearchJson 1.26.0.2.1.7.1000-46

Invalid

SETTINGS
SCHEDULING
PROPERTIES
RELATIONSHIPS
COMMENTS

Required field

Property	Value
Identifier Attribute	No value set
Index Operation	index
Index	#{Elasticsearch Index}
Type	No value set
Script	No value set
Scripted Upsert	false
Dynamic Templates	No value set
Batch Size	100
Character Set	UTF-8
Client Service	No value set
Log Error Responses	false
Output Error Responses	false

CANCEL
APPLY

Your goal is to migrate the variables used in Cloudera Flow Management 2.1.7.1000 powered by NiFi 1 to parameters used in Cloudera Flow Management 4.0.0 powered by NiFi 2.



Note: NiFi 2 supports only parameters, not variables, so variable migration is required to ensure compatibility with the target version.

The example guides you through variable migration one process group at a time, simplifying the process and maintaining a clear activity log. While this example flow is simple, the step-by-step approach shows how this method improves clarity for more complex migrations. In real-world scenarios, you have to define a migration strategy based on your flow's structure, individual process group migration may not always be necessary. Migrating a parent process group automatically applies the changes recursively to its descendants.

Before you begin

1. Stop all processors and disable all controller services in NiFi.
2. Stop NiFi.
3. Copy the flow.json.gz file from NiFi's conf directory to the Migration Tool's input folder (/etc/migration-tool-input).
4. Unzip the file to obtain flow.json.

Procedure

1. Run Stage 1 variable migration on the TCP Listener process group, using the following command.

```
bin/migration.sh nifi migrate-variables \
-i /etc/migration-tool-input/flow.json \
-od /etc/migration-tool-output/variables \
```

```
-pgid <process_group_id> \  
--stage STAGE_1
```

This generates a v1 folder that contains the output files of the Stage 1 migration.

```
variables  
### v1  
### activity_log.json  
### migrated_flow.json
```

activity_log.json

- Log of all actions performed during this stage of the migration.
- The following were changes made during the template migration:
 - A new parameter context was created with a new parameter called **TCP Listener Port**, which replaces the corresponding variable.
 - The **TCP Listener Port** variable was removed.

migrated_flow.json

- A modified NiFi 1 flow, which is not compatible with NiFi 2 yet.
- It contains everything the original flow did, except the **TCP Listener** process group now references a parameter instead of the removed variable.

2. Validate the Stage 1 variable migration output for the **TCP Listener** process group.

a) Load the migrated_flow.json into a NiFi 1 instance and check the flow.

1. Ensure NiFi is not running.
2. Go to NiFi's conf directory and back up the flow.json.gz and flow.xml.gz files.
3. Delete the original flow.json.gz and flow.xml.gz files.
4. Rename the migrated_flow.json to flow.json and compress it using gzip to create flow.json.gz.
5. Copy the newly created flow.json.gz file to NiFi's conf directory.
6. Start NiFi and check the flow.

b) Review the activity_log.json file and check for any manual-change-requests or manual-inspection-requests. If none are present, proceed to running Stage 1 variable migration on the **Elastic** process group.

1. If manual changes are necessary, load the migrated_flow.json to NiFi, update and validate it.
2. Once the flow is validated and meets expectations, continue with the next step using the new flow.json.gz file (unzipped to produce flow.json).

3. Run Stage 1 variable migration on the **Elastic** process group.

a) Move migrated_flow.json from Step 2 into the input folder (/etc/migration-tool-input) and rename it to flow.json for clarity.

b) Make a backup of the output folder (/etc/migration-tool-output/variables) before running the next migration step.



Note: This step is crucial because the same output folder is used as in the previous migration, and its contents will be overwritten. To retain a record of changes made at each stage, it is better to create a backup.

c) Run Stage 1 variable migration on the **Elastic** process group using the following command.

```
bin/migration.sh nifi migrate-variables \  
-i /etc/migration-tool-input/flow.json \  
-od /etc/migration-tool-output/variables \  
-pgid b42881c7-0194-1000-3cdf-1bd453a0ed0f \  

```



```
--stage STAGE_1
```

This generates a v1 folder that contains the two output files of the Stage 1 migration.

activity_log.json

- Log of all actions performed during this stage of the migration.
- The following were changes made during the template migration:
 - A new parameter context was created with a new parameter called **Elasticsearch Index**, which replaces the corresponding variable of the same name.
 - The new parameter is referenced from the PutElasticsearchJson processor.
 - The **Elasticsearch Index** variable was removed.

migrated_flow.json

- A modified NiFi 1 flow, which is not compatible with NiFi 2 yet.
 - It contains everything the original flow did, but the **Elastic** process group now references a parameter instead of the removed variable.
4. Validate the Stage 1 variable migration output for the **Elastic** process group.
 - a) Load the migrated_flow.json into a NiFi 1 instance and check the flow.
 - b) Review the activity_log.json file and check for any manual-change-requests or manual-inspection-requests. If none are present, proceed to the next step.
 - c) If manual changes are necessary, update the migrated_flow.json on the NiFi canvas after loading it. Once the flow is validated and meets expectations, continue with the next step.

At this stage, both process groups no longer contain variables and use parameters instead. If the flow meets your expectations, you can either run a full variable migration to validate your flow in NiFi 2 or proceed with migrating the components.

5. Run full variable migration (Stage 1 and 2) by using the following command.

```
bin/migration.sh nifi migrate-variables \
-i /etc/migration-tool-input/flow.json \
-od /etc/migration-tool-output/variables
```

The contents of the previously generated v1 folder will be overwritten, but this is not a concern as the tool generates the same v1 output as before.

