



Cisco Nexus Dashboard User Guide,
Release 2.3.x

Table of Contents

Trademarks	2
New and Changed Information	3
Latest Version of This Document	3
First Time Setup	4
Configuring Proxy	4
Adding Sites	5
Configuring Network Scale	8
Platform Overview	9
Hardware vs Software Stack	9
Services	9
Available Form Factors	10
Cluster Sizing Guidelines	11
Supported Services	11
Requirements and Guidelines	11
Network Time Protocol (NTP) and Domain Name System (DNS)	11
Nexus Dashboard External Networks	11
Nexus Dashboard Internal Networks	17
BGP Configuration and Persistent IPs	17
Communication Ports: Nexus Dashboard	18
Communication Ports: Nexus Dashboard Insights	20
Communication Ports: Nexus Dashboard Fabric Controller	20
Communication Ports: Nexus Dashboard Fabric Controller for SAN Deployments	24
Fabric Connectivity	25
Physical Node Cabling	25
Connecting via External Layer 3 Network	26
Connecting the Nodes Directly to Leaf Switches	28
GUI Overview	32
Navigation Bar and User Preferences	32
One View Page	33
Admin Console Page	33
Sites Page	35
Services Page	35
System Resources Pages	35
Operations Pages	36
Infrastructure Pages	37
Administrative Pages	37
Site Management	38
Adding Sites	38
Editing Sites	41
Deleting Sites	42
Services Management	43

Installing Services Using App Store	43
Installing Services Manually	44
Enabling Services	45
Updating Services	46
Disabling Services	46
Restarting Services	46
Uninstalling Services	46
Operations	48
Firmware Management (Cluster Upgrades)	48
Prerequisites and Guidelines	48
Adding Images	49
Upgrading the Cluster	50
Deleting Images	52
Tech Support	53
Backup and Restore	54
Creating Configuration Backups	54
Restoring Configuration	54
Event Analytics	55
Events	55
Audit Logs	56
Exporting Events	57
Infrastructure Management	58
Cluster Configuration	58
Persistent IP Addresses	61
Multi-Cluster Connectivity	64
Guidelines and Limitations	64
Connecting Multiple Clusters	65
Central Dashboard	67
Navigating Between Clusters	68
Disconnecting Clusters	69
Deploying Additional Physical Nodes	69
Prerequisites and Guidelines for Physical Nodes	69
Deploying Physical Nodes	71
Deploying Additional Virtual Nodes in VMware ESX	71
Prerequisites and Guidelines for ESX Nodes	72
Deploying ESX Node Using vCenter	73
Deploying ESX Node Directly in ESXi	79
Deploying Additional Virtual Nodes in Linux KVM	81
Prerequisites and Guidelines for KVM Nodes	81
Deploying KVM Nodes	82
Managing Worker Nodes	86
Adding Worker Nodes	86
Deleting a Worker node	87

Managing Standby Nodes	87
Adding Standby Nodes	88
Replacing Single Master Node with Standby Node	89
Replacing Two Master Nodes with Standby Nodes	90
Deleting Standby Nodes	92
Administrative	93
Roles and Permissions	93
Nexus Dashboard and Orchestrator Roles	93
Nexus Dashboard Insights	94
Nexus Dashboard Fabric Controller Roles	94
Choosing Default Authentication Domain	97
Remote Authentication	98
Configuring Remote Authentication Server	98
Adding LDAP as Remote Authentication Provider	100
Adding Radius or TACACS as Remote Authentication Provider	102
Validating Remote User Logins	103
Editing Remote Authentication Domains	104
Deleting Remote Authentication Domains	104
Multi-Factor Authentication	105
Configuring Okta Account as MFA Provider	105
Configuring MFA Client	109
Adding Okta as Remote Authentication Provider	111
Logging In To Nexus Dashboard Using MFA	112
Users	113
Adding Local Users	113
Editing Local Users	113
Security	114
Security Configuration	114
Security Domains	115
Validating Peer Certificates	116
Cisco Intersight	121
Configuring Device Connector Settings	121
Target Claim	122
Unclaiming the Device	124
Troubleshooting	125
Useful Commands	125
Upgrading CIMC	127
Manual Cluster Upgrades	132
Re-Imaging Nodes	133
Installing Nexus Dashboard Using Remotely-Hosted Image	134
Rebuilding Existing Cluster	137
AppStore Errors	138
Event Export	138

