

# Upgrading from 3.7(x) or Earlier Releases


- Overview, on page 1
- Prerequisites and Guidelines, on page 3
- Validate Existing Configuration and Create Backup, on page 5
- Uninstall Existing Nexus Dashboard Orchestrator, on page 9
- Upgrade Nexus Dashboard Cluster, on page 10
- Install Nexus Dashboard Orchestrator Release 4.0(x), on page 14
- Restore Configuration, on page 16
- Resolve Configuration Drifts, on page 20


## Overview

The following sections describe how to upgrade Cisco Nexus Dashboard Orchestrator that is deployed in Cisco Nexus Dashboard from release 3.2(x) or later to release 4.0(1) or later.

**Note** If you are already running 4.0(1) or later release, skip this section and follow the instructions described in Upgrading from Existing 4.0(x) Release instead.

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 migration 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 to be migrated to release 4.0(1). 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. Back up existing Nexus Dashboard Orchestrator configuration and download the backup to your local machine.
3. Disable and completely uninstall the Nexus Dashboard Orchestrator service from your Nexus Dashboard cluster.

   This is mandatory for physical Nexus Dashboard clusters because you will deploy release 4.0(1) on the same cluster.

   However, if your Nexus Dashboard cluster is virtual, you can choose to deploy a brand new cluster and install Nexus Dashboard Orchestrator release 4.0(x) in it. After the new cluster is up and running, you can disconnect the old cluster's VMs and complete the migration process on the new cluster. This allows you to preserve your existing cluster and easily bring it back in service in case of any issue with the migration procedure.
4. If necessary, upgrade the Nexus Dashboard cluster.

   Similarly to the previous step, if your cluster is virtual, you can choose to preserve it and deploy a new virtual cluster to complete the migration.
5. Install the target 4.0(x) release of Nexus Dashboard Orchestrator.
6. 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 service.

# Prerequisites and Guidelines

Before you upgrade your Cisco Nexus Dashboard Orchestrator cluster:

- Ensure that you are running Nexus Dashboard Orchestrator release 3.2(1) or later.

**Note** If you are already running a 4.0(x) or later release, skip this section and follow the instructions described in Upgrading from 3.7(x) or Earlier Releases, on page 1 instead.

If you are running a release prior to Release 3.2(1), you must migrate your Nexus Dashboard Orchestrator to Nexus Dashboard before upgrading to this release. We recommend that you migrate to Release 3.7(1), as described in *Migrating Existing Cluster to Nexus Dashboard*, then come back to this document to upgrade to release 4.0(x).

- Ensure that your current Nexus Dashboard cluster is healthy.

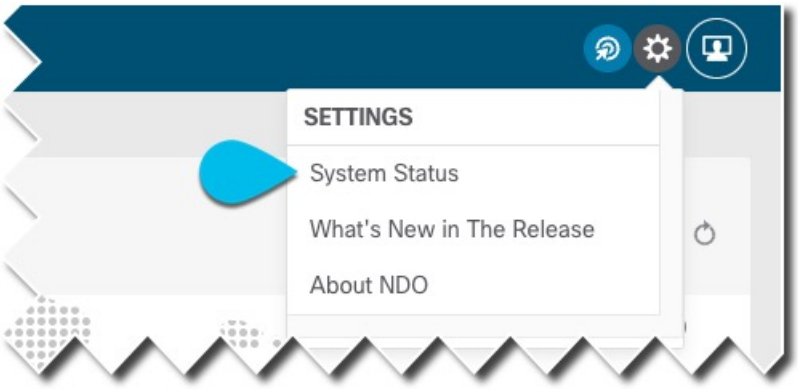
  You can check the Nexus Dashboard cluster health in one of two ways:

  - By logging into your Nexus Dashboard GUI and verifying system status in the **System Overview** page.
  - By logging into any one of the nodes directly as `rescue-user` and running the following command:

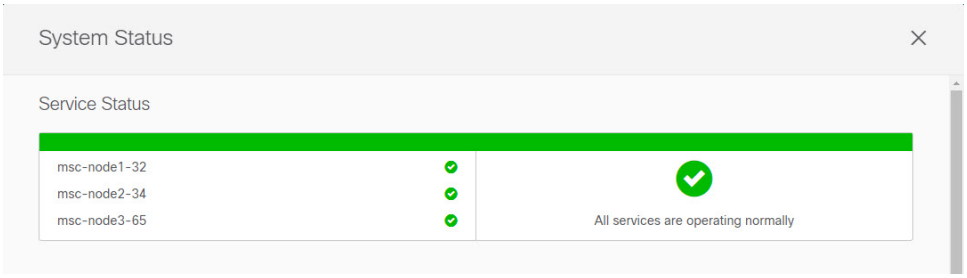
    ```
    # acs health
    All components are healthy
    ```

- Ensure that your current Cisco Nexus Dashboard Orchestrator is healthy.