Additionally, a v2 directory is created, containing the output files of the Stage 2 part of the migration.

migrated_flow.json

- NiFi 2-compatible flow in terms of variables
- You can load it into NiFi 2 and validate it

activity_log.json

- List of all actions performed during this stage of the migration.

Stage 2 migration allows you to validate the result of the variable migration process in NiFi 2. However, component migration should not be performed on the v2/migrated_flow.json because the input of component migration must be a NiFi 1 flow.

6. Proceed with component migration using the migrated_flow.json output from step 3.

Migrating components

Learn how to use the Cloudera Flow Management Migration Tool to migrate components to parameters and parameter contexts.

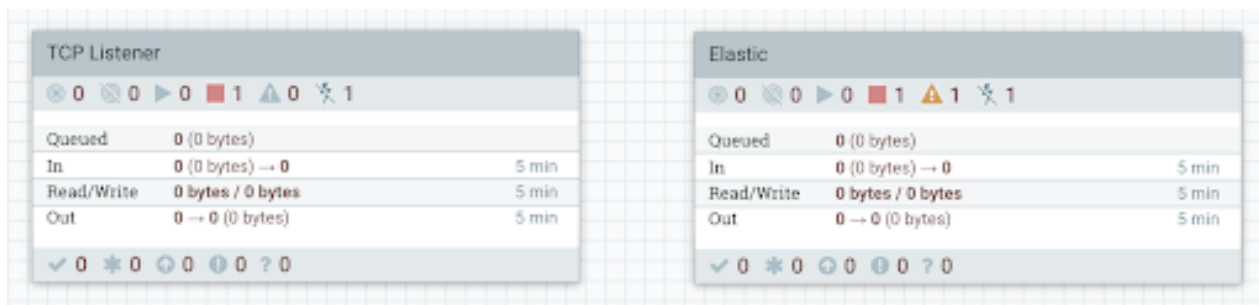
Example flow for migrating components

The following NiFi flow is used to demonstrate both variable and component migration.

It consists of two process groups:

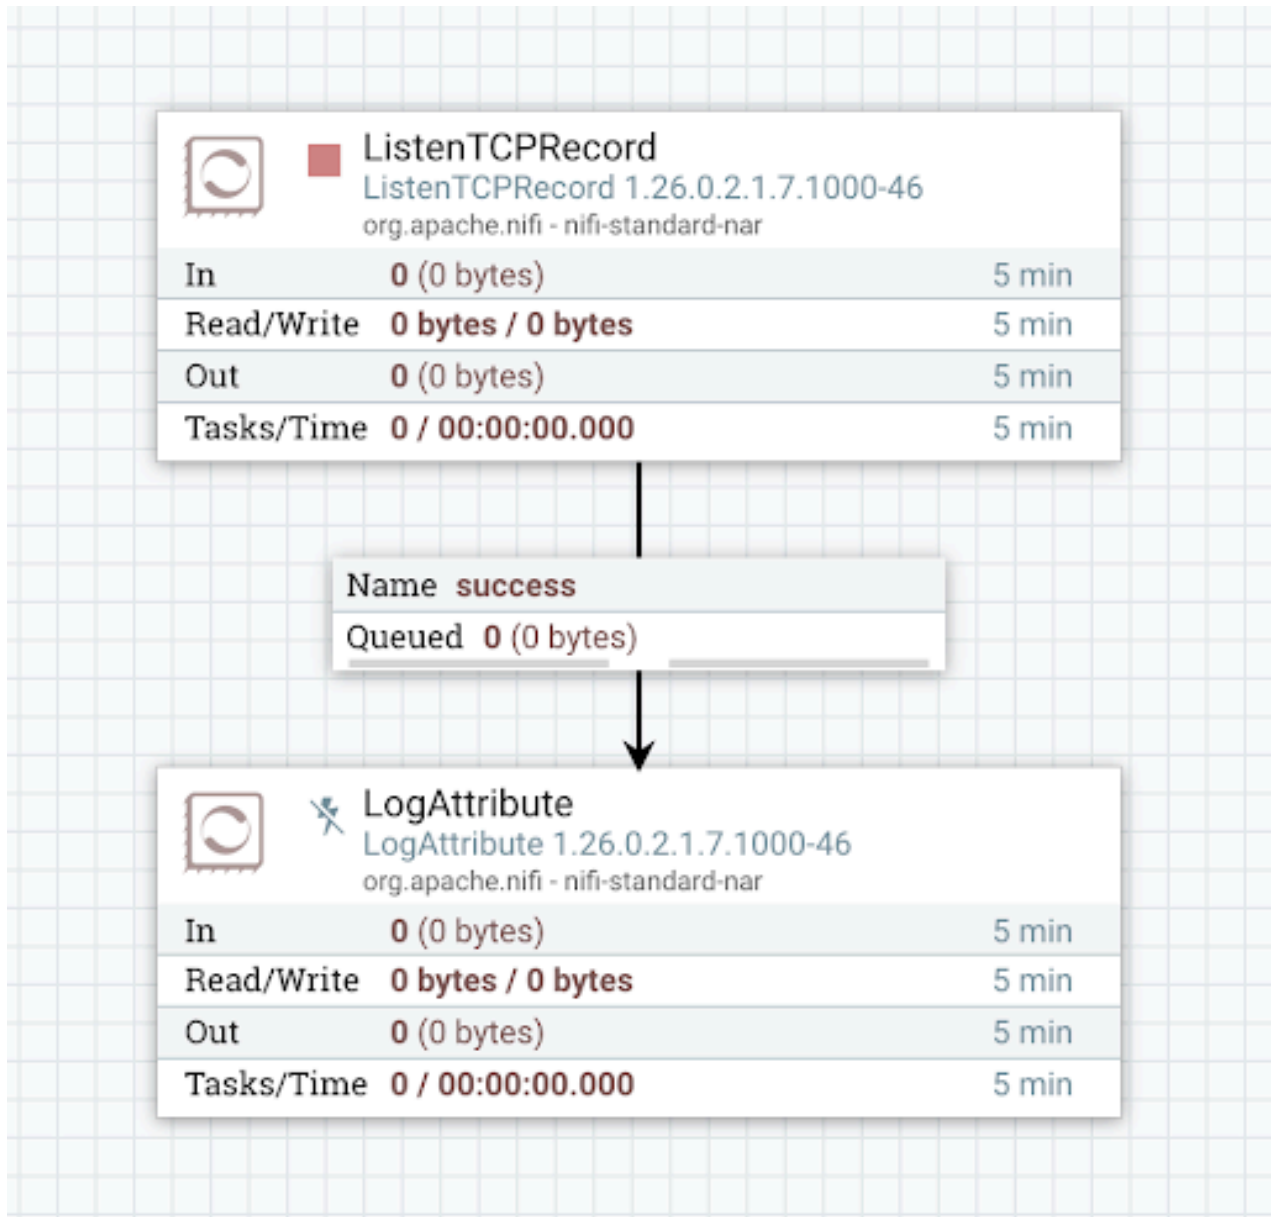
- **TCP Listener** (ID: b41940d7-0194-1000-42fc-458834630567)
- **Elastic** (ID: b42881c7-0194-1000-3cdf-1bd453a0ed0f)

