

Upgrading Manually Using Configuration Restore

- Overview, on page 1
- Prerequisites and Guidelines, on page 3
- Validate Existing Configuration and Create Backup, on page 5
- Deploy Nexus Dashboard and Install Nexus Dashboard Orchestrator, on page 9
- Restore Configuration, on page 11
- Resolve Configuration Drifts, on page 15

Overview

There are two approaches when it comes to upgrading your Nexus Dashboard Orchestrator:

- Upgrading in-place by upgrading each component (such as the Nexus Dashboard platform and the Orchestrator service) in sequence.

  This approach is described in Upgrading Automatically Via Service Catalog and is recommended in the following cases:

  - If you are using a physical Nexus Dashboard cluster.
  - If you are running a recent release of Nexus Dashboard (2.2.2 or later) and Nexus Dashboard Orchestrator (3.7.1 or later).

    While you can use this approach to upgrade any Orchestrator release 3.3(1) or later, it may require upgrading the underlying Nexus Dashboard platform before you can upgrade the Orchestrator service. In those cases, an upgrade via configuration restore described below may be faster and simpler.

- Deploy a brand new Nexus Dashboard cluster, installing a new NDO service instance in it and transferring existing Orchestrator configuration via the configuration restore workflow

  This approach is described in this chapter and is recommended in the following cases:

  - If you are running any release of Nexus Dashboard Orchestrator or Multi-Site Orchestrator prior to release 3.3(1).

    In this case you must upgrade using configuration restore because in-place upgrade is not supported.

  - If you are using a virtual Nexus Dashboard cluster and running an older release of Nexus Dashboard Orchestrator.

Upgrading from an old Nexus Dashboard Orchestrator release requires upgrading the underlying Nexus Dashboard platform as well, in which case deploying a new cluster and restoring configuration may shorten the required maintenance window.

This also allows you to simply disconnect the existing cluster and keep the existing VMs until the upgrade is complete in case you want to revert to the previous version or the upgrade does not succeed.

Changes in Release 4.0(1) and Later

Beginning with Release 4.0(1), Nexus Dashboard Orchestrator will validate and enforce a number of best practices when it comes to template design and deployment:

- All policy objects must be **deployed** in order according to their dependencies.

  For example, when creating a bridge domain (BD), you must associate it with a VRF. In this case, the BD has a VRF dependency so the VRF must be deployed to the fabric before or together with the BD. If these two objects are defined in the same template, then the Orchestrator will ensure that during deployment, the VRF is created first and associate it with the bridge domain.

  However, if you define these two objects in separate templates and attempt to deploy the template with the BD first, the Orchestrator will return a validation error as the associated VRF is not yet deployed. In this case you must deploy the VRF template first, followed by the BD template.

- All policy objects must be **undeployed** in order according to their dependencies, or in other words in the opposite order in which they were deployed.

  As a corollary to the point above, when you undeploy templates, you must not undeploy objects on which other objects depend. For example, you cannot undeploy a VRF before undeploying the BD with which the VRF is associated.

- No cyclical dependencies are allowed across multiple templates.

  Consider a case of a VRF (`vrf1`) associated with a bridge domain (`bd1`), which is in turn associated with an EPG (`epg1`). If you create `vrf1` in `template1` and deploy that template, then create `bd1` in `template2` and deploy that template, there will be no validation errors since the objects are deployed in correct order. However, if you then attempt to create `epg1` in `template1`, it would create a circular dependency between the two template, so the Orchestrator will not allow you to save `template1` addition of the EPG.

Due to these additional rules and requirements, an upgrade to release 4.0(1) or later from an earlier release requires an analysis of all existing templates and conversion of any template that does not satisfy the new requirements. This is done automatically during the upgrade process described in the following sections and you will receive a detailed report of all the changes that had to be applied to your existing templates to make them compliant with the new best practices.



Note You must ensure that you complete all the requirements described in the following "Prerequisites and Guidelines" section before you back up your existing configuration for the upgrade. Failure to do so may result in template conversion to fail for one or more templates and require you to manually resolve the issues or restart the migration process.

Upgrade Workflow

The following list provides a high level overview of the migration process and the order of tasks you will need to perform.

1. Review the upgrade guidelines and complete all prerequisites.
2. Validate existing configuration using a Cisco-provided validation script, then create a backup of the existing Nexus Dashboard Orchestrator configuration and download the backup to your local machine.
3. Disconnect or bring down your existing cluster.

   If your existing cluster is virtual, you can simply disconnect it from the network until you've deployed a new cluster and restored the configuration backup in it. This allows you to preserve your existing cluster and easily bring it back in service in case of any issues with the migration procedure.
4. Deploy a brand new Nexus Dashboard cluster release 2.3(b) or later and install Nexus Dashboard Orchestrator release 4.1(2) or later.
5. Add a remote location for backups to the fresh Nexus Dashboard Orchestrator instance, upload the backup you took on your previous release, and restore the configuration backup in the new NDO installation.
6. Resolve any configuration drifts.

