



Upgrading an Existing Nexus Dashboard Cluster to This Release

- [Prerequisites and guidelines for upgrading an existing Nexus Dashboard cluster, on page 1](#)
- [Supported upgrade paths, on page 4](#)
- [Upgrade Nexus Dashboard, on page 5](#)
- [Post-upgrade information and tasks, on page 7](#)
- [Troubleshooting upgrades, on page 11](#)

Prerequisites and guidelines for upgrading an existing Nexus Dashboard cluster

Before you upgrade your existing 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.
- Ensure that you perform a backup of your Nexus Dashboard cluster before upgrading and you store the backup file in a safe place. To perform a backup, refer to [Unified Backup and Restore for Nexus Dashboard and Services](#). Note that you will not be able to restore this backup directly to a Nexus Dashboard cluster running the 4.1.1 release.

In addition, the upgrade process might fail if the most recent backup had a failure. In that case, perform another backup and make sure it succeeds before attempting another upgrade of the Nexus Dashboard cluster. If you are unable to perform a successful backup and cannot upgrade, contact [Cisco Technical Assistance Center \(TAC\)](#) for support.

- If you have either NDI or NDFC, with NDI performing telemetry for remotely-created NDFC clusters, or if you have multi-cluster connectivity with multiple ND clusters, all the clusters must be upgraded to Nexus Dashboard 4.1.1. Having a mix of Nexus Dashboard release 3.2x and 4.1.1 clusters using multi-cluster connectivity is not supported.
- If you are upgrading a physical Nexus Dashboard cluster, ensure that the nodes have the minimum supported CIMC version for the target Nexus Dashboard release.

Supported CIMC versions are listed in the [Nexus Dashboard Release Notes](#) for the target release.

The CIMC upgrade is described in detail in the "Troubleshooting" article in the Nexus Dashboard [documentation library](#).

- If you are upgrading a virtual Nexus Dashboard cluster, Nexus Dashboard will enforce a check of the HDD latency to verify that it is <30ms. If the HDD has a higher latency, the upgrade will fail.
- If you are upgrading a virtual Nexus Dashboard cluster deployed in VMware ESX, ensure that the ESX version is still supported by the target release.

This release supports VMware ESXi 7.0, 7.0.1, 7.0.2, 7.0.3, 8.0, 8.0.2, 8.0.3.



Note

If you need to upgrade the ESX server, you must do that before upgrading your Nexus Dashboard. ESX upgrades are outside the scope of this document, but in short:

1. Upgrade one of the ESX hosts as you typically would with your existing Nexus Dashboard node VM running.
2. After the host is upgraded, ensure that the Nexus Dashboard cluster is still operational and healthy.
3. Repeat the upgrade on the other ESX hosts one at a time.
4. After all ESX hosts are upgraded and the existing Nexus Dashboard cluster is healthy, proceed with upgrading your Nexus Dashboard to the target release as described in this document.

- Ensure that your current Nexus Dashboard cluster is healthy.

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

- Nexus Dashboard does not support platform downgrades.

If you want to downgrade to an earlier release, you will need to deploy a new cluster.

- If you have a user who only has the `Dashboard User` (app-user) user role in Nexus Dashboard release 3.2.1, after the upgrade to Nexus Dashboard release 4.1.1, delete the user with the `Dashboard User` user role or use the `Observer` role instead for that user in Nexus Dashboard release 4.1.1.
- The number of persistent IP addresses and how they are mapped out has changed from previous releases to Nexus Dashboard release 4.1.1. See [Nexus Dashboard persistent IP addresses](#) to understand the number of persistent IP addresses that were needed in previous releases, and how you will have to make certain updates in the number of persistent IP addresses that you will need before you can upgrade to Nexus Dashboard release 4.1.1.
- If you have a Nexus Dashboard (ND) in any persona with any services, specifically Nexus Dashboard Fabric Controller (NDFC) and Nexus Dashboard Insights (NDI), and you have upgraded from previous releases (such as ND 2.2.x and below) to ND 3.2.x, be aware of the following important information regarding Elasticsearch (ES) indices size:

If the cluster was deployed on ND 2.2.x or earlier and was upgraded since then, the upgrade from ND 3.2.x to ND 4.1.1 may be indexing the time series database. The process will prompt you to contact [Cisco Technical Assistance Center \(TAC\)](#) if the sizes exceed the sizes of the ability to reindex on this cluster.