Factory Reset	139
Changing Node IP Addresses	139
Cluster Configuration Errors	140
Two-Factor Authentication (2FA) Not Prompting for Login Info	140
Red Hat Enterprise Linux (RHEL) Deployments	140
Unable to Connect to Site After APIC Configuration Import	141
Re-Adding Same Master Node to Physical Cluster	141
Replacing a Single Virtual Master Node Without a Standby Node	142
Replacing a Single Physical Master Node Without a Standby Node	143
Replacing Worker or Standby Nodes	144
Initial Cluster Bootstrap Issues	144
Multi-Cluster Connectivity Issues	146
Non-Primary Cluster Unable to Reconnect	146
Non-Primary Cluster Redeployed with Older Version	147
Generating Private Key, Creating CSR, and Obtaining CA-Signed Certificate	147
Generating Private Key and Self-Signed Certificate	149
Updating NDO Configuration After Replacing Switch Devices Managed by NDFC	152
Replacing a Core or Route Server (RS) Device	152
Replacing a Leaf Switch	152
Replacing Border Gateway (BGW) Devices	153

First Published: 2023-01-31

Last Modified: 2023-05-05

Americas Headquarters

Cisco Systems, Inc.

170 West Tasman Drive

San Jose, CA 95134-1706

USA

<http://www.cisco.com>

Tel: 408 526-4000

800 553-NETS (6387)

Fax: 408 527-0883

Trademarks

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <http://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

© 2017–2022 Cisco Systems, Inc. All rights reserved.

New and Changed Information

The following table provides an overview of the significant changes to the organization and features in this guide from the release the guide was first published to the current release. The table does not provide an exhaustive list of all changes made to the guide.

Table 1. Latest Updates

Release	Change	Where Documented
2.3.1	First release of this document	–

Latest Version of This Document


This document is available from your Nexus Dashboard GUI as well as online at www.cisco.com. For the latest version of this document, visit [Nexus Dashboard User Guide](#).

First Time Setup

When you first log in to your new Nexus Dashboard cluster, you will be able to configure basic settings from the first time setup wizard.

1. In the **What's New** screen, click **Begin Setup**.
2. Add sites and configure proxy as described in one of the following two sections.

If at any time you exit out of the first time setup wizard, you can return to it using the intent menu in your Nexus Dashboard's **Admin Console**:



If you have already finished the first time setup wizard, you can skip the next two sections.

Otherwise, you will be prompted for the following two things:

- Configure proxy—allows you to provide a proxy server, which can be used to connect to the Internet.

This may be required when adding sites managed by Cloud Network Controller, in which case it must be configured prior to onboarding those sites.

- Add sites—allows you to onboard one or more fabrics, which you will use with the services running in your cluster.

Configuring Proxy

In certain deployment scenarios, such as with a combination of on-premises and cloud sites and the Nexus Dashboard cluster deployed inside a corporate network, you may have to access the internet and the cloud sites through a proxy.



This release supports adding a single proxy server.

To add a proxy server:

1. In the **Proxy Configuration** tile, click **Begin**.
2. In the **Setup - Proxy Configuration** page, click **+Add Server**.
 - a. From the **Type** dropdown, select the type of traffic that you want to be proxied.
 - b. In the **Server** field, provide the full address for the proxy server.

You can also choose to provide the port, for example <http://proxy.company.com:80>.

- c. If the server requires login credentials, provide the **Username** and **Password**.
3. (Optional) Click **Add Ignore Host** to provide any hosts that will ignore the proxy.

You can add one or more hosts with which the cluster will communicate directly bypassing the proxy.

Adding Sites

Before you begin

- Fabric connectivity must be already configured.
- If adding a Cisco APIC or Cloud Network Controller site, the site must be running Release 4.2(4) or later.
- If adding a Cisco APIC site, EPG/L3Out for Cisco Nexus Dashboard data network IP connectivity must be pre-configured.

Refer to [Fabric Connectivity](#) for more information.

- If adding a Cisco APIC site and planning to deploy Cisco NIR application:
 - IP connectivity from Cisco Nexus Dashboard to Cisco APIC Inband IP over data network must be configured.
 - IP connectivity from Cisco Nexus Dashboard to the leaf nodes and spine nodes in-band IPs must be configured.
- If adding a Cisco NDFC site:
 - The site must be running Release 11.5(1) or later.
 - You must configure Layer 3 connectivity to the fabric and switches.
 - If you cluster is deployed in AWS or Azure, you must configure inbound rules on the data interface.