  You can check the status of you Nexus Dashboard Orchestrator service by navigating to **Settings** > **System Status**:




  Then ensure that the status of all nodes and services is healthy:




- Ensure that there are no configuration drifts before you back up your existing configuration.

  Resolving configuration drifts is described in the "Configuration Drifts" section of the*Nexus Dashboard Orchestrator Configuration Guide*.

- Back up 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 for AWS 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.

- Ensure that you have an existing remote location configured for backups in your existing Orchestrator.

  The upgrade process requires configuration backup and restore, so you must have a remote location where you can save the backup already set up in your existing cluster.

  Note that the remote locations where you store configuration backups are not preserved during the upgrade process, so you will need to re-add the same remote location after you deploy this release in order to restore the configuration.

- 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.

- 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.

- 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.

- 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.

- Depending on your Nexus Dashboard cluster's form factor, choose the appropriate upgrade approach:

  - For physical Nexus Dashboard clusters, you will need to back up your configuration, remove the existing Orchestrator instance, deploy this release of the service, and then restore the configuration backup from your existing cluster.
  - For virtual Nexus Dashboard clusters, you can choose to follow the same workflow as the physical clusters. You also have the option to disconnect your existing cluster and preserve its VMs until the upgrade is complete while deploying a brand new virtual cluster with Orchestrator 4.0(1) and restoring configuration there.

  In either case, you can follow the instructions in the following sections which will call out any differences based on the chosen approach.

- Downgrading from this release is not supported.

  You will create a full backup of the configuration before upgrading to Release 4.0(1), so that if you ever want to downgrade, you can deploy a brand-new cluster using the earlier version and then restore your configuration in it.

# 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 upgrade the Nexus Dashboard Orchestrator service.

### Before you begin

You must have the following completed:

- Familiarized yourself with the migration workflow described in the Overview
- Reviewed and completed the prerequisites described in Prerequisites and Guidelines.
- Set up a remote location for configuration backups.

### Procedure

**Step 1** Log in to your Nexus Dashboard and open the Nexus Dashboard Orchestrator service.

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

You can verify that your configuration backup will be compatible with release 4.0(1) upgrade by running a local Python validation script. 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) 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.

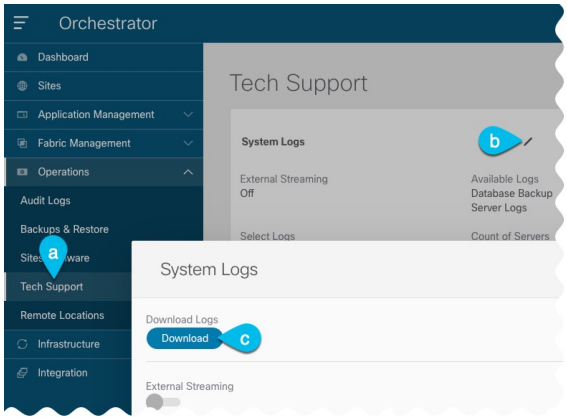
You can download the script tarball from https://software.cisco.com/download/home/285968390/type/286317465/release/4.0(1h) and then extract it using any tool of your choice, for example:

```
% unzip <name>.zip
```

c) 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 **Operations** > **Tech Support** page in the Orchestrator UI. Then click the **Edit** icon in the **System Logs** tile, and finally click the **Download** button:



This downloads the `msc_report_<date>.zip` archive to your machine.

d) 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.

e) 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.

f) 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.

# Uninstall Existing Nexus Dashboard Orchestrator

This section describes how to completely remove the existing Nexus Dashboard Orchestrator service from your Nexus Dashboard cluster, which is required as part of the Orchestrator upgrade to release 4.0(1) or later.

**Note** As mentioned in the Overview, if your Nexus Dashboard cluster is virtual, you can choose to preserve the existing VMs and deploy a brand new cluster and restore configuration there. If you choose that approach, you can skip this section and simply disconnect the existing cluster's VMs.