Prerequisites and Guidelines

Before you upgrade your Cisco Nexus Dashboard Orchestrator cluster:

- Note that downgrading from this release is not supported.

  If you ever want to downgrade, you can deploy a brand-new cluster using the earlier version and then restore configuration from the earlier release. Note that you cannot restore a backup created on a newer version in an older version, in other words restoring a backup from release 4.2(1) in release 3.7(1) is not supported.
- The backup/restore upgrade workflow supports upgrades from any Multi-Site Orchestrator (MSO) release 2.x and 3.x as well as Nexus Dashboard Orchestrator (NDO) release 3.x and 4.x to this release of NDO.
- Ensure that there are no configuration drifts before you back up your existing configuration.

  This applies to all template types available in your existing release, such as application, tenant policies, fabric policies, and fabric resource policies templates.

  If your existing Nexus Dashboard Orchestrator is release 3.7(1) or later, you can use the drift reconciliation workflow for application templates, as described in the "Configuration Drifts" section of the *Nexus Dashboard Orchestrator Configuration Guide*.
- Back up and download your existing Orchestrator configurations.

  Configuration backups are described in the "Backup and Restore" chapter of the *Nexus Dashboard Orchestrator Configuration Guide* for your release.
- Back up and download your existing fabrics' configurations.

  We recommend running configuration drift reconciliation after you upgrade your Nexus Dashboard Orchestrator, which may require you to redeploy configurations to your fabrics. As such, we recommend creating configuration backups of all fabrics managed by your Nexus Dashboard Orchestrator.

For more information on creating Cisco APIC configuration backups, see the "Management" chapter of the *Cisco APIC Basic Configuration Guide* for your release.

For more information on creating Cisco Cloud Network Controller configuration backups, see the "Configuring Cisco Cloud Network Controller Components" chapter of the *Cisco Cloud Network Controller User Guide* for your release.

For more information on creating Cisco Nexus Dashboard Fabric Controller configuration backups, see the "Backup and Restore" chapter of the *Cisco NDFC Fabric Controller Configuration Guide* for your release.

- Note that if versioning enabled (supported since release 3.4(1)), only the latest versions of the templates are preserved during the upgrade.

  All other existing versions of templates, including older versions that are tagged `Golden`, will not be transferred.

- Ensure that all templates are in a supported state before creating the configuration backup of the existing cluster:
  - Templates that are **undeployed** or were **never deployed** after they were created require no additional attention and will be migrated during the upgrade.
  - All **deployed** templates must have no pending configuration changes.

    If you have one or more templates that have been modified since they were last deployed, you must either deploy the latest version of the template or undo the changes to the template since it was deployed by reverting to the last-deployed version and re-deploying it.

Note Attempting to restore a backup that contains invalid templates will fail and you would need to revert to your existing release, restore your backup, resolve any existing issues, and then restart the migration process. So we strongly recommend that you validate your backup locally using the provided Python script before proceeding with the upgrade, as described in the Validate Existing Configuration and Create Backup, on page 5 section below. If for any reason you are unable to run the script, we recommend contacting Cisco support to have them validate your configuration backup before proceeding with the upgrade.

- When installing the Orchestrator service, you can do so in one of two ways:
  - Using the Nexus Dashboard's App Store, in which case, the Cisco DC App Center must be reachable from the Nexus Dashboard via the Management Network directly or using a proxy configuration. Nexus Dashboard proxy configuration is described in the *Nexus Dashboard User Guide*.

Note The App Store allows you to upgrade to the latest available version of the service only. If you want to upgrade to a different release, you must use the manual upgrade process as described below.

  - By manually uploading the new app image, which you can do if you are unable to establish the connection to the DC App Center or if you want to upgrade to a version of the service that is not the latest available release.

- SR-MPLS and SDA integration configurations are not transferred during the upgrade.

  If you have either of these integrations in your deployment, it will not affect the migration, but you will receive a notification and will need to reconfigure them after you complete the upgrade.

- If you plan to add and manage new Cloud Network Controller sites after you upgrade your Nexus Dashboard Orchestrator to this release, ensure that they are running Cloud Network Controller release 5.2(1) or later.

  On-boarding and managing Cloud Network Controller sites running earlier releases is not supported.

- Ensure that you have a remote location for backups that you can add to the Nexus Dashboard Orchestrator after the upgrade.

  Backing up and restoring configuration in release 4.1(2) and later requires the backup to be stored on a remote location, which must be configured in NDO UI. Detailed information about backups and remote locations is available in the **Operations** > **Backup and Restore** chapter of the *Cisco Nexus Dashboard Orchestrator Configuration Guide*.

