

Deployment Overview and Requirements

- Nexus Dashboard deployment overview, on page 1

Nexus Dashboard deployment overview

Nexus Dashboard platform

Cisco Nexus Dashboard is a central management console for multiple data center fabrics that provides real-time analytics, visibility, assurance for network policies and operations, as well as policy orchestration for the data center fabrics, such as Cisco ACI and NX-OS.

Nexus Dashboard is the comprehensive management solution for ACI as well as NX-OS deployments spanning LAN fabric, SAN fabric, and IP Fabric for Media (IPFM) networks in data centers powered by Cisco. Nexus Dashboard also supports other devices, such as IOS-XE switches, IOS-XR routers, and non-Cisco devices. Being a multi-fabric controller, Nexus Dashboard manages multiple deployment models such as VXLAN EVPN, classic 3-tier LAN, FabricPath, and routed-based fabrics for LAN while providing ready-to-use control, management, monitoring, and automation capabilities for all these environments. In addition, Nexus Dashboard, when you select SAN installation, Cisco Nexus Dashboard automates Cisco MDS switches and Cisco Nexus-family infrastructure in NX-OS mode with a focus on storage-specific features and analytics capabilities.

Note This document describes how to deploy a Nexus Dashboard cluster initially and onboard the fabrics. After your cluster is up and running, see the Nexus Dashboard *configuration and operation articles* for day-to-day operation.

Unified Nexus Dashboard deployment

The Nexus Dashboard (ND) platform and related services were available in the following ways previously:

- Prior to ND release 3.1, Nexus Dashboard shipped with only the platform software and no services included. You would download, install, and enable the services separately (NDI, NDO, and/or NDFC) after the initial ND platform deployment.
- For ND releases 3.1 and 3.2, Nexus Dashboard packaged the ND platform software and the individual services' software in a unified packaging form; however, you still enabled the services separately. Management and Insights of the fabrics were still two independent pieces that were not unified.

In addition, there existed a concept of "deployment mode" in Nexus Dashboard releases 3.1 and 3.2, where you would statically enable specific services in Nexus Dashboard by selecting the deployment mode. However, changing the deployment mode was a disruptive exercise that would wipe out the entire service, including data and reinstalls. And finally, you could not run all the services in a single Nexus Dashboard cluster in Nexus Dashboard releases 3.1 and 3.2.

Beginning with ND release 4.1, the platform and the individual services have been unified into a single product. You no longer deploy and configure the services separately, and you do not have to activate individual services or statically configure deployment modes. In addition, depending on the form factor, Nexus Dashboard allows you to consume any capabilities that were shipped as services in previous releases. The user experience is now unified, as there is no more concept of independent services; instead, all capabilities are now available from a single dashboard view.

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 in the unified Nexus Dashboard product. Review the information in the Nexus Dashboard Capacity Planning tool to verify what features would be available for your cluster installation.

Hardware vs software stack

Nexus Dashboard is offered as a cluster of specialized Cisco UCS servers (Nexus Dashboard platform) with the software framework (Nexus Dashboard) pre-installed on it. The Cisco Nexus Dashboard software stack can be decoupled from the hardware and deployed in a number of virtual form factors. For the purposes of this document, we will use "Nexus Dashboard hardware" specifically to refer to the hardware and "Nexus Dashboard" to refer to the software stack and the GUI console. In addition, we will use the term "physical Nexus Dashboard" or "pND" to refer to the Cisco UCS physical appliance hardware with the Nexus Dashboard software stack pre-installed on it, and "virtual Nexus Dashboard" or "vND" to refer to the virtual form factor that supports data and app nodes.

This guide describes the initial deployment of the Nexus Dashboard software, which is common for physical and virtual form factors. If you are deploying a physical cluster, see *Nexus Dashboard Hardware Setup Guide* for the UCS servers' hardware overview, specification, and racking instructions.

Note Root access to the Nexus Dashboard software is restricted to Cisco TAC only. A special user `rescue-user` is created for all Nexus Dashboard deployments to enable a set of operations and troubleshooting commands. For additional information about the available `rescue-user` commands, see the "Troubleshooting" article in the Nexus Dashboard *documentation library*.

Available form factors

This release of Cisco Nexus Dashboard can be deployed using a number of different form factors. However, you must use the same form factor for all nodes, mixing nodes of different form factors within the same cluster is not supported. The physical form factor currently supports four different Cisco UCS servers (`SE-NODE-G2`, `ND-NODE-L4`, `ND-NODE-G5S`, and `ND-NODE-G5L`). You can mix `SE-NODE-G2` and `ND-NODE-L4` servers in the same cluster, but you cannot mix a `ND-NODE-G5S` or `ND-NODE-G5L` server in the same cluster as `SE-NODE-G2` and `ND-NODE-L4` servers.

- Physical appliance (`.iso`) – This form factor refers to the Cisco UCS physical appliance hardware with the Nexus Dashboard software stack pre-installed on it.

  The later sections in this document describe how to configure the software stack on the existing physical appliance hardware to deploy the cluster. Setting up the Nexus Dashboard hardware is described in *Nexus Dashboard Hardware Setup Guide* for the specific UCS model.