### Before you begin

You must have the following completed:

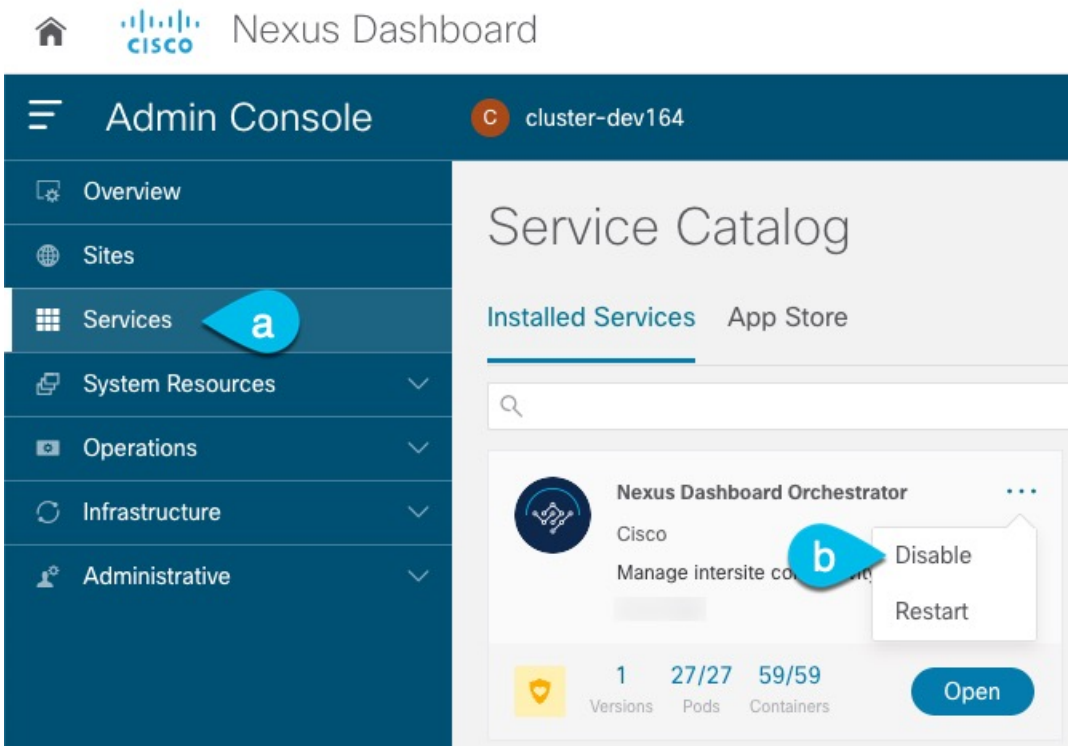
- Backed up and downloaded the existing configuration as described in Validate Existing Configuration and Create Backup, on page 5.

### Procedure

**Step 1** Navigate to your Nexus Dashboard's **Admin Console**.

**Step 2** From the main navigation menu, select **Services**.

**Step 3** Disable the existing Orchestrator service.




a) From the main navigation menu, select **Services**.
b) From the Nexus Dashboard Orchestrator tile's **Actions** (**…**) menu, choose `Disable`.

   Wait until the service is disabled.

**Step 4** From the Nexus Dashboard Orchestrator tile's **Actions** (**…**) menu, choose `Delete`.

# Upgrade Nexus Dashboard Cluster

Before you upgrade to release 4.0(x) of Nexus Dashboard Orchestrator, you must upgrade the Nexus Dashboard cluster to release 2.1(2d)) or later as described in this section.

**Note** We recommend upgrading the cluster to Release 2.2(1h) or later.

### Before you begin

If you are already running Nexus Dashboard, Release 2.1(2d)) or later, you can skip this section. Otherwise before proceeding with instructions in this section, ensure that you have:

- Backed up and downloaded the existing Orchestrator configuration as described in Validate Existing Configuration and Create Backup, on page 5.

- Uninstalled the existing Orchestrator service as described in Uninstall Existing Nexus Dashboard Orchestrator, on page 9.

Before upgrading the Nexus Dashboard cluster:

- Ensure that you have read the target release's *Release Notes* for any changes in behavior, guidelines, and issues that may affect your upgrade.

  The upgrade process is the same for all Nexus Dashboard form factors. Regardless of whether you deployed your cluster using physical servers, VMware ESX, Linux KVM, Azure, or AWS, you will use the target release's ISO image to upgrade.

- Ensure that you have read the *Release Notes* for any services you run in the existing cluster and plan to run on the target release for service-specific changes in behavior, guidelines, and issues that may affect your upgrade.

- You must be running Cisco Nexus Dashboard release 2.0(1d) or later to upgrade to release 2.1(2d)) or later.

  If you are running Cisco Application Services Engine, you must upgrade to Nexus Dashboard as described in *Cisco Nexus Dashboard Deployment Guide, Release 2.0(x)* before upgrading to release 2.0(1d) or later. In this case, we recommend upgrading your Application Services Engine cluster to Nexus Dashboard release 2.0(2h) and then to release 2.2(1h) or later.

- You must have valid DNS and NTP servers configured and reachable by all cluster nodes.

- Ensure that your current Nexus Dashboard cluster is healthy.