This is typically done during initial cluster deployment and described in detail in the [Cisco Nexus Dashboard Deployment Guide](#).

To add a site:

1. In the **Add Sites** tile, click **Begin**.
2. In the **Setup - Add Sites** page, click **Add Site**.
3. Select the type of site you want to add.



While Cisco Nexus Dashboard supports on-boarding all three types of fabrics, for specific fabric types and versions compatible with your services, see the [Services Compatibility Matrix](#).

- **ACI**—for on-premises ACI sites managed by Cisco APIC
- **Cloud Network Controller**—for cloud sites managed by Cisco Cloud Network Controller
- **NDFC**—for on-premises sites managed by Cisco NDFC

4. Provide the site's information.

a. If adding an **ACI** site, provide the following:

- **Site Name**—used throughout the Nexus Dashboard GUI when referring to this site.
- **Host Name/IP Address**—used to communicate with the Cisco APIC.

If you will use the site with Nexus Dashboard Orchestrator service only, you can provide either the in-band or out-of-band IP address of the APIC. If you will use the site with Nexus Dashboard Insights as well, you must provide the in-band IP address.



When providing the address, do not include the protocol (**http://** or **https://**) as part of the URL string or site addition will fail.

- **User Name** and **Password**—login credentials for a user with **admin** privileges on the site you are adding.
- (Optional) **Login Domain**—if you leave this field empty, the site's local login is used.
- (Optional) **Validate Peer Certificate**—allows Nexus Dashboard to verify that the certificates of hosts to which it connects (such as site controllers) are valid and are signed by a trusted Certificate Authority (CA).



You must have the certificate for this site already imported into your Nexus Dashboard before you can add a site using this option. If you have not yet added the certificates, cancel the **Add Site** wizard and follow the instructions described in [Validating Peer Certificates](#) first; then after you have imported the certificates, add the site as described here. If you enable the **Verify Peer Certificate** option but don't import the valid certificate, site onboarding will fail.

- (Optional) **In-Band EPG**—required when connecting to an ACI fabric via an EPG and bridge domain. For more information on fabric connectivity, see [Fabric Connectivity](#).

If you plan to use this site with the Nexus Dashboard Insights service, you must provide the node management In-Band EPG.

b. If adding a **Cloud Network Controller** site, provide the following:

- **Site Name**—used throughout the Nexus Dashboard GUI when referring to this site.
- **Host Name/IP Address**—used to communicate with the Cloud Network Controller.



When providing the address, do not include the protocol ([http://](#) or [https://](#)) as part of the URL string or site addition will fail.

- **User Name** and **Password**—login credentials for a user with **admin** privileges on the site you are adding.
- (Optional) **Login Domain**—if you leave this field empty, the site’s local login is used.
- (Optional) **Validate Peer Certificate**—allows Nexus Dashboard to verify that the certificates of hosts to which it connects (such as site controllers) are valid and are signed by a trusted Certificate Authority (CA).



You must have the certificate for this site already imported into your Nexus Dashboard before you can add a site using this option. If you have not yet added the certificates, cancel the **Add Site** wizard and follow the instructions described in [Validating Peer Certificates](#) first; then after you have imported the certificates, add the site as described here. If you enable the **Verify Peer Certificate** option but don’t import the valid certificate, site onboarding will fail.

- (Optional) **Enable Proxy**—enable this setting if your cloud site is reachable via a proxy.



Proxy must be already configured in your Nexus Dashboard’s cluster settings. For more information, see [Configuring Proxy](#).

c. If adding an **NDFC** site, provide the following:

- **Host Name/IP Address**—used to communicate with the Cisco NDFC.

This must be the in-band IP address of NDFC.



When providing the address, do not include the protocol ([http://](#) or [https://](#)) as part of the URL string or site addition will fail.

- **User Name** and **Password**—login credentials for a user with **admin** privileges on the site you are adding.
- (Optional) **Login Domain**—if you leave this field empty, the site’s local login is used.
- (Optional) **Validate Peer Certificate**—allows Nexus Dashboard to verify that the certificates of hosts to which it connects (such as site controllers) are valid and are signed by a trusted Certificate Authority (CA).