  Note that if you had a remote location already configured in your existing installation, it is not preserved during configuration restore; so you will need to add the same remote location after you deploy this release in order to restore the configuration.

Validate Existing Configuration and Create Backup

This section describes how to create a backup of the existing configuration, which you will then restore after you re-deploy a fresh instance of Nexus Dashboard Orchestrator.

Before you begin

You must have the following completed:

- Familiarized yourself with the migration workflow described in the Overview, on page 1
- Reviewed and completed the prerequisites described in Prerequisites and Guidelines.

Procedure

Step 1 Download and verify the configuration validation script.

Note
If you are upgrading from release 4.0(1) or later, you can skip this step.

You will use this script to validate your existing configuration before creating a backup and upgrading the Orchestrator service to this release.

a) Ensure that you have Python installed on your local machine.

The script requires Python 3 to run. You can check if Python is installed on your machine using the following command:

```
$ python3 --version
Python 3.9.6
```

b) Download and extract the validation script tarball.

Navigate to https://software.cisco.com/download/home/285968390/type/286317465, select the target NDO version to which you want to upgrade, download the upgrade validation script (`Final_ndo<version>-UpgradeValidationScript.tgz`), then extract it, for example:

```
$ tar -xzf Final_ndo<version>-UpgradeValidationScript.tgz
$ ls
ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM
Final_ndo4.1.2h-UpgradeValidationScript.tgz
UpgradeValidationScript.tgz
UpgradeValidationScript.tgz.signature
cisco_x509_verify_release.py3
```

c) Verify the validation script tarball signature.

You can use the following command to verify the Cisco signature on the configuration validation script.

```
$ ./cisco_x509_verify_release.py3 -e ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM \
  -i UpgradeValidationScript.tgz -s UpgradeValidationScript.tgz.signature -v dgst -sha512
Retrieving CA certificate from https://www.cisco.com/security/pki/certs/crcam2.cer ...
Successfully retrieved and verified crcam2.cer.
Retrieving SubCA certificate from https://www.cisco.com/security/pki/certs/innerspace.cer ...
Successfully retrieved and verified innerspace.cer.
Successfully verified root, subca and end-entity certificate chain.
Successfully fetched a public key from ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM.
Successfully verified the signature of UpgradeValidationScript.tgz using
ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM
```

Note
If signature verification fails, you will receive the following error:

```
$ ./cisco_x509_verify_release.py3 -e ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM\
-i UpgradeValidationScript.tgz -s UpgradeValidationScript.tgz.signature.fail -v dgst -sha512
Retrieving CA certificate from https://www.cisco.com/security/pki/certs/crcam2.cer ...
Successfully retrieved and verified crcam2.cer.
Retrieving SubCA certificate from https://www.cisco.com/security/pki/certs/innerspace.cer ...
Successfully retrieved and verified innerspace.cer.
Successfully verified root, subca and end-entity certificate chain.
Successfully fetched a public key from ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM.
Error log: Failed to verify dgst signature of UpgradeValidationScript.tgz.
Error log: Verification Failure
```

In this case, we recommend you re-download the `<ndo-version>-UpgradeValidationScript.tgz` tarball from the Cisco Software Download portal.

d) Once the validation script signature is verified, extract the script itself.

```
% tar -xzf UpgradeValidationScript.tgz
$ ls
ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM
Final_ndo4.1.2h-UpgradeValidationScript.tgz
README.md
UpgradeValidationScript.tgz
UpgradeValidationScript.tgz.signature
cisco_x509_verify_release.py3
ndo
ndoCmd.py
ndoCopy.py
requirements.txt
```

Step 2 Validate your existing configuration before you create the backup.

Note

If you are upgrading from release 4.0(1) or later, you can skip this step.

You can verify that your configuration backup will be compatible with upgrade to this release by running the validation script you have downloaded in the previous step. If for any reason you are unable to run the script, we recommend contacting Cisco support to have them validate your configuration backup before proceeding with the upgrade.

a) Log in to your Nexus Dashboard and open the Nexus Dashboard Orchestrator service.
b) Download the tech support logs from your existing Orchestrator.

   While for the migration you will create and download the configuration backup using the standard procedure, the validation is done on the tech support information. Note that it is normal for the tech support archive to be significantly larger than your typical configuration backup.

   You can generate the tech support logs by navigating to **Admin** > **Tech Support** page in the Orchestrator UI. Then click the **Download** icon in the **System Logs** tile. This downloads the `msc_report_<date>.zip` archive to your machine.

c) Extract the tech support archive you downloaded.