At the root level, the flow is structured as follows:



TCP Listener Process Group

This process group contains the following simple flow:



The process group defines a variable called **TCP Listener Port**, which is referenced by the ListenTCPRecord processor.

Configure Processor | ListenTCPRecord 1.26.0.2.1.7.1000-46

■ Stopped

SETTINGS
SCHEDULING
PROPERTIES
RELATIONSHIPS
COMMENTS


Required field
☑
+

Property	Value	
Local Network Interface	🔍 No value set	
Port	🔍 <code>\${TCP Listener Port}</code>	
Max Size of Socket Buffer	🔍 1 MB	
Max Number of TCP Connections	🔍 2	
Read Timeout	🔍 10 seconds	
Record Reader	🔍 JsonTreeReader	→
Record Writer	🔍 AvroRecordSetWriter	→
Read Error Strategy	🔍 Transfer	
Record Batch Size	🔍 1000	
SSL Context Service	🔍 No value set	
Client Auth	🔍 NONE	

CANCEL
APPLY

Elastic Process Group

This process group also contains a simple flow:



The process group defines a variable called **Elasticsearch Index**. This variable is referenced in the PutElasticsearchJson processor's "Index" property.

Configure Processor | PutElasticsearchJson 1.26.0.2.1.7.1000-46

Invalid

SETTINGS
SCHEDULING
PROPERTIES
RELATIONSHIPS
COMMENTS

Required field

Property	Value
Identifier Attribute	No value set
Index Operation	index
Index	#{Elasticsearch Index}
Type	No value set
Script	No value set
Scripted Upsert	false
Dynamic Templates	No value set
Batch Size	100
Character Set	UTF-8
Client Service	No value set
Log Error Responses	false
Output Error Responses	false

CANCEL
APPLY

Your goal is to migrate the components used in Cloudera Flow Management 2.1.7.1000 powered by NiFi 1 to NiFi 2-compatible components used in Cloudera Flow Management 4.0.0 powered by NiFi 2.

The example guides you through component migration one process group at a time, simplifying the process and maintaining a clear activity log. While this example flow is simple, the step-by-step approach shows how this method improves clarity for more complex migrations. In real-world scenarios, you have to define a migration strategy based on your flow's structure, individual process group migration may not always be necessary. Migrating a parent process group automatically applies the changes recursively to its descendants.

Procedure

1. Move the migrated_flow.json file, the output of variable migration from Step 3, into the input folder (/etc/migration-tool-input) and rename it to flow.json for clarity.

This is a NiFi-1 compatible flow that no longer contains variables.

2. Run Stage 1 component migration on the **TCP Listener** process group using the following command.

```
bin/migration.sh nifi migrate-components \
-i /etc/migration-tool-input/flow.json \
-od /etc/migration-tool-output/components \
-pgid b41940d7-0194-1000-42fc-458834630567 \
--stage STAGE_1
```

This generates a v1 folder that contains the output files of the Stage 1 migration.

```
components
### v1
### activity_log.json
```

```
### migrated_flow.json
```

activity_log.json

- The log describes all the actions that were performed for this stage of the process group migration.
- Log of all actions performed during this stage of the process group migration.

migrated_flow.json

- A modified NiFi 1 flow, which is not compatible with NiFi 2 yet.
- It contains everything the original flow did, except the **TCP Listener** process group was modified as described in the activity log.

3. Validate the Stage 1 component migration output for the **TCP Listener** process group.

- Load the migrated_flow.json into a NiFi 1 instance and check the flow.
- Review the activity_log.json file and check for any manual-change-requests or manual-inspection-requests. If none are present, proceed to the next step.

In this example, you can see the following information in the activity log:

```
{
  "sequence" : 2,
  "type" : "change-info",
  "subject" : "b41966ad-0194-1000-a08d-a92489457356",
  "message" : "Component [org.apache.nifi.processors.standard.ListenTCPRecord] has been deprecated (NIFI-13509)",
  "context" : {
    "rule" : "36227b60-75f0-40dc-8caf-a2ec577aa54c"
  }
}
```

Search for the value of the “subject” element in NiFi’s search box.