If the process fails to reindex, you can use the `acs recover` command that's provided in the UI to proceed with the upgrade after the failure.

- If you had multi-cluster connectivity configured and you had NDFC and NDI co-located in your Nexus Dashboard release 3.2.x system, where:
 - NDFC was running on one cluster, and
 - NDI was running on another cluster

Then we mandate that you disconnect the clusters and delete the federation before you begin the upgrade process to Nexus Dashboard release 4.1.1. See the sections "Disconnecting Clusters" and "Deleting the Federation" in the [Nexus Dashboard Infrastructure Management](#) for those procedures. After you have completed the upgrade, you will re-enable multi-cluster connectivity as part of the post-upgrade tasks.

- If you have Nexus Dashboard Orchestrator as part of your pre-ND release 4.1.1 cluster with different data subnets, before upgrading to Nexus Dashboard release 4.1.1, you must make the following configurations:
 - Add BGP configurations on all the nodes
 - Add persistent IP addresses

The validation of the image during the upgrade will fail if you do not have these items configured before upgrading.

Auto reconcile configuration drifts on upgrade

When upgrading to ND release 4.1.1 from any pre-4.2.(1) NDO release, you must perform a multi-step upgrade:

1. First upgrade the pre-4.2(1) NDO release to ND 3.0.1i, then to ND release 3.2.x, as described in the ["Supported Upgrade Paths"](#) section in *Cisco Nexus Dashboard and Services Deployment and Upgrade Guide, Release 3.2.x*.
2. Then upgrade from ND release 3.2.x to ND release 4.1.1, as described in this chapter.

When upgrading a pre-4.2(1) NDO release to ND release 3.2.x, you may observe some configuration drifts in application templates. These drifts occur because NDO release 4.2(1) and later support managing new application template object properties that were not managed in earlier versions. For a list of new properties managed by NDO 4.2(1), refer to the [Nexus Dashboard Orchestrator 4.2\(1\) release notes](#).

ND release 4.1.1 supports automatic reconciliation of configuration drifts as part of the upgrade process. Nexus Dashboard will check whether the templates are in sync with the fabrics and, if necessary, will import fabric values into Nexus Dashboard to automatically resolve the drifts.

When following the multi-step upgrade path from a pre-4.2(1) NDO version, we recommend that you ignore any drifts detected in ND release 3.2.x and continue the upgrade to ND release 4.1.1.

Any drifts remaining after upgrading to ND release 4.1.1 will need to be resolved manually. For example, a drift that cannot be auto-resolved occurs when a configured object is part of a stretched template across multiple fabrics, and the template-level property has different values configured on different fabrics.

Supported upgrade paths

As described in [Nexus Dashboard deployment overview](#), in earlier releases, Nexus Dashboard shipped with only the platform software and no services included, which you would then download, install, and enable separately after the initial platform deployment. In addition, Nexus Dashboard release 3.1(1) introduced a tighter coupling between the Nexus Dashboard and individual services with only a single version of each service compatible with each version of the platform. As a result, as long as you were on the minimum required version of the Nexus Dashboard software, you could upgrade both the platform and all currently enabled services directly to Nexus Dashboard release 3.1x and 3.2x.

Now, the platform and the individual services have been unified into a single product, which means that you no longer deploy, configure, or upgrade the services separately.