   The tech support archive comes in a standard `.zip` format, so you can use any tool of your choice to extract the contents, for example:

   ```
   $ unzip msc_report_<date>.zip
   ```

   After you extract the archive, copy the `msc-db-json-<date>_temp.tar.gz` file inside into the directory where you extracted the validation script.

d) Run the validation script.

   The script requires a number of dependencies, which are all defined in the `requirements.txt` file that comes with the script, so we recommend creating a Python virtual environment before installing the dependencies and running the script:

   ```
   $ python -m venv ndo-upgrade
   $ source ndo-upgrade/bin/activate
   $ pip install -r requirements.txt
   ```

   After the virtual environment is set up and the required modules are installed, run the script using the tech support file you downloaded and extracted in a previous step, for example:

   - `-f` allows you to provide the file on which to run the validation.
   - `-N` specifies that no configuration will be deployed to any live system.
   - `-C` generates the JSON-formatted output at the end of the script.

```
(ndo-upgrade)ndoCmd $ ./ndoCopy.py -f
msc_report_20220617_181529/msc-db-json-20220617181553_temp.tar.gz -N -C
11:49:56 Loading collection site2...4
11:49:56 Loading collection tenant...12
[...]
11:49:56 Checking template versions
11:49:56 Checking policy deployment dependencies
11:49:56 Fixing template policy flow loops
11:49:56 Fixing template dependency loops
11:49:56 Fixing policies for upgrade
11:49:56 Determine template ordering
11:49:56 Analysis completed
{
  "summaryStats": {
    "appTemplatePoliciesConverted": 139,
    "appTemplateSiteAssocMods": 7,
```

```
    "appTemplatePolicyEvictions": 2,
    "appTemplateSchemasConverted": 11,
    "appTemplatesConverted": 38,
    "appTemplatesCreated": 1,
    "tenantMods": 1
  },
[...]
}
```

After the output is generated:

- If there are no `errors` or `warnings` blocks at the end of the generated JSON, then your configuration is compliant with the migration requirements and you can proceed to the "Back up existing deployment configuration" step.
- If there is only a number of warnings but no errors, it means the migration will complete successfully, but there's a number of things that you may want to resolve before or after the upgrade. We recommend reviewing any warnings before continuing with the next step.

```
  "warnings": [
    "dropped DHCP Relay policy dhcp-tn-epgOnRL-policy: invalid provider ip address:
141.1.141.2/24",
    "dropped Route Map policy sameContract: fromPrefixLen and toPrefixLen must be larger than
 prefix",
    "dropped Mulicast Route Map policy mCastRt.map: invalid RP ip: 12.13.14.15/23",
    "dropped DHCP Option policy dhcpBdMso-option: duplicate option id: 1; duplicate option
id: 1",
    "removed dhcpLabels.0 from bd[tn-epgOnRL::Template 1::bdDhcpClient] for
unresolved policy ref key[dhcpRelayPolicies::tn-epgOnRL::dhcp-tn-epgOnRL-policy]",
    "removed dhcpLabels.0 from bd[dhcp-msite-mso::bd::bd-client-l3out] for
unresolved policy ref
key[dhcpRelayPolicies::dhcp-msite-mso::dhcp-msite-mso-relay-policy-epg-cleint-l3out]"
  ],
```

- If there is 1 or more errors listed in the JSON, the migration would fail if you continue with the current configuration.

  Note
  You must resolve any existing errors before creating the backup and proceeding with the upgrade. We recommend re-running the validation script after you resolve any existing errors to ensure that the backup will be ready for the migration.

  For example, the following sample shows 2 possible errors that can come up during validation:

```
"errors": [
  "template appTemplate[<template>] version 6 is in state EDIT_CONFIG",
  "deployed policy bd[<bd>] requires vrf[<vrf>] which is not deployed",
]
```

  - As mentioned in the Prerequisites and Guidelines section, any deployed templates must not have undeployed changes. You must either deploy the latest version of that template or revert to the deployed version (so it is the latest version) and re-deploy the template.
  - Objects must be deployed in order of their dependencies. In other words, you must not have a deployed bridge domain if the required VRF is not deployed.

e) Resolve any shown errors and repeat this step to re-validate the configuration.

Step 3 Back up existing deployment configuration.

a) From the left navigation pane, select **Operations** > **Backups & Restore**.
b) In the main window, click **New Backup**.

A **New Backup** window opens.

c) In the **Name** field, provide the name for the backup file.

The name can contain up to 10 alphanumeric characters, but no spaces or underscores (_).

d) From the **Remote Location** dropdown, choose the remote location you have configured previously.
e) In the **Remote Path** field, provide the path on the remote server where to store the backup.
f) Click **Save** to create the backup.

Step 4 Download the backup file.

In the main window, click the actions ( ) icon next to the backup and select **Download**. This will download the backup file to your system.