You are directed to the ListenTCPRecord processor. No manual modifications are needed at this stage. However, it is important to note that the ListenTCPRecord processor is deprecated and is not available in NiFi 2. Once the full component migration (Stage 1 and Stage 2) is complete, instructions will be provided on how to handle this deprecation.

- If manual changes are necessary, update the migrated_flow.json on the NiFi canvas after loading it. Once the flow is validated and meets expectations, continue with the next step.

4. Run Stage 1 component migration on the **Elastic** process group.

- Move migrated_flow.json from Step 2 into the input folder (/etc/migration-tool-input) and rename it to flow.json for clarity.
- Make a backup of the output folder (/etc/migration-tool-output/components) before running the next migration step.



Note:

This step is crucial because the same output folder is used as in the previous migration, and its contents will be overwritten. To retain a record of changes made at each stage, it is better to create a backup.

- Run Stage 1 component migration on the **Elastic** process group using the following command.

```
bin/migration.sh nifi migrate-components \
-i /etc/migration-tool-input/flow.json \
-od /etc/migration-tool-output/components \
-pgid b42881c7-0194-1000-3cdf-1bd453a0ed0f \
```

```
--stage STAGE_1
```

This generates a v1 folder that contains the two output files of the Stage 1 migration.

activity_log.json

- Log of all actions performed during this stage of the migration.

migrated_flow.json

- A modified NiFi 1 flow, which is not compatible with NiFi 2 yet.
- It contains everything the original flow did, except the Elastic process group was modified with the actions described in the activity log.

5. Validate the Stage 1 component migration output for the **Elastic** process group.

- Load the migrated_flow.json into a NiFi 1 instance and check the flow.
- Review the activity_log.json file.

It contains a change-info entry that informs you of changes to a NiFi 2 processor, identified by the subject ID.

```
{
  "sequence" : 2,
  "type" : "change-info",
  "subject" : "b428aad6-0194-1000-d73d-9f2332a59f04",
  "message" : "Property [Max JSON Field String Length] has been added (NIFI-12343); Property [put-es-json-not_found-is-error] has been renamed to [put-es-not_found-is-error] (NIFI-12255); Property [put-es-json-error-documents] has been removed (NIFI-12255); Relationships [success] has been renamed to [original]; Relationship [successful] has been added (NIFI-12255);",
  "context" : {
    "rule" : "6501a693-b524-4d1a-b11e-69747c2651e8"
  }
}
```

To identify the processor affected by these changes, search for the value of the “subject” element in NiFi’s search box.

You are directed to the PutElasticsearchJson processor. No manual modifications are needed at this stage. However, it is important to note that changes will be applied to the PutElasticsearchJson processor during Stage 2 of the migration. Once the full component migration (Stage 1 and Stage 2) is complete, instructions will be provided on how to handle this change.

At this stage, you have completed Stage 1 of both variable and component migration for the process groups in your flow. After reviewing the logs, you confirmed that no manual changes were needed. You can proceed with a full component migration using the migrated_flow.json from Step 4.

6. Run full component migration (Stage 1 and 2).

- Move migrated_flow.json from Step 4 into the input folder (/etc/migration-tool-input) and rename it to flow.json for clarity.
- Make a backup of the output folder (/etc/migration-tool-output/components) before running the next migration step.



Note:

This step is crucial because the same output folder is used as in the previous migration, and its contents will be overwritten. To retain a record of changes made at each stage, it is better to create a backup.

- Run the full component migration using the following command.

```
bin/migration.sh nifi migrate-components \
-i /etc/migration-tool-input/flow.json \
```



```
-od /etc/migration-tool-output/components
```

This generates the following output:

```
components
### v1
#   ### activity_log.json
#   ### migrated_flow.json
### v2
    ### activity_log.json
    ### migrated_flow.json
```

The contents of the previously generated v1 folder will be overwritten. The contents represent the NiFi 1-compatible version of the migrated flow. Since the root process group only contained the two process groups on which you already performed Stage 1 component migration, the activity log will only include the same entries as those in Steps 2 and 4.

Additionally, a v2 directory is created, containing the output files of the Stage 2 part of the migration:

migrated_flow.json

- NiFi 2-compatible flow in terms of variables
- You can load it into NiFi 2 and validate it

activity_log.json

- List of all actions performed during this stage of the migration, and any manual steps that you need to perform.

7. Validate the Stage 2 component migration output.

- Load the v2/migrated_flow.json into a NiFi 2 instance and check the flow.
- Review the v2/activity_log.json file.

In this example, you can see the following manual-change-request entry.

```
{
  "sequence" : 3,
  "type" : "manual-change-request",
  "subject" : "b428aad6-0194-1000-d73d-9f2332a59f04",
  "message" : "New relationship [original] has been added, it has to be
connected to a downstream processor or terminated.",
  "context" : {
    "rule" : "6501a693-b524-4d1a-b11e-69747c2651e8"
  }
}
```