You must have the certificate for this site already imported into your Nexus Dashboard before you can add a site using this option. If you have not yet added the certificates, cancel the **Add Site** wizard and follow the instructions described in [Validating Peer Certificates](#) first; then after you have imported the certificates, add the site as described here. If you enable the **Verify Peer Certificate** option but don’t import the valid certificate, site onboarding will fail.

- **Sites**—click **Select Sites** to select the NDFC fabrics managed by the controller you provided.
5. Click **Add** to finish adding the site.
 6. (Optional) Click on the **Geographical Location** map to specify where the site is located.
 7. (Optional) Repeat these steps for any additional sites.

Configuring Network Scale

Starting with Release 2.2(1), you can configure the target scale for your services and the Nexus Dashboard cluster will automatically allocate the appropriate amount of resources and limits.

To configure network scale:

1. In the **Network Scale** tile, click **Begin**.
2. In the **Setup - Network Scale** page, provide the required information.



Modifying the network scale requires a restart of your services for the changes to be applied.

- a. In the **Number of Sites** field, provide the target number of sites for your deployment that this Nexus Dashboard cluster will manage.

By default, network scale is set to 10 sites.

- b. In the **Number of Fabric Nodes** field, provide the target number of switch nodes for your deployment
- c. From the **Flows per second** drop-down menu, select the target number of flows for your Nexus Dashboard Insights service.

Platform Overview

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 Nexus Dashboard Insights and Nexus Dashboard Orchestrator. 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.

Nexus Dashboard cluster typically consists of 1 or 3 **master** nodes. For 3-node clusters, you can also provision a number of **worker** nodes to enable horizontal scaling and **standby** nodes for easy cluster recovery in case of a **master** node failure. For maximum number of **worker** and **standby** nodes supported in this release, see the "Verified Scalability Limits" sections of the [Cisco Nexus Dashboard Release Notes](#). For more information about extending your cluster with additional nodes, see [Infrastructure Management](#).

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 platform" specifically to refer to the hardware and "Nexus Dashboard" to refer to the software stack and the GUI console.

This guide describes how to use the Nexus Dashboard; for hardware installation, see [Nexus Dashboard Hardware Setup Guide](#) and for deployment planning and Nexus Dashboard software installation, see the [Nexus Dashboard Deployment Guide](#).

Services

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. You can subscribe and consume services like Insights, Orchestrator, Fabric Controller, and Data Broker with the Nexus Dashboard platform providing the necessary capacity and life cycle management operations for these services.

Typically, the Nexus Dashboard platform is shipped with only the software required for managing the lifecycle of these services, but no actual services are packaged with the appliance. If you allow public network connectivity from your data centers, you can download and install the services with a few clicks. However, without public network connectivity, you will need to manually download these services, upload them to the platform, and perform installation operations before you can use them.

If you are ordering the physical Nexus Dashboard servers, you have the option to choose some services to be pre-installed on the hardware before it is shipped to you. For more information, see the Nexus Dashboard Ordering Guide. Note that if you are deploying the virtual or cloud form factors of

the Nexus Dashboard, there are no changes to service installation and you will need to deploy the services separately after the cluster is ready.

If you are ordering the physical Nexus Dashboard servers, you have the option to choose Nexus Dashboard Insights and Nexus Dashboard Orchestrator services to be pre-installed on the hardware before it is shipped to you. For more information, see the [Nexus Dashboard Ordering Guide](#). Note that if you are deploying the virtual or cloud form factors of the Nexus Dashboard, you will need to deploy the services separately after the cluster is ready.

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 different form factors within the same cluster is not supported.



Not all services are supported on all form factors. When planning your deployment, ensure to check the [Nexus Dashboard Cluster Sizing](#) tool for form factor and cluster size requirements.

- Cisco Nexus Dashboard physical appliance (.iso)

This form factor refers to the original physical appliance hardware that you purchased with the Cisco 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 original Cisco Nexus Dashboard platform hardware is described in [Cisco Nexus Dashboard Hardware Setup Guide](#).

- VMware ESX (.ova)

Virtual form factor that allows you to deploy a Nexus Dashboard cluster using VMware ESX virtual machines with one of two resource 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

- Linux KVM (.qcow2)

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

- 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.

- In an existing Red Hat Enterprise Linux (RHEL) system

Beginning with Release 2.2(1), you can run Nexus Dashboard node in an existing Red Hat Enterprise Linux server.