Deploy Nexus Dashboard and Install Nexus Dashboard Orchestrator

Because you are deploying a fresh cluster, the steps are identical to the ones described in the Deploying Nexus Dashboard Orchestrator chapter of this guide. This section summarizes the steps and provides specific links for each one.

Before you begin

You must have the following completed:

- Backed up and downloaded the existing Nexus Dashboard Orchestrator configuration.

Procedure

Step 1 Deploy a new Nexus Dashboard cluster.

Detailed information about deployment requirements, available form factors, and installation instructions are available from the *Cisco Nexus Dashboard Deployment Guide*.

Note
You must deploy in Nexus Dashboard release 3.0.1 or later.

Step 2 After Nexus Dashboard and Nexus Dashboard Orchestrator have been successfully deployed, navigate to your Nexus Dashboard's **Admin Console**.

You can log in to Nexus Dashboard by opening your browser and navigating to the management IP address of any one of the Nexus Dashboard cluster's nodes, then selecting **Admin Console** from the drop down menu in the top navigation bar.

Step 3 Onboard the fabrics managed by the Orchestrator service to Nexus Dashboard.

Note

You must on-board all the fabrics previously managed by your Nexus Dashboard Orchestrator (or Multi-Site Orchestrator) to Nexus Dashboard before you can proceed with restoring configuration and completing the upgrade process.

In this release, fabric on-boarding is done in the common Nexus Dashboard screen. The process is described in detail in the "Site Management" chapter of the *Cisco Nexus Dashboard User Guide*, but in short:

a) From the main navigation menu, choose **Sites**.
b) In the main pane, click **Add Site**.
c) Choose the type of site you want to on-board and provide the site's information, such as controller's IP address, username, and password.

- For on-premises sites managed by Cisco APIC, choose `ACI`.
- For cloud sites managed by Cloud Network Controller (previously Cloud APIC), choose `Cloud Network Controller`.

  Use this option for all Cloud APIC sites that were managed by your Orchestrator.
- For on-premises sites managed by NDFC (previously DCNM), choose `NDFC`.

  Use this option for all DCNM sites that were managed by your Orchestrator.

Note
You must use the same site **Name** as you did when on-boarding the site to your Orchestrator in the past. Adding a site with a different name will cause configuration restore to fail.

d) Click **Save** to add the site.
e) Wait for site to come up and show as `Up` in the Nexus Dashboard UI.
f) Repeat this step for all sites that were previously managed by your Orchestrator.

Step 4 Install Nexus Dashboard Orchestrator service.

This is described in detail in the Deploying Nexus Dashboard Orchestrator chapter of this guide, but if you are already familiar with Orchestrator installation, the following steps summarize the process.

If you are installing the service using the App Store:

a) In the **Services** screen, select the **App Store** tab.
b) In the **Nexus Dashboard Orchestrator** tile, click **Upgrade**.
c) In the License Agreement window that opens, click **Agree and Download**.

If you are installing the service manually:

a) Browse to the Nexus Dashboard Orchestrator page on DC App Center:
https://dcappcenter.cisco.com/nexus-dashboard-orchestrator.html
b) From the **Version** dropdown, choose the version you want to install and click **Download**.
c) Click **Agree and download** to accept the license agreement and download the image.
d) From the left navigation menu in Nexus Dashboard, select **Services**.
e) In the Nexus Dashboard's **Services** screen, select the **Installed Services** tab.
f) From the **Actions** menu in the top right of the main pane, select **Upload Service**.
g) In the **Upload Service** window, choose the location of the image

If you downloaded the application image to your system, choose **Local**.

If you are hosting the image on a server, choose **Remote**.

h) Choose the file.

If you chose **Local** in the previous substep, click **Select File** and select the app image you downloaded.

If you chose **Remote**, provide the full URL to the image file, for example `http://<ip-address>:<port>/<full-path>/cisco-mso-<version>.nap`.

i) Click **Upload** to add the app to the cluster.

It may take a few minutes for the image to be uploaded to the cluster and initialized.

Step 5 Wait for the new image to initialize.

Step 6 In the Nexus Dashboard Orchestrator tile, click **Enable**.

It may take a few minutes for all the application services to start and the GUI to become available.

Step 7 Launch the Orchestrator service.

Simply click **Open** on the service tile.

The single sign-on (SSO) feature allows you to log in to the application using the same credentials as you used for the Nexus Dashboard.

What to do next

After Nexus Dashboard and Nexus Dashboard Orchestrator are installed and all sites are on-boarded, proceed to restore the configuration as described in Restore Configuration, on page 11.

Restore Configuration

This section describes how to deploy and configure the new Nexus Dashboard cluster and the NDO service, which you will use to restore your previous configuration.