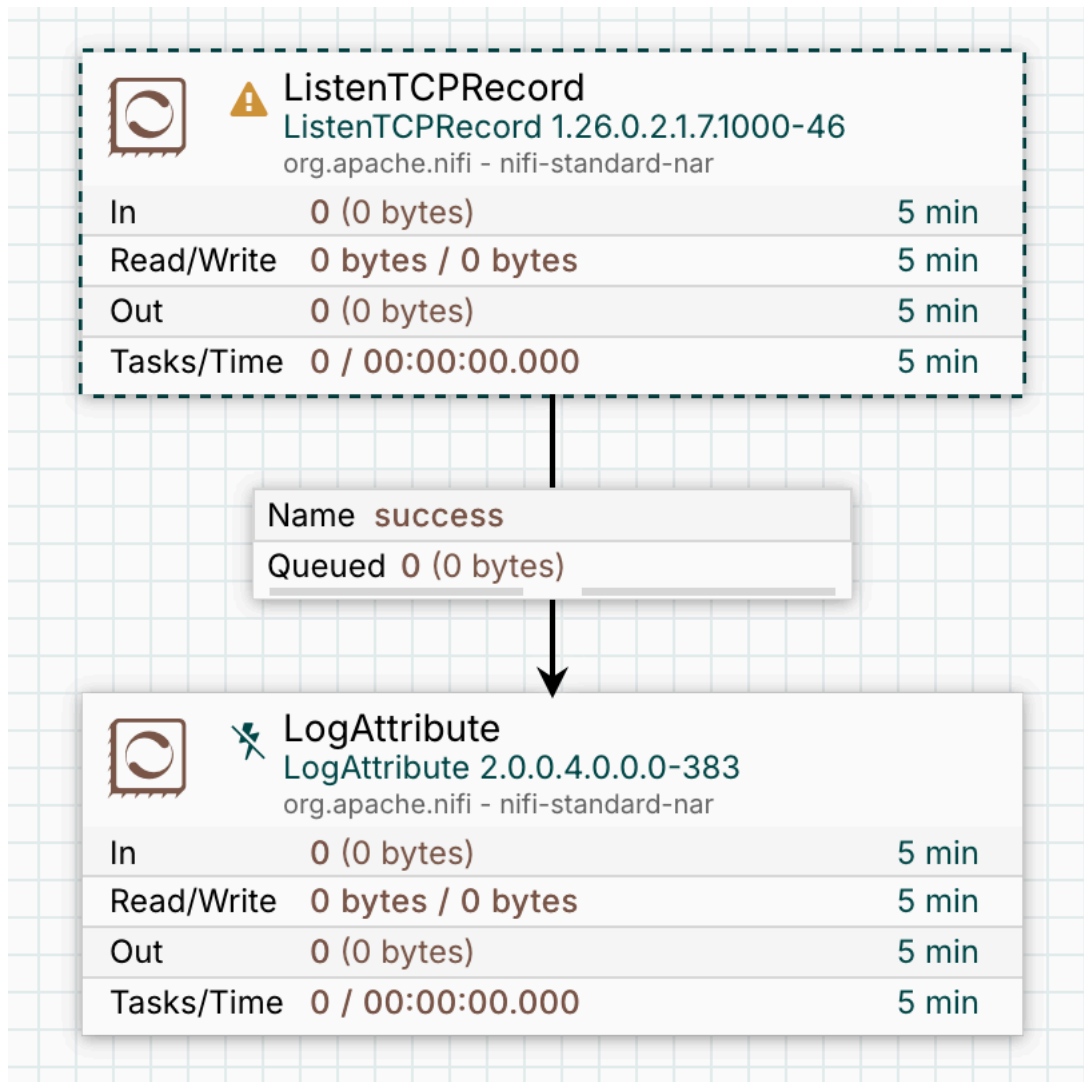
To identify the processor affected by these changes, search for the value of the “subject” element in NiFi’s search box. It refers to the PutElasticsearchJson processor. You will find an unbound “original” relationship that needs to be connected to a downstream processor or terminated. Make the necessary modifications manually. Once completed, this change request is resolved.

You can see another manual-change-request entry in the Activity Log.

```
{
  "sequence" : 6,
  "type" : "manual-change-request",
  "subject" : "b41966ad-0194-1000-a08d-a92489457356",
  "message" : "The component has been deprecated. It is suggested to r
eplace it with a combination of [org.apache.nifi.processors.standard.Lis
tenTCP] and [org.apache.nifi.processors.standard.ConvertRecord] processo
rs",
  "context" : {
    "rule" : "36227b60-75f0-40dc-8caf-a2ec577aa54c"
  }
}
```

}

To identify the processor affected by these changes, search for the value of the “subject” element in NiFi’s search box. It refers to the ListenTCPRecord processor. Open the TCP Listener process group. You can see that the processor is marked with dashed borders and its version number still shows that of the NiFi 1 instance. It means that it is a “ghost processor”, not available in NiFi 2. The log message suggests replacing it with a combination of ListenTCP and ConvertRecord processors.



Instantiate and configure the ListenTCP and ConvertRecord processors, connect them appropriately, and then remove the ghost processor.

Besides these changes there are no further manual-change-requests or manual-inspection-requests in the activity log.

8. Wrap up the migration.
 - a) Review the completed flow to ensure it meets your use cases and that all necessary changes have been made.
 - b) Enable the controller services associated with the flow in the NiFi 2 instance.
 - c) Once the controller services are enabled, start the processors in your migrated NiFi 2 flow.

Using migrate-all

Learn how to use the migrate-all command to migrate templates, variables, and components in a single operation.

This command combines the functionality of the `migrate-templates`, `migrate-variables`, and `migrate-components` commands for ease and convenience. While `migrate-all` simplifies the migration process, for larger flows, it is usually better to run the three migration commands separately. You may also want to migrate process groups individually to ensure activity logs and manual validation remain manageable.

The `migrate-all` command can be useful in the following scenarios:

- The flow you want to migrate is simple and of manageable size.
- The flow is large and complex, and you need an initial high-level overview of the migration outcome before performing a step-by-step migration. Running `migrate-all` provides a preview of the final migrated flow. The Activity Log shows you the actions taken during migration and helps assess the required manual intervention by listing `manual-change-requests` and `manual-inspection-requests`.