  You can check the system status on the **System Overview** page of the Nexus Dashboard GUI or by logging in to one of the nodes as `rescue-user` and ensuring that the `acs health` command returns `All components are healthy`.

- We recommend that you create a backup of the existing cluster configuration prior to the upgrade.

- Ensure that no configuration changes are made to the cluster, such as adding worker or standby nodes, while the upgrade is in progress.

- If you are upgrading Nexus Dashboard from release 2.1(1) or earlier, you may need to clear your browser cache after the upgrade is completed for the new event monitoring page to properly show in the UI.

### Procedure

**Step 1** Download the Nexus Dashboard image.

a) Browse to the Software Download page.

https://software.cisco.com/download/home/286327743/type/286328258

b) Choose the Nexus Dashboard version you want to download.

c) Download the Cisco Nexus Dashboard image (`nd-dk9.<version>.iso`).

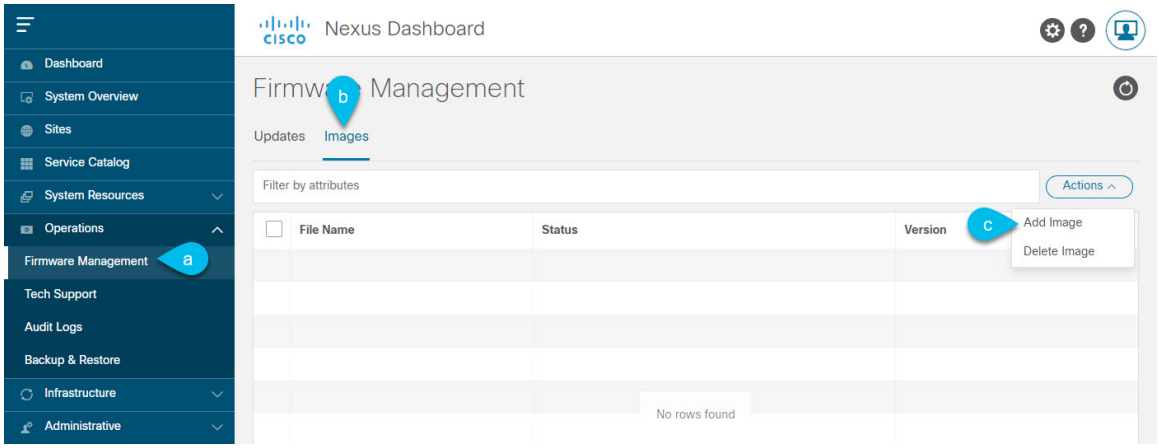
**Note**
You must download the `.iso` image for all upgrades, even if you used the VMware ESX `.ova` image or a cloud provider's marketplace for initial cluster deployment.

d) (Optional) Host the image on a web server in your environment.

When you upload the image to your Nexus Dashboard cluster, you will have an option to provide a direct URL to the image.

**Step 2** Log in to your current Nexus Dashboard GUI as an `Administrator` user.

**Step 3** Upload the new image to the cluster.




a) Navigate to **Operations** > **Firmware Management**.
b) Select the **Images** tab.
c) From the **Actions** menu, select **Add Image**.

**Step 4** Select the new image.

a) In the **Add Firmware Image** window, select **Local**.

Alternatively, if you hosted the image on a web server, choose **Remote** instead.

b) Click **Select file** and select the ISO image you downloaded in the first step.

If you chose to upload a remote image, provide the file path for the image on the remote server.

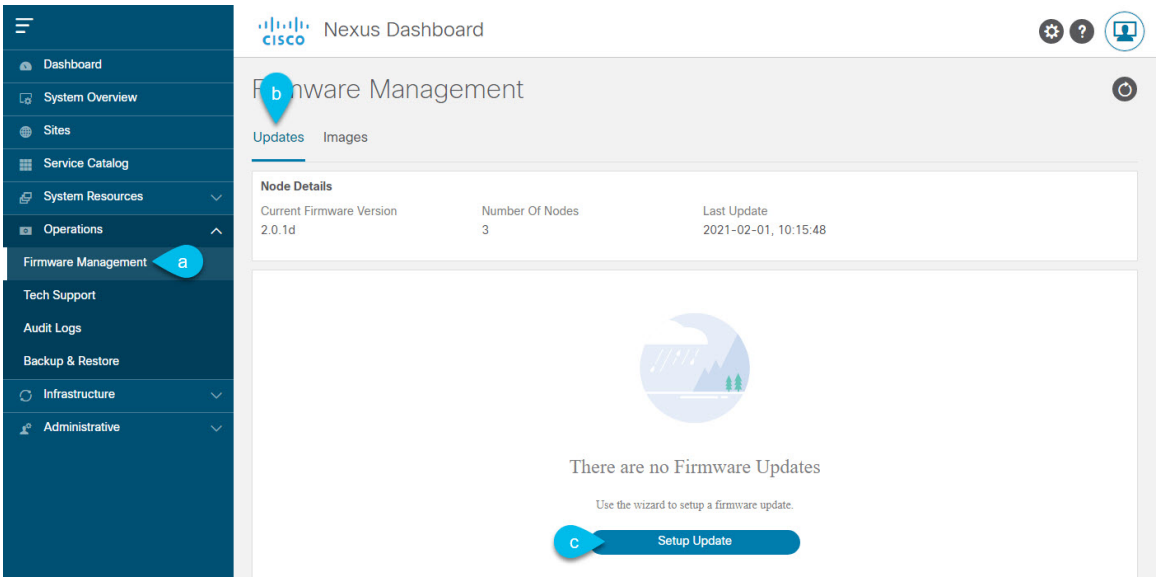
c) Click **Upload** to add the image.