Cluster Sizing Guidelines

As mentioned previously, Nexus Dashboard cluster is first deployed using 1 or 3 master nodes. Depending on the type and number of services you choose to run, you may be required to deploy additional **worker** nodes in your cluster after the initial deployment. For cluster sizing information and recommended number of nodes based on specific use cases, see the [Nexus Dashboard Capacity Planning](#) tool.



Single-node clusters are supported for a limited number of services and cannot be extended to a 3-node cluster after the initial deployment.

Only 3-node clusters support additional worker nodes.

If you deploy a single-node cluster and want to extend it to a 3-node cluster or add worker nodes, you will need to redeploy it as a base 3-node cluster.

For 3-node clusters, at least 2 master nodes are required for the cluster to remain operational. If 2 master nodes fail, the cluster will go offline and cannot be used until you recover it as described in this guide.

Adding worker nodes to your cluster is described in [Managing Worker Nodes](#).

Adding standby nodes to your cluster is described in [Managing Standby Nodes](#).

Supported Services

For the full list of supported applications and the associated compatibility information, see the [Data Center Networking Services Compatibility Matrix](#).

Requirements and Guidelines

Network Time Protocol (NTP) and Domain Name System (DNS)

The Nexus Dashboard nodes require valid DNS and NTP servers for all deployments and upgrades.

Lack of valid DNS connectivity (such as if using an unreachable or a placeholder IP address) can prevent the system from deploying or upgrading successfully.



Nexus Dashboard acts as both a DNS client and resolver. It uses an internal Core DNS server which acts as DNS resolver for internal services. It also acts as a DNS client to reach external hosts within the intranet or the Internet, hence it requires an external DNS server to be configured.

Additionally, Nexus Dashboard does not support DNS servers with wildcard records.

Nexus Dashboard External Networks

Cisco Nexus Dashboard is deployed as a cluster, connecting each service node to two networks. When first configuring Nexus Dashboard, you will need to provide two IP addresses for the two Nexus Dashboard interfaces—one connected to the Data Network and the other to the Management

Network.

Individual services installed in the Nexus Dashboard may utilize the two networks for additional purposes, so we recommend consulting the specific service’s documentation in addition to this document for your deployment planning.

Data Network	Management Network
<ul style="list-style-type: none">▪ Nexus Dashboard node clustering▪ Application to application communication▪ Nexus Dashboard nodes to Cisco APIC nodes communication <p>For example, the network traffic for the Nexus Dashboard Insights service.</p>	<ul style="list-style-type: none">▪ Accessing the Nexus Dashboard GUI▪ Accessing the Nexus Dashboard CLI via SSH▪ DNS and NTP communication▪ Nexus Dashboard firmware upload▪ Cisco DC App Center (AppStore) <p>If you want to use the Nexus Dashboard App Store to install applications as described in Services Management, the https://dcappcenter.cisco.com page must be reachable via the Management Network.</p> <ul style="list-style-type: none">▪ Intersight device connector

The two networks have the following requirements:

- For all new Nexus Dashboard deployments, the management network and data network must be in different subnets.
- For physical clusters, the management network must provide IP reachability to each node’s CIMC via TCP ports 22/443.

Nexus Dashboard cluster configuration uses each node’s CIMC IP address to configure the node.

- For Nexus Dashboard Insights service, the data network must provide IP reachability to the in-band network of each fabric and of the APIC.
- For Nexus Dashboard Insights and AppDynamics integration, the data network must provide IP reachability to the AppDynamics controller.
- For Nexus Dashboard Orchestrator service, the data network can have in-band and/or out-of-band IP reachability for Cisco APIC sites but must have in-band reachability for Cisco NDFC sites.
- The data network interface requires a minimum MTU of 1500 to be available for the Nexus Dashboard traffic.

Higher MTU can be configured if desired.

- The following table summarizes service-specific requirements for the management and data networks.



Changing the data subnet requires redeploying the cluster, so we recommend using a larger subnet than the bare minimum required by the nodes and services to account for any additional services in the future. In addition to the requirements listed in this section, ensure that you consult the Release Notes for

the specific service you plan to deploy.

Allocating persistent IP addresses for both Layer 2 and Layer 3 connectivity is done after the cluster is deployed using the External Service Pools configuration in the UI, as described in the [Cisco Nexus Dashboard User Guide](#).

We recommend consulting the specific service's documentation for any additional requirements and caveats related to persistent IP configuration.

