

Deployment Overview and Requirements

• Deployment Overview, on page 1

Deployment Overview

Nexus Dashboard Platform

Cisco Nexus Dashboard is a central management console for multiple data center sites and a common platform for hosting Cisco data center operation services, such as Insights, Orchestrator, and Fabric Controller. These services are available for all the data center sites and provide real time analytics, visibility, assurance for network policies and operations, as well as policy orchestration for the data center fabrics, such as Cisco ACI or Cisco NDFC.

Nexus Dashboard provides a common platform and modern technology stack for the above-mentioned micro-services-based applications, simplifying the life cycle management of the different modern applications and reducing the operational overhead to run and maintain these applications. It also provides a central integration point for external 3rd party applications with the locally hosted applications.

Note This document describes how to initially deploy a Nexus Dashboard cluster, enable one or more services, and onboard the fabrics to be managed by those services. After your cluster is up and running, see the Nexus Dashboard *configuration and operation articles* as well as the service-specific documentation for day-to-day operation.

Nexus Dashboard Services and Unified Installation in Release 3.1(1)

Nexus Dashboard is a standard appliance platform to build and deploy services that would allow you to consume all Nexus Dashboard products in a consistent and uniform manner. Nexus Dashboard services (such as Insights, Orchestrator, and Fabric Controller) can be enabled with the Nexus Dashboard platform to provide real time analytics, visibility, assurance for network policies and operations, as well as policy orchestration for the data center fabrics, such as Cisco ACI, Cisco NDFC, or Standalone NX-OS switches without a controller.

Prior to release 3.1(1), 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. Beginning with release 3.1(1), the platform and the individual services have been unified into a single image and can be deployed and enabled during the initial cluster configuration for a simpler, more streamlined experience.

Note While prior releases of Nexus Dashboard platform often supported multiple versions of the individual services, because of the unified installation, each Nexus Dashboard release supports a single specific version of each service, which you can choose to enable during cluster deployment.

In addition, the unified installation ensures that only the supported combinations of services can be enabled during cluster deployment, which prevents the possibility of unsupported service combinations. Detailed information about service cohosting is available in the Nexus Dashboard Cluster Sizing tool.

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.

This guide describes the initial deployment of the Nexus Dashboard software, which is common for physical as well as virtual or cloud 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. Keep in mind 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 two different Cisco UCS servers (`SE-NODE-G2` and `ND-NODE-L4`) for the cluster nodes, which can be mixed within the same cluster.

- 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
  - VMware ESX (`.ova`)

    Virtual form factor that allows you to deploy a Nexus Dashboard cluster using VMware ESX virtual machines.

  - Linux KVM (`.qcow2`)

Virtual form factor that allows you to deploy a Nexus Dashboard cluster using Linux KVM virtual machines.

- Public cloud
  - Amazon Web Services (`.ami`)

    Cloud form factor that allows you to deploy a Nexus Dashboard cluster using AWS instances.
  - Microsoft Azure (`.arm`)

    Cloud form factor that allows you to deploy a Nexus Dashboard cluster using Azure instances.



Note Not all services are supported on all form factors, and some form factors support only certain services. 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, services, 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 services and scale requirements, 3-node clusters can be extended with additional `secondary` nodes to support cohosting of services and higher scale. For physical clusters, you can also add up to 2 `standby` nodes for easy cluster recovery in case of a primary node failure.



Note

- Single-node deployments are supported for a limited number of services and cannot be extended to a 3-node cluster after the initial deployment.
- Single-node deployments do not support additional `secondary` or `standby` nodes.
- If you deploy a single-node cluster and want to extended it to a 3-node cluster or add `secondary` nodes, you will first need to redeploy it as a 3-node base cluster.
- For 3-node clusters, at least two `primary` nodes are required for the cluster to remain operational.

  If two `primary` nodes fail, the cluster will go offline and cannot be used until you recover it as described in the "Troubleshooting" article in the Nexus Dashboard *documentation library*.
- 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.

  An exception applies when a cluster size change is explicitly specified in Cisco documentation. For example, during an upgrade, you may need to change the cluster size for supportability reasons.

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