The image will be uploaded to the Nexus Dashboard cluster, unpacked, processed, and made available for the upgrade. The whole process may take several minutes and you will be able to see the status of the process in the **Images** tab.

**Step 5** Wait for the image status to change to `Downloaded`.

You can check the status of the image download progress in the **Images**.

**Step 6** Set up the update.




a) Navigate to **Operations** > **Firmware Management**.
b) Select the **Updates** tab.
c) Click **Setup Update**.

   The **Firmware Update** screen opens.

**Step 7** Choose the upgrade image.

a) In the **Firmware Update** > **Version selection** screen, select the firmware version you uploaded and click **Next**.
b) In the **Firmware Update** > **Confirmation** screen, verify the details and click **Begin Install**.

   The installation progress window is displayed. You can navigate away from this screen while the update is in progress. To check on the update status at a later time, navigate to the **Firmware Management** screen and click **View Details** in the **Last Update Status** tile.

   This will set up the required Kubernetes images and services but will not switch the cluster to the new version. The cluster will continue to run the existing version until you activate the new image in the next step. The entire process may take up to 20 minutes.

**Step 8** Activate the new image.

a) Navigate back to the **Operations** > **Firmware Management** screen
b) In the **Last Update Status** tile, click **View Details**.
c) Click **Activate**.
d) In the **Activation Confirmation** window, click **Continue**.

   It may take up to 20 additional minutes for all the cluster services to start and the GUI to become available. The page will automatically reload when the process is completed.

**Step 9** If you upgraded a virtual cluster deployed in VMware ESX, convert the nodes to the new profile.

**Note**
If you upgraded a physical cluster, skip this step.

Starting with Release 2.1(1), Nexus Dashboard supports two different node profiles for virtual nodes deployed in VMware ESX. After the upgrade, you must convert all the nodes of the existing cluster to one of the new profiles:

- **Data node**—node profile designed for data-intensive applications, such as Nexus Dashboard Insights
- **App node**—node profile designed for non-data-intensive applications, such as Nexus Dashboard Orchestrator

The profile you choose depends on your use case scenario:

- If you plan to run only the Nexus Dashboard Orchestrator service, convert all nodes to the `App` node profile.
- If you plan to run Nexus Dashboard Insights or co-host applications, you must convert the nodes to the `Data` profile.

You convert the nodes to the new profile by deploying brand new nodes using that profile and replacing existing nodes with them one at a time.

a) Bring down one of the nodes.

You must replace one node at a time.

b) Deploy a new node in VMware ESX using the `App` or `Data` profile OVA.

When deploying the new node, you must use the same exact network configuration parameters as the node you are replacing.

c) Log in to the existing Nexus Dashboard GUI.

You can use the management IP address of one of the remaining healthy master nodes.

d) From the left navigation pane, select **System Resources** > **Nodes**.

The node you are replacing will be listed as `Inactive`.

e) Click the **(...)** menu next to the inactive master node you want to replace and select **Replace**.

The **Replace** window will open.

f) Provide the **Management IP Address** and **Password** for the node, then click **Verify**.

The cluster will connect to the new node's management IP address to verify connectivity.

g) Click **Replace**.

It may take up to 20 minutes for the node to be configured and join the cluster.

h) Wait for the cluster to become healthy, then repeat this step for the other two nodes.

**Step 10** If you are hosting multiple applications in the same cluster, configure deployment profiles for the App Infra Services.

If you are hosting only a single application in your Nexus Dashboard cluster, skip this step.

If you are co-hosting multiple applications in the same cluster, you must configure the App Infra Services with deployment profiles appropriate for your combination of applications and fabric sizes.