The following table provides a few example scenarios for specific deployment combinations:

Table 1:

Current Nexus Dashboard Release	Compatible Services (depending on form factor and cluster size, you may have one or more of these services currently enabled)	Upgrade Workflow
3.2(2)	Fabric Controller: 12.2(3) Orchestrator: 4.4(2) Insights: 6.5(2)	Upgrade directly to release 4.1(1) as described in the following section. All services are unified under a single Nexus Dashboard product in release 4.1(1).
3.2(1)	Fabric Controller: 12.2(2) Orchestrator: 4.4(1) Insights: 6.5(1)	Upgrade directly to release 4.1(1) as described in the following section. All services are unified under a single Nexus Dashboard product in release 4.1(1).
3.1(1)	Fabric Controller: 12.2(1) Orchestrator: 4.3(x) Insights: 6.4(1)	<ol style="list-style-type: none"> 1. Upgrade the Nexus Dashboard platform to release 3.2(x) as described in Nexus Dashboard Deployment Guide, Release 3.2.x All services will be automatically upgraded along with the platform. 2. Upgrade from release 3.2(x) to release 4.1(1) . All services are unified under a single Nexus Dashboard product in release 4.1(1).
3.0(1)	Fabric Controller: 12.1(3) Orchestrator: 4.2(x) Insights: 6.3(1)	<ol style="list-style-type: none"> 1. Upgrade the Nexus Dashboard platform to release 3.2(x) as described in Nexus Dashboard Deployment Guide, Release 3.2.x All services will be automatically upgraded along with the platform. 2. Upgrade from release 3.2(x) to release 4.1(1) . All services are unified under a single Nexus Dashboard product in release 4.1(1).

Current Nexus Dashboard Release	Compatible Services (depending on form factor and cluster size, you may have one or more of these services currently enabled)	Upgrade Workflow
2.3(2) and earlier	Fabric Controller: 12.1(2) or earlier Orchestrator: 4.1(x) or earlier Insights: 6.2(x) or earlier	<ol style="list-style-type: none"> 1. Upgrade the Nexus Dashboard platform to release 3.1(1) as described in Nexus Dashboard Deployment Guide, Release 3.1.x. All services will be automatically upgraded along with the platform. 2. Upgrade from release 3.1(1) to release 3.2(x) as described in Nexus Dashboard Deployment Guide, Release 3.2.x. All services will be automatically upgraded along with the platform. 3. Upgrade from release 3.2(x) to release 4.1(1). All services are unified under a single Nexus Dashboard product in release 4.1(1).

Upgrade Nexus Dashboard

This section describes how to upgrade an existing Nexus Dashboard cluster.

As described in [Supported upgrade paths, on page 4](#), you must be on Nexus Dashboard release 3.2(x) to upgrade directly to Nexus Dashboard release 4.1(1). Note these important points on these releases:

- **Nexus Dashboard release 3.2(x):** As described in [Nexus Dashboard deployment overview](#), for these Nexus Dashboard releases, Nexus Dashboard was available as one product, and individual services (such as Nexus Dashboard Insights, Orchestrator, and Fabric Controller) were available as individual products, separate from Nexus Dashboard.
- **Nexus Dashboard release 4.1(1):** This is the first Nexus Dashboard release where Nexus Dashboard and the individual services listed above are packaged together as a single, unified product.

This means that your upgrade process will go through these phases:

1. You will begin the upgrade process in Nexus Dashboard release 3.2(x), where Nexus Dashboard and the individual services are separate products.
2. You will then upgrade to Nexus Dashboard 4.1(1), where you will end the upgrade process with Nexus Dashboard and the individual services packaged together as a single, unified product.

Before you begin

Ensure that you have completed the prerequisites described in [Prerequisites and guidelines for upgrading an existing Nexus Dashboard cluster, on page 1](#)

Procedure

Step 1 In your Nexus Dashboard release 3.2(x) system, download the Nexus Dashboard 4.1(1) image.

- a) Browse to the Software Download page.

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

- b) Choose the Nexus Dashboard 4.1(1) release to download.
c) Download the Nexus Dashboard image for the 4.1(1) release.

Note

The upgrade process is the same for all Nexus Dashboard form factors and uses the Nexus Dashboard ISO image (`nd-dk9.<version>.iso`). In other words, even if you used the virtual form factors (such as the ESX `.ova`) or a cloud provider's marketplace for initial cluster deployment, you must still use the `.iso` image for upgrades.

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

Note

We recommend hosting the image on a 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, which can significantly speed up the process.

Step 2 Log in to your current Nexus Dashboard's **Admin Console** as an `Administrator` user.