Before you begin

You must have the following completed:

- Backed up and downloaded the existing Nexus Dashboard Orchestrator configuration.
- Installed the target Orchestrator release as described in Deploy Nexus Dashboard and Install Nexus Dashboard Orchestrator, on page 9.

Procedure

Step 1 Ensure that the new Nexus dashboard cluster is up and running and the NDO service is installed.

The NDO service must be a fresh install with no configuration changes to the sites or policies.

Step 2 Open your new Nexus Dashboard Orchestrator service.

Step 3 Add remote location for configuration backups.

This release of Nexus Dashboard Orchestrator does not support configuration backups stored on the cluster's local disk. So before you can import the backup you saved before the migration, you need to configure a remote location in Nexus Dashboard Orchestrator to which you can then import your configuration backups.

a) From the left navigation pane, select **Admin** > **Backup & Restore**.
b) Choose the **Remote Locations** tab.
c) Choose **Create Remote Location**.

The **Create Remote Location** screen appears.

d) Provide the name for the remote location and an optional description.

Two protocols are currently supported for remote export of configuration backups:

- `SCP`
- `SFTP`

Note
`SCP` is supported for non-Windows servers only. If your remote location is a Windows server, you must use the `SFTP` protocol

e) Specify the host name or IP address of the remote server.

Based on your **Protocol** selection, the server you specify must allow `SCP` or `SFTP` connections.

f) Provide the full path to a directory on the remote server where you will save the backups.

The path must start with a slash (`/`) characters and must not contain periods (`.`) or backslashes (`\`). For example, */backups/ndo*.

Note
The directory must already exist on the remote server.

g) Specify the port used to connect to the remote server.

By default, port is set to `22`.

h) Specify the authentication type used when connecting to the remote server.

You can configure one of the following two authentication methods:

- `Password`—provide the username and password used to log in to the remote server.
- `SSH Private Files`—provide the username and the SSH Key/Passphrase pair used to log in to the remote server.

i) Click **Save** to add the remote server.

Step 4 Import the backup file to your new Nexus Dashboard Orchestrator cluster.

a) From the left navigation pane, select **Operations** > **Backups & Restore**.
b) In the main pane, click **Upload**.
c) In the **Upload to Remote** window that opens, click **Select File** and choose the configuration backup file you created before the upgrade.
d) From the **Remote Location** dropdown menu, select the remote location.
e) (Optional) Update the remote location path.

The target directory on the remote server, which you configured when creating the remote backup location, will be displayed in the **Remote Path** field.

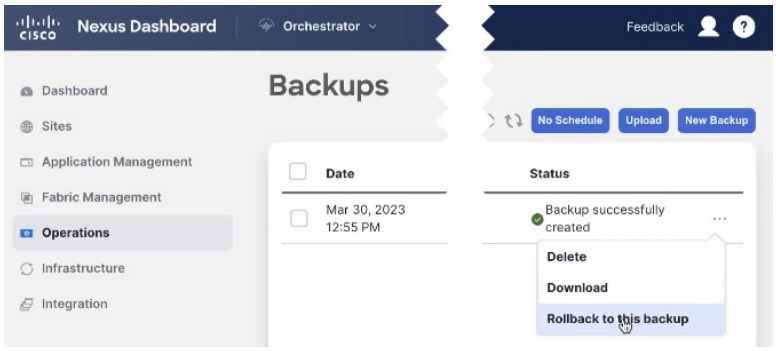
You can choose to append additional subdirectories to the path. However, the directories must be under the default configured path and must have been already created on the remote server.

f) Click **Import** to upload the file.

Importing a backup will add it to the list of the backups displayed the **Backups** page. Note that even though the backups are shown on the NDO UI, the files are stored only on the remote server and not directly on the cluster nodes.

Step 5 Restore the configuration.

a) In the main window, click the actions (**…**) icon next to the backup you want to restore and select **Rollback to this backup**.



b) In the **Restore from this backup** dialog, read the warning and click **Restore** to confirm that you want to restore the backup you selected.

The restore process imports the backup and checks for any issues, which may take several minutes to complete. After the initial backup import, you will be prompted for additional validation in the next step, which is required for database upgrades from releases prior to release 4.0(1).

c) After the backup import is complete, ensure there are no failures listed in the report, then click **Restore Validation Required** to proceed.

Before the configuration database is updated for this release, the upgrade process performs a number of validations. The validation provides a summary of template and policy changes that will be performed during this final upgrade stage in the next step and includes the following:

- Implicit template stretching – if one or more objects are implicitly stretched, the upgrade process will create new explicitly-stretched templates and move the objects into those templates.

  For example, if you have a template (`t1`) that contains `vrf1` and is associated to `site1` and another template (`t2`) that contains a BD that references `vrf1` but is associated to two sites (`site1` and `site2`), then `vrf1` will be implicitly stretched between the two sites.