After the cluster upgrade is completed, follow the instructions described in the "App Infra Services" section of the *Cisco Nexus Dashboard User Guide*, which is also available in the products GUI.

# Install Nexus Dashboard Orchestrator Release 4.0(x)

This section describes how to install this release of Nexus Dashboard Orchestrator.

### Before you begin

You must have the following completed:

- Backed up and downloaded the existing configuration as described in Validate Existing Configuration and Create Backup, on page 5.
- Upgraded your Nexus Dashboard cluster to release 2.1(2d)) or later as described in Upgrade Nexus Dashboard Cluster, on page 10

### Procedure

**Step 1** Navigate to your Nexus Dashboard's **Admin Console**.

**Step 2** From the left navigation menu, select **Services**.

**Step 3** Install Nexus Dashboard Orchestrator.

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 4** Wait for the new image to initialize.

**Step 5** 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 6** 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.

# 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 configuration from your old Orchestrator instance as described in Validate Existing Configuration and Create Backup, on page 5.
- Upgraded your Nexus Dashboard cluster to release 2.1(2d)) or later as described in Upgrade Nexus Dashboard Cluster, on page 10.
- Installed the target Orchestrator release as described in Install Nexus Dashboard Orchestrator Release 4.0(x), on page 14.

### 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 **Operations** > **Remote Locations**.

b) In the top right of the main window, click **Add Remote Location**.

An **Add New Remote Location** screen appears.

c) 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

d) 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.

e) 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.

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

By default, port is set to `22`.

g) 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.

h) 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 from file** window that opens, click **Select File** and choose the backup file you want to import.
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.

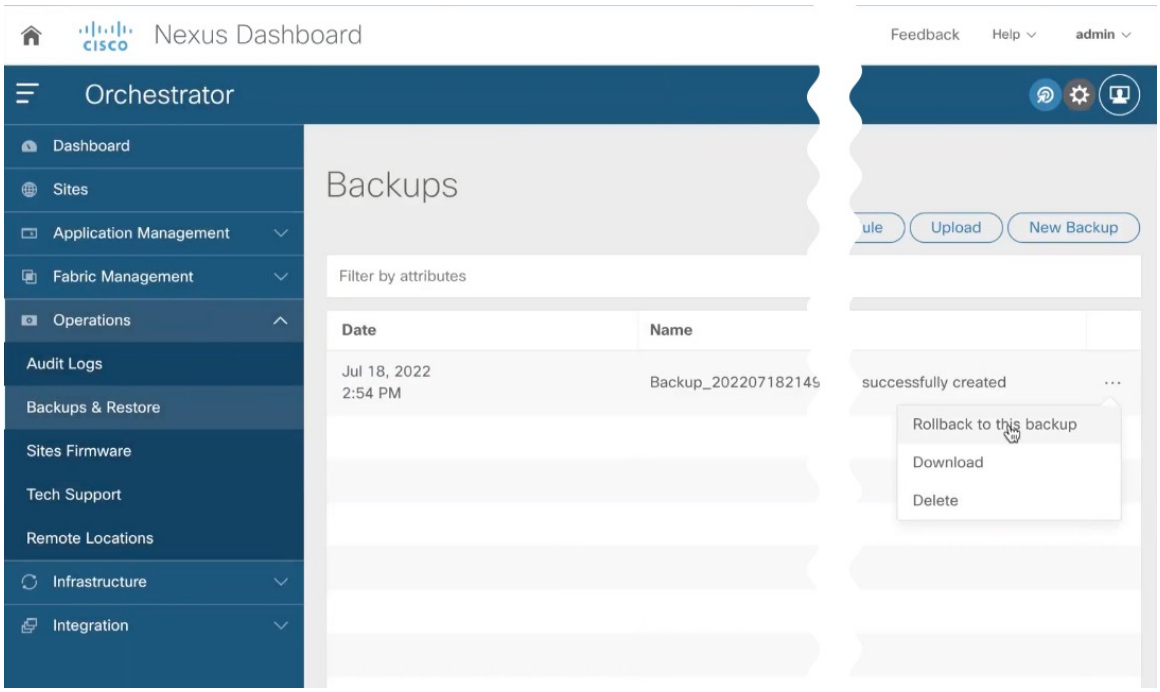
You can choose to append additional subdirectories to the path. However, the directories must be under the default configured path and must have been already created on the remote server.

f) Click **Upload** to import 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**.