Step 3 Delete any older, non-active upgrade images from your cluster.

If this is the first time you're upgrading your cluster, you can skip this step.

- a) Navigate to **Manage > Software Management**.
b) Click the trash icon on an upgrade image's tile to delete any older, non-active upgrade **Images**.
c) Repeat this step for all older, non-active upgrade images.

Step 4 Upload the new image to the cluster.

- a) Navigate to **Manage > Software Management**.
b) Click **Add Image**.
c) In the **Add Software Image** window, select whether the image is **Remote** on a web server or **Local** on your machine.

In both cases, the image will be a file ending with `.iso`.

- **Remote:** Provide the **URL** to the image you downloaded in the first step.
- **Local:** Click **Choose file** and navigate to the local folder where you downloaded the image.

- d) Click **Add** to add the image.

Nexus Dashboard then downloads the upgrade image and starts processing the image, and goes through a number of preparation and validation stages to ensure successful upgrade. This may take several minutes to complete.

Note

See [Troubleshooting upgrades, on page 11](#) for more information on the validation checks that occur during this point of the upgrade and how to deal with upgrade issues that might arise.

- e) After the validation is complete, then the **Install** button appears in the card in the **Software Management** page. Click **Install** to install the software and go through the upgrade process.

The installation progress window is displayed. You can navigate away from this screen while the update is in progress.

This step may take up to 60 minutes or more, depending on the number of nodes in the cluster, during which the nodes will reboot and the GUI will not be accessible. Nexus Dashboard goes through several stages:

- Install Release Firmware
- Disable Services
- Shutdown Infrastructure services
- Update Platform Services
- Enable Infrastructure Services
- Enable Services

You can click on the **Details** link to see the progress and the various stages of the upgrade.

Note

If you see any issues during the upgrade process, such as a possible indexing issue, refer to [Prerequisites and guidelines for upgrading an existing Nexus Dashboard cluster, on page 1](#) for more information and possible workarounds.

After the process above is complete, you should be upgraded to Nexus Dashboard 4.1(1), where Nexus Dashboard and the individual services are packaged together as a single, unified product.

Note

Depending on the cluster format and the number of cluster nodes that you have deployed, certain features (such as controller, orchestrator, or telemetry) might not be available. Review the information in the [Nexus Dashboard Capacity Planning tool](#) to verify what features would be available for your cluster installation.

Step 5 After the node upgrade tasks are completed, verify that the nodes are healthy and you can log into the UI.

Once the upgrade process completes, you can view the Nexus Dashboard UI as you typically would.

You can check the **Overview** page for overall system health and the **Admin > System Software** page to see the current *Running version*.

What to do next

Go to [Post-upgrade information and tasks, on page 7](#) to perform the necessary post-upgrade tasks.

Post-upgrade information and tasks

This section provides information about changes and tasks that you must complete after upgrading from Nexus Dashboard release 3.2.x to Nexus Dashboard 4.1.1.

- [Multi-cluster connectivity, on page 8](#)
- [Address Anomalies issues, on page 8](#)
- [Set device credentials, on page 8](#)
- [Migrate older cluster types to new 3 node virtual cluster \(data\) cluster type, on page 9](#)

- [Redeploy the telemetry configuration, on page 9](#)
- [Changes in SNMP server users, on page 11](#)
- [View historical Performance Manager \(PM\) data, on page 11](#)

Multi-cluster connectivity

If you had any of these configurations in pre-4.1.1. Nexus Dashboard:

- If you had Nexus Dashboard Insights onboarding remotely-created NDFC fabrics
- If you were using One Manage to manage and monitor multiple NDFC or NDI clusters
- If you had multi-cluster fabric groups with multiple NDFC clusters

If these clusters were already connected in the previous release using multi-cluster connectivity, then, on the primary cluster, you will have to re-register all the clusters that you upgraded to Nexus Dashboard release 4.1.1.

In addition, if you had NDI and NDO integrations in the previous release, this is not supported in Nexus Dashboard release 4.1.1, and in order to leverage the NDO integration, you will have to federate NDI and NDO clusters in Nexus Dashboard release 4.1.1.