Migrating templates, variables, and components together using `migrate-all`

Before you begin

1. Offload all flowfiles from NiFi.
2. Stop all processors and disable all controller services in NiFi.
3. Stop NiFi.
4. Copy the `flow.json.gz` file from NiFi's `conf` directory to the Migration Tool's input folder (`/etc/migration-tool-input`).
5. Unzip the file to obtain `flow.json`.

Procedure

1. Run the following command to migrate the flow.

```
bin/migration.sh nifi migrate-all \  
-i /etc/migration-tool-input/flow.json \  
-od /etc/migration-tool-output
```

This generates the following output:

```
v1  
### NiFi_Flow_b4175f57-0194-1000-8470-9251a24519b4  
#   ### elastic_template.json  
#   ### elastic_template.xml  
### activity_log.json  
### migrated_flow.json  
  
v2  
### NiFi_Flow_b4175f57-0194-1000-8470-9251a24519b4  
#   ### elastic_template.json  
### activity_log.json  
### migrated_flow.json
```

2. Review the migration output.
 - The `v1` directory contains the NiFi 1-compatible version of the flow, including the exported template and its process group counterpart, along with the Activity Log.
 - The `v2` directory contains the NiFi 2-compatible version of the flow and its corresponding Activity Log.
3. Address issues in NiFi 1 (if needed).

If `v1/activity_log.json` contains `manual-change-requests`, follow these steps:

- a) Load `v1/migrated_flow.json` into your NiFi 1 instance.
- b) Apply the required manual changes.
- c) Run the `migrate-all` command again using the manually modified flow.

4. Validate the migrated flow.
 - a) Rename v2/migrated_flow.json to flow.json.
 - b) Compress it into flow.json.gz.
 - c) Load flow.json.gz into your NiFi 2 instance to inspect the migrated flow.
 - d) Open v2/activity_log.json and review any manual-change-requests and manual-inspection-requests.
 - e) Apply the required manual modifications to the flow.

Migrating a process group using migrate-all

For complex flows, you can migrate one process group and on one stage at a time. See below for an example workflow.

Procedure

1. Migrate the TCP Listener process group using the following command.

```
bin/migration.sh nifi migrate-all \
-i /etc/migration-tool-input/flow.json \
-od /etc/migration-tool-output -pgid b41940d7-0194-1000-42fc-458834630567 \
--stage STAGE_1
```

2. Use the output v1/migrated_flow.json (the result of the previous command moved to the input folder) as the input for migrating the **Elastic** process group.

```
bin/migration.sh nifi migrate-all \
-i /etc/migration-tool-input/flow.json \
-od /etc/migration-tool-output -pgid b42881c7-0194-1000-3cdf-1bd453a0ed0f \
--stage STAGE_1
```

3. Use the output v1/migrated_flow.json (the result of the previous command moved to the input folder) as the input for migrating the root process group and completing the Stage 2 migration.

```
bin/migration.sh nifi migrate-all \
-i /etc/migration-tool-input/flow.json \
-od /etc/migration-tool-output
```

This produces a NiFi 2-compatible version of the flow.



Important:

Always ensure that the input for a migration command is a NiFi 1 flow, not a NiFi 2 flow.

Cloudera Flow Management Migration Tool Reference

The reference documentation provides a comprehensive guide to the Migration Tool commands, which support partial or full migration of incoming data flows. Each command addresses a specific part of the migration process, offering flexibility and control to meet a variety of migration needs. Use this reference to understand the function and configuration of each command so that you can effectively tailor the migration process to your requirements.

Additionally, the reference documentation covers the Activity Log, which provides a detailed record of migration-related events, changes, and manual interventions. The log helps you track modifications, review necessary actions, and ensure a smooth transition during the migration process. Use this reference to gain insights into how your flows were migrated, identify any manual steps required, and ensure that all necessary changes were applied correctly.

Commands

This section contains information about the commands you can use with the Migration Tool when migrating your data flows from NiFi 1 to NiFi 2.

The command list includes:

- **Command descriptions:** A detailed explanation of each command and its role within the migration workflow.
- **Input arguments:** A list of common arguments shared by all commands, along with those specific to individual commands or use cases.
- **Syntax and usage examples:** Practical examples showcasing the correct syntax and usage for each command, helping you implement them in your environment.
- **Output artifacts:** Information on the output generated by each command, ensuring traceability and transparency throughout the migration process.

migrate-templates

Migrates NiFi templates using the Cloudera Flow Management Migration Tool.

Syntax

```
migrate-templates --input {filePath} --outputDirectory {directoryPath} [--stage {stageName}] [--processGroupId {pgId}]
```

Required arguments

--input

- Short format: `-i {filePath}`
- Specifies the path to a local file containing the flow in JSON format for processing.
- The file must be accessible to the Migration Tool, and reading privileges are required.
- The original file remains unmodified.

--outputDirectory

- Short format: `-od {directoryPath}`
- Specifies the directory where migration results will be saved.
- The directory must be accessible to the Migration Tool, and writing privileges are required.
- Existing content in the specified directory may be overwritten by subsequent runs.

Supported arguments

--stage

- Short format: `-s {stageName}`
- Specifies the stage(s) to process.
- To process only Stage 1 migration, use `STAGE_1`
- If omitted, both Stage 1 and Stage 2 are processed sequentially, using Stage 1's output as input for Stage 2.
- Providing unsupported or custom stage names may result in invalid stage error.

--processGroupId

- Short format: `-pgid {pgId}`
- Defines the starting point for migration, applying transformations to the specified group and its child groups.
- If omitted, the root group is processed.

**Note:**

The input flow.json file must represent the entire flow, regardless of whether this argument is used.