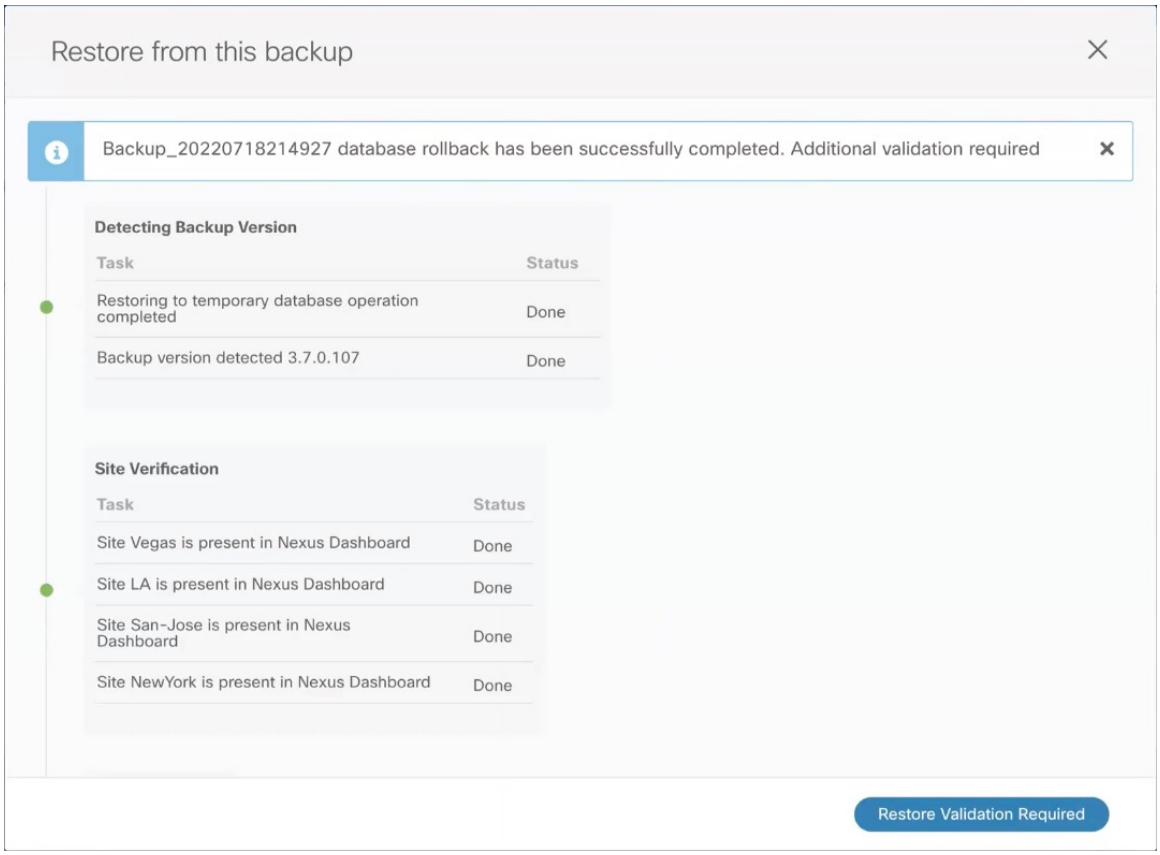

b) In the **Restore from this backup** window, click **Restore** to confirm that you want to restore the backup you selected.

This may take a few minutes to complete. After the initial backup import, you will be prompted for additional validation, which is required for database upgrades to release 4.0(1).

c) In the **Restore from this backup** window, click **Restore Validation Required** to proceed.

Before the configuration database is updated for this release, the upgrade process will perform a number of validations. A validation report is compiled and presented for you to review.

Note that 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.

Note that if validation finds any errors, the **Restore and Continue** button will be disabled and you will need to resolve any issues with your existing configuration as described in Validate Existing Configuration and Create Backup, on page 5 section and then restart the restore workflow.

The validation report generated in the previous step provides a summary of template and policy changes that will be performed during the final upgrade stage 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).

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 4.0(1) 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 4.0(1) deployment requirements.

**Note**
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. After all templates are saved, you will need to verify that all objects were imported and re-created successfully and then check for any configuration drifts between the newly created configuration and what is actually deployed in the fabrics as described in the next section.