If you had multi-cluster connectivity configured in your Nexus Dashboard release 3.2.x system, after you have completed the upgrade, you will need to re-enable multi-cluster connectivity. This task is applicable only when you have NX-OS fabrics in a co-location environment. See [Connecting Clusters](#) for more information.

Address Anomalies issues

After you have completed the upgrade, check the **Anomalies** area for any issues and check for requests to complete the **Recalculate and Deploy** in some NX-OS fabrics..

Set device credentials

If you had co-hosted NX-OS fabrics configured in the Nexus Dashboard release 3.2.x system, follow these procedures to make sure the appropriate device credentials are configured after you've upgraded to Nexus Dashboard release 4.1.1.

1. In your upgraded Nexus Dashboard 4.1.1 system, navigate to **Manage > Device Credentials**.
2. Review the information displayed in the **Device Credentials** area.
You should see **Not Set** displayed in red text in the **Device Credentials** area.
3. In the **Device Credentials** area, click **Set**.
4. In the **Set Default Credentials** page, enter the necessary information.
 - Username, Password, and Confirm password: Enter the necessary username and password information.
 - Robot: Choose the **Robot** checkbox to set the robot credentials, if necessary.

Then click **Save**.

You are returned to the **Set Default Credentials** page, and the text **Default Set** is now displayed in blue.

Migrate older cluster types to new 3 node virtual cluster (data) cluster type

The pre-4.1.1 cluster type `App node with 1.5TB disk` is not supported in Nexus Dashboard release 4.1.1; after upgrading to Nexus Dashboard release 4.1.1, it will show up as `SE-VIRTUAL-APP-LARGE`.

In addition, these pre-4.1.1 cluster types are not supported as greenfield deployments in Nexus Dashboard release 4.1.1 but are supported when upgrading to release 4.1.1:

- 3-node virtual cluster (app node with 1.5TB storage)
- 5-node virtual cluster (app 500G or 1.5TB storage)

If you want to migrate from those older cluster types to the new cluster type `3 node virtual cluster (data)`, you can migrate to the new cluster type using these procedures:

1. Upgrade to Nexus Dashboard release 4.1.1, keeping the older cluster type as-is, using the procedures provided in this chapter.
2. After you have upgraded to Nexus Dashboard release 4.1.1, perform a backup of your older cluster type. See [Backing Up and Restoring Your Nexus Dashboard](#) for more information.
 - You can use either a **Config-only** or **Full** backup option when backing up the older cluster type.
 - Verify that you have successfully retrieved the backup from the older cluster before proceeding.
3. Shut down the cluster with the older cluster type.
4. Perform a greenfield deployment of the new cluster type `3 node virtual cluster (data)`.
 - Reuse the name of the older cluster for the new cluster.
 - You can deploy a virtual data cluster or a physical cluster.
 - You can reuse the IP addresses from the older cluster type, or you can use new IP addresses in the new cluster and restore with a check in the **Ignore External Service IP Configuration** check box, as described in [Backing Up and Restoring Your Nexus Dashboard](#).
5. Restore the backup of the older `3 node virtual cluster (app)` or `5 node virtual cluster (app)` to the newer `3 node virtual cluster (data)`.
See [Backing Up and Restoring Your Nexus Dashboard](#) for more information.

Redeploy the telemetry configuration

After upgrading to Nexus Dashboard release 4.1.1, the persistent IP addresses for handling software telemetry streaming from the NX-OS fabrics would have changed. To resume telemetry operations after the upgrade to Nexus Dashboard 4.1.1, you must redeploy the telemetry configuration.

1. Navigate to the **System Status** page.
Admin > System Status
2. Click **Telemetry**, then click the **Fabrics** tab in the **Telemetry status** area.
3. Review the information displayed in the **Fabrics** table.

For NX-OS fabrics, you will notice that one or more fabrics listed in the **Fabrics** table will show a status of **Pending updates** in the **Telemetry config status** column, even though those same fabrics had a status of **OK** prior to the upgrade.

In this state, telemetry streaming will not be functional, and the status of individual features (such as software telemetry, flow collection, and switch status) are not valid and should be ignored.