  This is no longer allowed starting with release 4.0(1) and the VRF must be explicitly stretched to both sites. In such cases during the upgrade, the VRF will be either moved to a different template which will be explicitly stretched between both sites or the original template will be associated with both sites, depending on whether the other policies in that template require stretching as well.

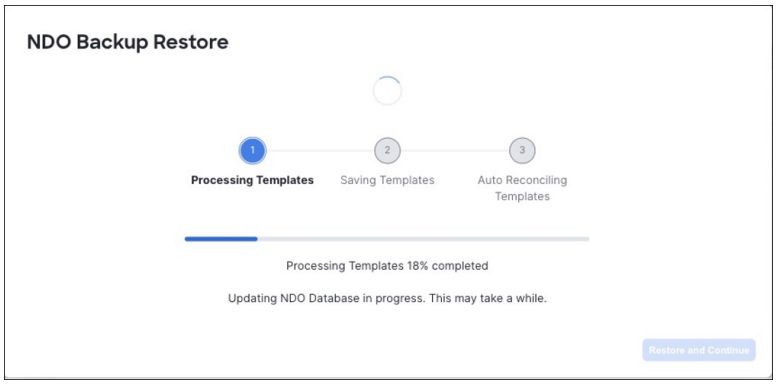
  Any templates that are created in this case will be named `UpgradeTemplate%d`, where `%d` is an incrementing number starting with 1 to ensure that all newly added templates are unique.

- Global policy migration – all global tenant policies (such as DHCP relay or route maps) and fabric policies (such as QoS) will be moved into the new tenant and fabric policy templates that have been added in release 4.0(1).

Note

At this stage, all the tenants have been imported from the backup and created in NDO, but the schemas and templates will be created in the next step.

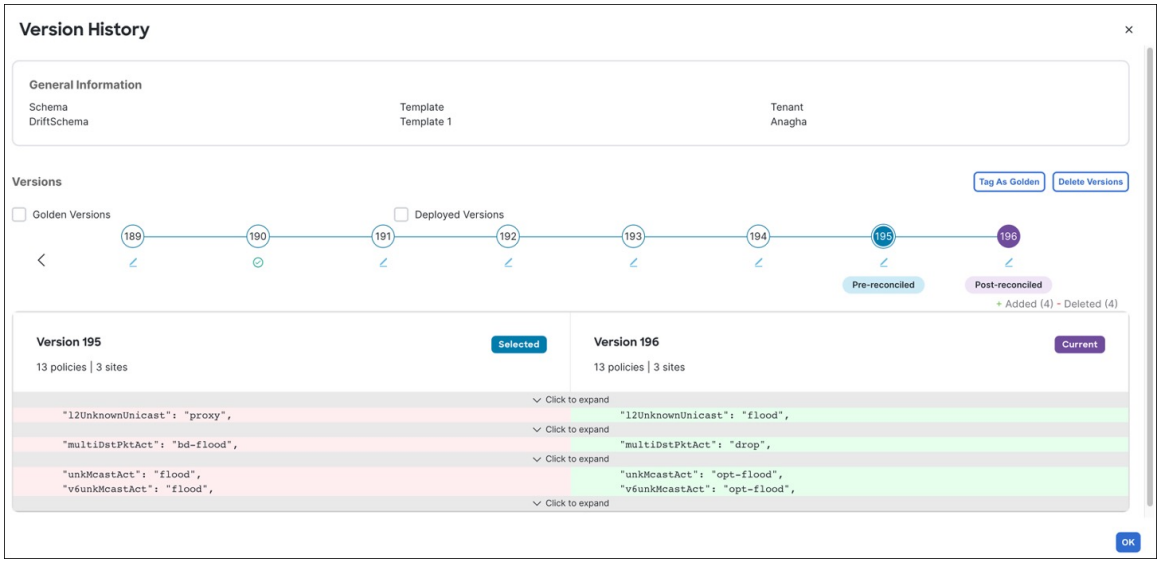
d) In the **Restore Validation Report** window, click **Restore and Continue** to proceed.



This is the stage of the upgrade where the schemas and the templates present in the backup are imported and recreated in your NDO configuration database according to the current best practices. These schemas and templates are then posted to the local NDO database as if it were a greenfield schema/template creation. Then the newly saved templates are deployed in the correct order that conforms to the current deployment requirements and best practices. The template deployment in this step uses a “local deploy” option to calculate the deployment plan and update the database, but does not send any configuration payload to the sites' controllers.

The upgrade process also checks for any configuration drifts between the local NDO database (configuration that is correct from NDO's point of view) and what is actually deployed in the fabrics. If this release of NDO supports additional objects or properties compared to the release from which you are upgrading, the upgrade will automatically reconcile those drifts by importing the existing configuration from the site's controller.