- Virtual Appliance – The virtual form factor allows you to deploy a Nexus Dashboard cluster using VMware ESX (`.ova`) or RHEL KVM (`.qcow2`).

  The virtual form factor supports the following two profiles:

  - Data node – This profile with higher system requirements is designed for higher scale and/or unified deployment.
  - App node – This profile with lower system requirements can be deployed as secondary nodes. Can also be deployed as primary nodes but does not support unified deployment.

  Support is also available for running a virtual Nexus Dashboard (vND) on the AWS public cloud. See Deploying a Virtual Nexus Dashboard (vND) in Amazon Web Services (AWS) for more information.

  Beginning with Nexus Dashboard release 4.2(1), support is also available for running a virtual Nexus Dashboard (vND) on the Nutanix. See Deploying a Virtual Nexus Dashboard (vND) in Nutanix for more information.

Note When planning your deployment, ensure to check the list of "Prerequisites and Guidelines" in one of the following sections of this document specific to the form factor you are deploying. A quick reference of the supported form factors, scale, and cluster sizing requirements are available in the Nexus Dashboard Cluster Sizing tool.

Scale and cluster sizing guidelines

A basic Nexus Dashboard deployment typically consists of 1 or 3 `primary` nodes, which are required for the cluster to come up. Depending on scale requirements, 3-node or larger clusters can be extended with up to 3 additional `secondary` nodes to support higher scale. Adding secondary nodes to the cluster must be done during cluster deployment and cannot be added once the cluster has been initialized.

- For physical clusters, you can also add up to 2 `standby` nodes for easy cluster recovery in case of a primary node failure.
- For virtual clusters, up to 2 `standby` nodes are also supported, but only with a 3-node vND (app) profile for a Controller-only or Orchestration-only deployment.

Exact number of additional secondary nodes required for your specific use case is available from the Nexus Dashboard Cluster Sizing tool.

Scale and cluster sizing limitations

These limitations apply to scale and cluster sizing:

- The `ND-NODE-G5L` (UCS-C225-M8) node is available only in a 3-node cluster configuration.
- Deploying a 4-cluster node, where the cluster consists of:

- 3 virtual nodes (data), and
- 1 standby node

is not a supported configuration. Redeploy this cluster without the standby node. If one of the nodes in the cluster fails, reinstall a new node and navigate to **Admin** > **System Status** > **Nodes**, then click **Actions** > **Re-Register** to add the reinstalled node back into the cluster.

- Single-node deployments cannot be extended to a 3-node cluster after the initial deployment.

  If you deploy a single-node cluster and want to extend it to a 3-node cluster or add `secondary` nodes, you will need to back it up, deploy a new 3-node base cluster, and then restore the backup on this later. For more information, see *Backing Up and Restoring Your Nexus Dashboard*
- Single-node deployments do not support additional `secondary` or `standby` nodes.
- For 3-node clusters, at least two `primary` nodes are required for the cluster to remain operational.

  For more information, see *Deploying Highly Available Services with Cisco Nexus Dashboard*.
- You cannot dynamically add or remove primary or secondary nodes to change the overall size of an existing cluster. To resize a cluster, you must back up the cluster, rebuild it to the desired size, and then restore the configuration. Ensure that the new cluster size meets the supported scale and feature requirements for your deployment.
- When a cluster size change is required, delete only one secondary node at a time. Wait until the cluster returns to a healthy state before you delete the next secondary node. Verify the cluster health in the Nexus Dashboard GUI or by running the acs health command. Do not delete multiple secondary nodes concurrently or while the cluster is converging.
- For validated scale, all required worker nodes should be included when the Nexus Dashboard cluster is initially deployed. Additional worker nodes can also be added to an existing cluster; however, in the current release, adding worker nodes after initial deployment does not dynamically increase the cluster's validated scale capacity. The additional nodes can still help distribute workloads across a larger set of resources, reducing concentration on individual nodes. If increased scale capacity is required after worker nodes have been added, the cluster can be backed up and restored so that the expanded topology is incorporated into the scale configuration.

About supported node types and features

- Several node types are available for Nexus Dashboard, including SE-NODE-G2 (UCS-C220-M5), ND-NODE-L4 (UCS-C225-M6), ND-NODE-G5S (UCS-C225-M8), and ND-NODE-G5L (UCS-C225-M8).
- Key features that can be enabled per fabric include **Controller**, **Telemetry**, and **Orchestration**.
- Node and feature compatibility is subject to guidelines and limitations, such as restrictions on mixing node types and enabling all features on certain clusters.

Supported Node Types and Feature Overview

These node types are available for Nexus Dashboard:

- `ND-NODE-G5L` (UCS-C225-M8). The product ID of the 3-node cluster is `ND-CLUSTERG5L`.

- `ND-NODE-G5S` (UCS-C225-M8). The product ID of the 3-node cluster is `ND-CLUSTERG5S`.
- `ND-NODE-L4` (UCS-C225-M6). The product ID of the 3-node cluster is `ND-CLUSTER-L4`.
- `SE-NODE-G2` (UCS-C220-M5). The product ID of the 3-node cluster is `SE-CL-L3`.