Output files of Stage 1 and Stage 2**Extracted templates**

The resulting templates from the migration process.

activity_log.json

It is a log file recording details of the migration process. For more information, see *Activity Log*.

flow_definitions.json

It is a flow definition representation of their respective template.

Examples**1Migrate templates processing Stage 1 and export templates**

```
./migration.sh nifi migrate-templates --input /etc/exported-flows/flow.json \
--outputDirectory /etc/output/ --stage STAGE_1 --processGroupId bb81df68-c
d6a-461a-b724-384265875b53
```

Processes the process group specified by the ID bb81df68-cd6a-461a-b724-384265875b53 only in Stage 1. Reads the input from /etc/exported-flows/flow.json, and saves the results in /etc/output/.

2Migrate templates processing all stages

```
./migration.sh nifi migrate-templates --input /etc/exported-flows/flow.json \
--outputDirectory /etc/output/
```

Processes both Stage 1 and Stage 2 sequentially, using the output of Stage 1 as input for Stage 2.

migrate-variables

Migrates variable configurations from NiFi 1 to NiFi 2 using the Cloudera Flow Management Migration Tool.

Syntax

```
migrate-templates --input {filePath} --outputDirectory {directoryPath} [--st
age {stageName}] [--processGroupId {pgId}]
```

Required arguments**--input**

- Short format: -i {filePath}
- Specifies the path to a local file containing the flow in JSON format for processing.
- The file must be accessible to the Migration Tool.
- The original file remains unmodified.

--outputDirectory

- Short format: -od {directoryPath}
- Specifies the directory where migration results will be saved.
- The directory must be accessible to the Migration Tool.
- Existing content in the specified directory may be overwritten by subsequent runs.

Supported arguments

--stage

- Short format: -s {stageName}
- Specifies the stage(s) to process.
- To process Stage 1 migration, use STAGE_1
- If omitted, both Stage 1 and Stage 2 are processed sequentially, using Stage 1's output as input for Stage 2.
- Providing unsupported or custom stage names may result in invalid stage error.

--processGroupId

- Short format: -pgid {pgId}
- Defines the starting point for migration, applying transformations to the specified group and its child groups.
- If omitted, the root group is processed.



Note:

The input flow.json file must represent the entire flow, regardless of whether this argument is used.

Output files of Stage 1 and Stage 2

NiFi flow

For example: migrated_flow.json.

activity_log.json

A log file recording the details of the migration process. For more information, see *Activity Log*.

Examples

1Migrate variables processing Stage 1

```
./bin/migration.sh nifi migrate-variables --input /etc/exported-flows/flow.json --outputDirectory /etc/output/ --stage STAGE_1
```

Processes the flow.json file only in Stage 1, located in /etc/exported-flows/. Saves the migrated variables and activity log in /etc/output/.

2Migrate variables processing all stages

```
./bin/migration.sh nifi migrate-variables --input /etc/exported-flows/flow.json --outputDirectory /etc/output/
```

Processes the flow.json file in all stages and saves the results in /etc/output/.

3Migrate variables using full customization with all arguments

```
./bin/migration.sh nifi migrate-variables --input /etc/exported-flows/flow.json --outputDirectory /etc/output/ --stage STAGE_1 --processGroupId bb81df68-cd6a-461a-b724-384265875b53
```

Runs Stage 1 of the migration, starting the processing from the process group ID bb81df68-cd6a-461a-b724-384265875b53. Saves the migrated flow and the Activity Log in /etc/output/.

migrate-components

Migrates components from NiFi 1 to NiFi 2 using the Cloudera Flow Management Migration Tool.

Syntax

```
migrate-components --input {filePath} --outputDirectory {directoryPath} [--stage {stageName}] [--processGroupId {pgId}]
```

Required arguments

--input

- Short format: -i {filePath}
- Specifies the path to a local file containing the flow in JSON format for processing.
- The file must be accessible to the Migration Tool.
- The original file remains unmodified.

--outputDirectory

- Short format: -od {directoryPath}
- Specifies the directory where migration results will be saved.
- The directory must be accessible to the Migration Tool.
- Existing content in the specified directory may be overwritten by subsequent runs.

Supported arguments

--stage

- Short format: -s {stageName}
- Specifies the stage(s) to process.
- To process Stage 1 migration, use `STAGE_1`
- If omitted, both Stage 1 and Stage 2 are processed sequentially, using Stage 1's output as input for Stage 2.
- Providing unsupported or custom stage names may result in invalid stage error.

--processGroupId

- Short format: -pgid {pgId}
- Defines the starting point for migration, applying transformations to the specified group and its child groups.
- If omitted, the root group is processed.
- Management-level components, such as Controller Services, Parameter Providers, and Reporting Tasks, are only included when this argument is not used.



Note:

The input flow.json file must represent the entire flow, regardless of whether this argument is used.

Output files of Stage 1 and Stage 2

Flow file

activity_log.json

A log file recording the details of the migration process. For more information, see *Activity Log*.

Examples

1 Migrate components processing Stage 1

```
./bin/migration.sh nifi migrate-components --input /etc/flow_original.json --outputDirectory /etc/output/ --stage STAGE_1
```


Processes the flow_original.json file located in /etc/. Applies transformation rules from Stage 1 on all components. Saves the migrated flow file and the Activity Log in /etc/output/.

2Migrate components in a specific Process Group

```
./bin/migration.sh nifi migrate-components --input /etc/flow_original.json --outputDirectory /etc/output/ --processGroupId 3f8d2cba-4d3b-4901-bd0b-4781ff5b5c9f
```

Starts the migration at the process group identified by 3f8d2cba-4d3b-4901-bd0b-4781ff5b5c9f. Transforms components within this group and its child groups. Outputs are saved in /etc/output/.