Note Some fabrics in the **Fabrics** table might show a status of **OK** in the **Telemetry config status** column, while other fabrics might show a status of **Pending updates**. If you were to click the **Switches** tab in the **Telemetry status** area, you might see switches with incorrect green **OK** or **Success** status entries in the telemetry columns, even though those switches are associated with the fabrics that show a status of **Pending updates** in the **Fabrics** table. This is incorrect status information displayed at the switch level for those fabrics and should be ignored.

4. With the **Fabrics** tab selected in the **Telemetry status** area, click **Redeploy telemetry**.

Click **Confirm** in the confirmation page.

After you click **Confirm**, the status of the individual features will change to **Enable in Progress**, and the cumulative **Telemetry Config Status** will update to **In Progress**.

5. After you click **Confirm**, you are returned to the **Fabrics** tab in the **Telemetry** page under **System Status**. Click **Refresh** in the upper right corner of the page.

The telemetry status might be shown incorrectly immediately after you click **Confirm** in the confirmation page, but will show the correct telemetry status after you click **Refresh**.

6. Review the information displayed in the **Fabrics** table again.

After clicking **Refresh**, you should see that the status change in these areas:

- The status of the individual features will change to **Enable in Progress**, and
- The status for the fabric in the **Telemetry config status** column will change from **Pending updates** to **In Progress**. After several minutes, the status for the fabric will change to **OK** in the **Telemetry config status** column, either automatically or after you click **Refresh** again.

After the operation completes, the individual feature statuses will be:

- `Enabled` if the configuration push is successful for all switches,
- `Enable Fail` if it fails for any switch, or
- `Enable Pending` if change control mode is enabled. In that case, apply the change control ticket explicitly through **Manage > Change control**. See [Using Change Control and Rollback in Your Nexus Dashboard](#) for more information.

The cumulative **Telemetry Configuration Status** will then display as:

- `OK` if the configuration is successful for all switches,
- `Not OK` if it fails or is pending in change control mode for all switches, or
- `Partial OK` if it succeeds for some switches and fails or is pending in change control mode for others.

Changes in SNMP server users

In releases prior to Nexus Dashboard release 4.1.1, SNMP server users that were configured on managed NX-OS switches without passwords as part of the intent, such as the example shown below:

```
snmp-server user Demo_CMDv5 vdc-operator
snmp-server user DemoOps_admin vdc-operator
snmp-server user DemoOps_admin network-admin
```

were not shown in the diffs during **Recalculate and Deploy** operations. These often went unnoticed, especially in environments where the switches relied on remote authentication methods such as TACACS+.

Once you upgrade to Nexus Dashboard release 4.1.1, these differences are now correctly detected and displayed in the GUI as expected diffs. After the upgrade to release 4.1.1, you must now either:

- Push these configurations to bring the switches into sync, or
- Remove these SNMP user entries from the intent if they are no longer needed.

View historical Performance Manager (PM) data

After you upgrade Nexus Dashboard from release 3.2.x to release 4.1.1, if you enable Telemetry on the fabrics, the historical Performance Manager (PM) data will not be displayed. Instead, the system will start collecting PM data fresh from that point forward.

To view the historical PM data, disable Telemetry on the affected fabrics. The older PM data will now re-appear.

Troubleshooting upgrades

After all the nodes restart during new image activation stage described in the previous section, you may log in to the GUI to check the status of the upgrade workflows. Initially, you can see the bootstrap process similar to the initial cluster deployment and once the nodes come up, you can see additional information about service activation in the GUI's **Overview** page.

In case the upgrade fails for any reason, the GUI will display the error and additional workaround steps. For example, you might then see an error message, along with the remedy, similar to this:

Failed to activate

Upgrade failed while shutting down the cluster: Operation Timedout, last status: Operation Timedout

Please login to one of the primary nodes as 'rescue-user' and follow the steps provided by the upgrade recovery helper by invoking following command: 'acs upgrade recover Cluster Shutdown'. If the issue persists, please contact Cisco TAC for assistance.

If an issue persists, click **Admin** to access Tech support. See [Working with Cisco Tech Support](#) for more information.