Note Support for the `SE-NODE-G2` (UCS-C220-M5) node will come to an end in an upcoming Nexus Dashboard release. The EOL milestones for this appliance can be found here:

End-of-Sale and End-of-Life Announcement for the Cisco UCS M5 Rack Servers (C220 M5, C240 M5, C480 M5).

Future Nexus Dashboard release notes will contain details about the last release where this appliance will be supported.

If you require an RMA replacement for a UCS-**C220-M5** node, you can continue to receive a UCS-**C225-M6** node as the replacement while M6 remains supported. As M6 approaches its End-of-Life (EOL) milestone, you should plan to upgrade to a 4.x release to ensure continued support for M5 node RMAs after M6 reaches its EOL.

Until the M6 EOL milestone, M5 RMA replacements will continue to be fulfilled with M6 nodes. We recommend that you use this transition period to plan and complete your upgrade to the 4.x release.

In addition, in LAN deployments, these are the available features that you can leverage:

- **Controller** : Also referred to as Fabric Management. This feature is used to manage NX-OS and non-NX-OS switches (such as Catalyst, ASR, and so on). This includes creating any non-ACI fabric types, as well as performing software upgrades and creating new configurations on those fabrics.
- **Telemetry** : This feature provides telemetry functionality, similar to the functionality provided by Nexus Dashboard Insights in prior releases. You can enable and use the **Telemetry** feature when you create or edit a fabric through **Manage** > **Fabrics** .
- **Orchestration** : You can use the **Orchestration** feature through Nexus Dashboard to connect multiple ACI fabrics together, and consolidate and deploy tenants, along with network and policy configurations, across multiple ACI fabrics. You can enable and use the **Orchestration** feature when you add an ACI through **Admin** > **System Settings** > **Multi-cluster connectivity** > **Connect Cluster**.

For each fabric managed by Nexus Dashboard, you can enable the following features independently within the same Nexus Dashboard cluster or across a multi-cluster federated Nexus Dashboard deployment:

- Controller
- Telemetry
- Orchestration

Note Enabling all capabilities on a single cluster might not be available in some cluster deployment formats. For a quick reference of the supported form factors, scale, and cluster sizing requirements specific to your Nexus Dashboard deployment, see the Nexus Dashboard Capacity Planning tool.

Guidelines and limitations:

- You cannot mix the newer `ND-NODE-G5L` or `ND-NODE-G5S` (UCS-C225-M8) nodes in a cluster with the older `SE-NODE-G2` (UCS-C220-M5) and `ND-NODE-L4` (UCS-C225-M6) nodes.

  In addition, you can only have homogeneous `ND-NODE-G5L` or `ND-NODE-G5S` clusters. In other words:

  - A cluster containing a `ND-NODE-G5L` node can only have additional `ND-NODE-G5L` nodes in that cluster, and
  - A cluster containing a `ND-NODE-G5S` node can only have additional `ND-NODE-G5S` nodes in that cluster.

- The `ND-NODE-G5L` node is only supported for LAN deployments.
- The virtual form factor does not support all features in many cluster sizes and types, as described in the *Cisco Nexus Dashboard Verified Scalability Guide* .
- If a Nexus Dashboard node becomes unavailable, some Telemetry functionality may be impacted. Telemetry dynamic jobs such as Assurance, Bug Scan, and Log Collection, are blocked and will not run until the node is restored.
- Telemetry functionality is fully supported during a node failure when the deployment is operating at 50% or less of the supported scale for its form factor. For example, a 3vND data form factor supports up to 100 switches at full scale. If a deployment contains 50 switches and a Nexus Dashboard node becomes unavailable, Telemetry services will continue to operate normally. For deployments operating above 50% of the supported scale, Telemetry behavior during a node failure is provided on a best-effort basis, and some functionality may be degraded.

Note Telemetry scalability for IOS-XE platform switches is validated only with standalone (non-stacked) switches.

Windows 11 VDI client sizing

In Windows 11 virtual desktop infrastructure (VDI) environments, allocate sufficient client resources to provide stable Nexus Dashboard UI performance when multiple users access the UI concurrently. The following configurations have been tested and validated for each concurrent VDI user session.

These sizing requirements apply to VDI client sessions used to access the Nexus Dashboard UI.

Component	Minimum supported configuration per user	Recommended configuration per user
CPU	8 vCPUs	16 vCPUs
Processor speed	2.1 GHz or higher	2.1 GHz or higher

Component	Minimum supported configuration per user	Recommended configuration per user
Processor	Compatible 64-bit processor with two or more cores; Intel Xeon Silver 4116 or equivalent	Compatible 64-bit processor with two or more cores; Intel Xeon Silver 4116 or equivalent
Memory	8 GB RAM	32 GB RAM
Hardware acceleration	Enabled or disabled	Enabled
Display resolution	Higher than 1280 × 720 pixels	Higher than 1280 × 720 pixels