Note that if a template is automatically reconciled, two template versions are created – one before the automatic reconciliation and one after:



e) Review the report from the previous substep and click **Ok** to finish.

The final stage of the database upgrade presents a full report of the performed actions for you to review. If you close the report but want to review it again, simply click the **View Restore Report** in the **Backups** page.

Step 6 Verify that backup was restored successfully and all objects and configurations are present.

a) In the **Sites** page, verify that all sites are listed as `Managed`.

b) In the **Tenants** and **Schemas** pages, confirm that all tenants and schemas from your previous Nexus Dashboard Orchestrator cluster are present.

c) Navigate to **Infrastructure** > **Site Connectivity** and confirm that intersite connectivity is intact.

In the main pane, click **Show Connectivity Status** next to each site and verify that the existing tunnels are up and connectivity was not interrupted.

d) In the main pane, click **Configure** to open **Fabric Connectivity Infra** screen and verify **External Subnet Pool** addresses.

You can view the external subnet pools by selecting **General Settings** > **IPSec Tunnel Subnet Pools** tab of the **Fabric Connectivity Infra** screen and verify that the External Subnet Pools previously configured in Cloud Network Controller have been imported from the cloud sites.

These subnets are used to address the IPsec tunnel interfaces and loopbacks of the Cloud Routers used for on-premises connectivity and had to be configured directly in the Cloud Network Controller in earlier Nexus Dashboard Orchestrator releases.

Resolve Configuration Drifts

In some cases you may run into a situation where the configuration actually deployed in the site's controller is different from the configuration defined in the Nexus Dashboard Orchestrator. These configuration discrepancies are referred to as **Configuration Drifts** and are indicated by an `Out of Sync` warning next to the site name in the template view page

After upgrading your Nexus Dashboard Orchestrator and restoring the previous configuration backup, we recommend that you check for and resolve any configuration drifts that were not automatically resolved by the upgrade process as described in this section.



Note Deploying any templates before resolving configuration drifts would push the configuration defined in the Orchestrator and overwrite the values defined in the fabrics' controllers.

Procedure

Step 1 In your Nexus Dashboard Orchestrator, navigate to **Operate** > **Tenant Templates**.

Step 2 Choose the **Applications** tab.

Step 3 Select the first schema and check its templates for configuration drifts.

You will repeat the following steps for every schema and template in your deployment

You can check for configuration drifts in one of the following two ways:

- Check the template deployment status icon for each site to which the template is assigned.
- Select the template and click **Deploy to sites** to bring up the configuration comparison screen to check which objects contain configuration drifts.

Step 4 If the template contains a configuration drift, resolve the conflicts.

For more information about configuration drifts, check the "Configuration Drifts" chapter in the *Cisco Nexus Dashboard Orchestrator Configuration Guide for ACI Fabrics*.

a) Close the template deployment dialog to return to the Schema view.

Deploying any templates at this point would push the values in the Orchestrator database and overwrite any existing settings in the fabrics.

b) From the template's **Actions** menu, select **Drift Reconciliation**.

The **Drift Reconciliation** wizard opens.

c) In the **Drift Reconciliation** screen, compare the template-level configurations for each site and choose the one you want.

Template-level properties are common across all sites associated to the template. You can compare the template level properties defined on Nexus Dashboard Orchestrator with the configuration rendered in each site and decide what should become the new configuration in the Nexus Dashboard Orchestrator template. Selecting the site configuration will modify those properties in the existing Nexus Dashboard Orchestrator template, whereas selecting the Nexus Dashboard Orchestrator configuration will keep the existing Nexus Dashboard Orchestrator template settings as is

d) Click **Go to Site Specific Properties** to switch to site-level configuration.

You can choose a site to compare that specific site's configuration. Unlike template-level configurations, you can choose either the Nexus Dashboard Orchestrator-defined or actual existing configurations for each site individually to be retained as the template's site-local properties for that site.

Even though in most scenarios you will make the same choice for both template-level and site-level configuration, the drift reconciliation wizard allows you to choose the configuration defined in the site's controller at the "Template Properties" level and the configuration defined in Nexus Dashboard Orchestrator at the "Site Local Properties" level or vice versa.

e) Click **Preview Changes** to verify your choices.

The preview will display full template configuration adjusted based on the choices picked in the **Drift Reconciliation** wizard

f) Save the schema.

g) Click **Deploy to sites** to deploy the configuration and finish reconciling the drift for that template

Step 5 Repeat the above steps for every schema and template in your Nexus Dashboard Orchestrator.

Step 6 Check audit logs to verify that all templates have been re-deployed.

You can view the audit logs in the **Operations** tab.

Audit Logs page and confirm that all templates show as `Redeployed` to ensure that full re-deployment successfully completed.