**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.

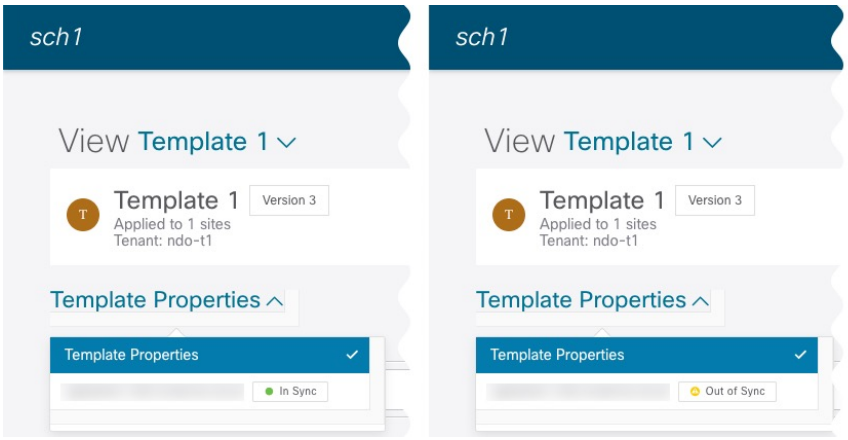
d) In the main pane, click **Configure** to open **Fabric Connectivity Infra** screen and verify **External Subnet Pool** addresses.

   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 as shown in the following figure:




After upgrading your Nexus Dashboard Orchestrator and restoring the previous configuration backup, we recommend that you check for and resolve any configuration drifts 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 **Application Management** > **Schemas**.

**Step 2** 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 3** 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*.

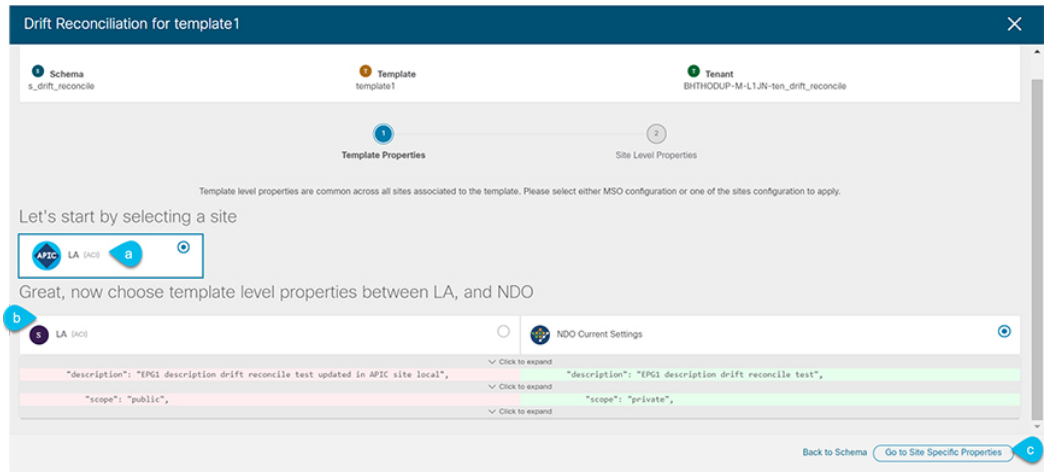
a) Close the template deployment dialog to return to the Schema view.

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 **Reconcile Drift**.

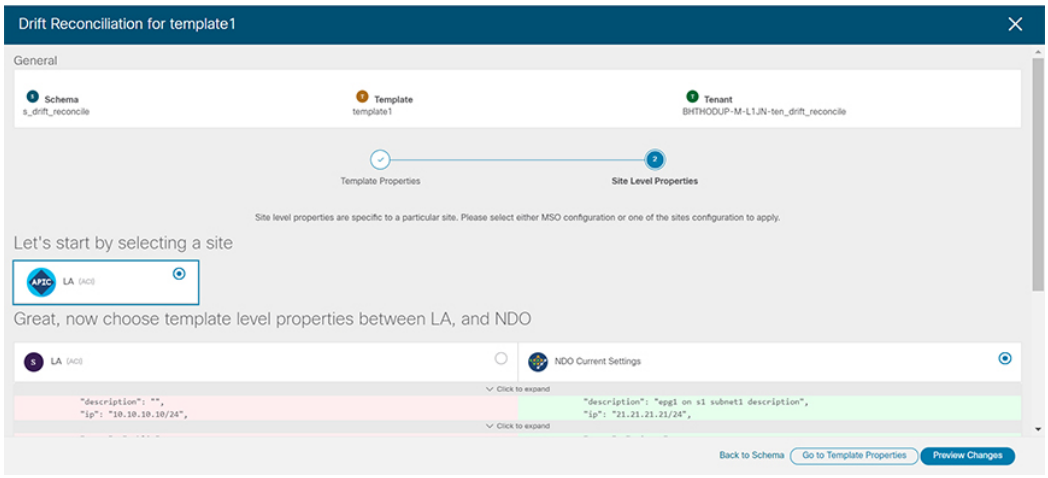
The **Drift Reconciliation** wizard opens.

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. You can then click **Deploy to sites** to deploy the configuration and reconcile the drift for that template.

**Step 4** Repeat the above steps for every schema and template in your Nexus Dashboard Orchestrator.

**Step 5** 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.