3Migrate components using full customization with all arguments

```
./bin/migration.sh nifi migrate-components --input /etc/flow_original.json --outputDirectory /etc/output/ --stage STAGE_1 --processGroupId 3f8d2cba-4d3b-4901-bd0b-4781ff5b5c9f
```

Runs Stage 1 only and the migration begins at the specified process group. Management level components are not migrated. Saves the migrated flow and the Activity Log in /etc/output/.

migrate-all

Aggregates the commands of all migration steps when using the Cloudera Flow Management Migration Tool.

Usage

- Combines all migration steps into a single command.
- Outputs of each step are used as inputs for the next step.
- The order of the included steps is the following:
 - Stage 1: Migrate Templates # Migrate Variables # Migrate Components
 - Stage 2: Migrate Templates # Migrate Variables # Migrate Components

help

Displays a list of available commands when using the Cloudera Flow Management Migration Tool.

Usage

- If you use it with a specific command (`[$[***COMMAND_NAME***] help`), it displays the available arguments for it.
- Specify the full command, including its group: `./migration.sh nifi migrate-all help`

Activity Log

This section contains information about activity_log.json. This log file provides a comprehensive record of the migration process, capturing detailed information about changes, events, and required interventions.

Summary

- File name: activity_log.json
- Serves a business-level record of changes and the context of the migration process.
- Provides a detailed log of the activities during migration in JSON format that follows a structured schema: it begins with metadata and consists of a list of entries.
- Designed for human readability and automated filtering, but not for fully automatic processing.

Metadata

- Provides context for understanding the log entries.
- Metadata fields:
 - recordingStarts: (number) UNIX epoch timestamp indicating when logging began, the approximate start of the stage run (not suitable for performance measurements)
 - recordingEnds: (number) UNIX epoch timestamp indicating when logging concluded, the approximate end of the stage run (not suitable for performance measurements)
 - stage: (string) Identifies the stage that is run. Possible values are "STAGE_1" or "STAGE_2"

Entry

- The Activity Log contains a list of entries, each representing a distinct event using different entry types.
- Each entry serves to provide insight into the actions performed during the migration process.

Entry properties

- sequence: A unique number within the log file, that represents the entry's sequence. It can be used to order and reference entries.
- type: A string denoting the purpose of the entry. For more information, see Entry types below.
- subject: A string representing the entity involved in the entry, such as a component, group, or other relevant entity.



Note: In the case of component migration, the entry's subject refers to the triggering component. This is important when a component is replaced, as the log will reference the original component's ID (which may have been deleted at this point) rather than the new component's ID.

- message: A string containing the main content or description of the entry.
- context: A JSON object with additional key-value pairs, providing further details about the entry (for example: applied rules, affected properties).

Entry types

- manual-change-request
 - Subject: Usually the ID of a component requiring manual intervention
 - Message: Describes changes that cannot be applied automatically and require human intervention. For example, deprecated components without replacements or decisions requiring domain expertise.
 - Context: Not commonly used
 - Example:

```
{
  "sequence": 5,
  "type": "manual-change-request",
  "subject": "69c4f8f3-549a-4103-b8f6-f8627b689e12",
  "message": "The value [null] for property [assume-role-sts-region] is not supported any more.",
  "context": { }
}
```

- manual-inspection-request
 - Subject: Typically the ID of a component under review.
 - Message: Highlights changes that should be explicitly reviewed to ensure correctness. These changes may or may not require adjustments.
 - Context: May include a rule key with a UUID for the applied migration rule, which helps identify errors.

- change-info
 - Subject: The component being migrated.
 - Message: Describes version differences and provides context for changes, without indicating any actual transformation.
 - Context: Includes a rule key with a UUID for the applied migration rule.
 - Example:

```
{
  "sequence": 4,
  "type": "change-info",
  "subject": "69c4f8f3-549a-4103-b8f6-f8627b689e12",
  "message": "Property [assume-role-sts-region] has been changed as follows: [Allowed values have been changed. Values removed: [eu-isoe-west-1]]",
  "context": {
    "rule": "6ca65082-09bd-4713-9fde-f5beb0722d1c"
  }
}
```

- change
 - Subject: The entity (for example: a component or template) that has been modified during migration.
 - Message: Records actual changes applied to the flow during migration. These entries provide specifics on transformations in response to version differences.
 - Context: Includes a rule key with a UUID and may contain additional transformation details. While change-info describes general version differences in NiFi (for example: Processor ConsumeKafka_1_0 has been removed from NiFi), change entry details the specific transformations applied to your flow to address those changes (For example: Deprecated processor ea9178e3-ad81-4c6f-b23c-88a01e08ddac has been replaced with processor d9730bae-4f88-430e-801d-8805f6069988).
 - Example:

```
{
  "sequence": 39,
  "type": "change",
  "subject": "0a8d5eb6-6791-1faa-2303-b0adcf300df1",
  "message": "Property [Kerberos Principal] has been removed",
  "context": {
    "rule": "65d775f9-dc70-45c9-b4d4-e18a21f8ccba"
  }
}
```

- control
 - Subject: Unspecified, often a Process Group
 - Message: Marks the progression of migration activities, such as the start or completion of a Process Group's migration.
 - Context: No widespread usage
 - Example:

```
{
  "sequence": 4,
  "type": "control",
  "subject": "0a8d5a99-6791-1faa-78ea-9582a48a4113",
  "message": "Group migration has started",
  "context": {
  }
}
```



Troubleshooting flow migration issues

Some operation systems may encounter issues such as “too many open files”. To resolve this, adjustments must be made at the operating system level.