Nexus Dashboard Service	Management Interface	Data Interface	Total Number of Persistent IPs
Nexus Dashboard Orchestrator	Layer 3 adjacent	Layer 3 adjacent	N/A
Nexus Dashboard Insights without SFLOW/NetFlow (ACI fabrics)	Layer 3 adjacent	Layer 3 adjacent	N/A
Nexus Dashboard Insights without SFLOW/NetFlow (NDFC fabrics)	Layer 3 adjacent	Layer 2 adjacent	6 IPs in data interface network if using IPv4 7 IPs in data interface network if using IPv6
Nexus Dashboard Insights with SFLOW/NetFlow (ACI or NDFC fabrics)	Layer 3 adjacent	Layer 2 adjacent	6 IPs in data interface network

Nexus Dashboard Service	Management Interface	Data Interface	Total Number of Persistent IPs
Nexus Dashboard Fabric Controller, Release 12.0(x)	Layer 2 adjacent	Layer 2 adjacent	<p>If LAN Device Management Connectivity is set to Management (default):</p> <ul style="list-style-type: none"> • 2 IPs in the management network for SNMP/Syslog and SCP services • If EPL is enabled, 1 additional IP in the data network for each fabric • If IP Fabric for Media is enabled, 1 additional IP in the management network for telemetry <p>If LAN Device Management Connectivity is set to Data:</p> <ul style="list-style-type: none"> • 2 IPs in the data network for SNMP/Syslog and SCP services • If EPL is enabled, 1 additional IP in the data network for each fabric • If IP Fabric for Media is enabled, 1 additional IP in the data network for telemetry

Nexus Dashboard Service	Management Interface	Data Interface	Total Number of Persistent IPs
Nexus Dashboard Fabric Controller, Release 12.1(1) and latter	Layer 2 or Layer 3 adjacent	Layer 2 or Layer 3 adjacent	<p>When operating in Layer 2 mode with LAN deployment type and LAN Device Management Connectivity set to Management (default)</p> <ul style="list-style-type: none"> ▪ 2 IPs in the management network for SNMP/Syslog and SCP services ▪ If EPL is enabled, 1 additional IP in the data network for each fabric ▪ If IP Fabric for Media is enabled, 1 additional IP in the management network for telemetry <p>When operating in Layer 2 mode with LAN deployment type and LAN Device Management Connectivity set to Data:</p> <ul style="list-style-type: none"> ▪ 2 IPs in the data network for SNMP/Syslog and SCP services ▪ If EPL is enabled, 1 additional IP in the data network for each fabric ▪ If IP Fabric for Media is enabled, 1 additional IP in the data network for telemetry

Nexus Dashboard Service	Management Interface	Data Interface	Total Number of Persistent IPs
Nexus Dashboard Fabric Controller, Release 12.1(1) and latter	Layer 2 or Layer 3 adjacent	Layer 2 or Layer 3 adjacent	<p>When operating in Layer 3 mode with LAN deployment type:</p> <ul style="list-style-type: none"> • LAN Device Management Connectivity must be set to Data • 2 IPs for SNMP/Syslog and SCP services • If EPL is enabled, 1 additional IP in the data network for each fabric • All persistent IPs must be part of a separate pool that must not overlap with the management or data subnets <p>For more information about Layer 3 mode for persistent IPs, see the Persistent IP Addresses</p> <p>When operating in Layer 3 mode with SAN Controller deployment type:</p> <ul style="list-style-type: none"> • 1 IP for SSH • 1 IP for SNMP/Syslog • 1 IP for SAN Insights functionality <p>IP Fabric for Media mode are not supported in Layer 3 mode</p>
Nexus Dashboard Data Broker	Layer 3 adjacent	N/A	N/A

- Connectivity between the nodes is required on both networks with the following additional round trip time (RTT) requirements.



You must always use the lowest RTT requirement when deploying the Nexus Dashboard cluster and applications. For example, if you plan to co-host the Insights and Orchestrator services, site connectivity RTT must not exceed 50ms.

Application	Connectivity	Maximum RTT
Nexus Dashboard cluster	Between nodes	150 ms
Nexus Dashboard Orchestrator	Between nodes	150 ms
	To sites	500 ms
Nexus Dashboard Insights	Between nodes	50 ms
	To sites	50 ms

Application	Connectivity	Maximum RTT
Nexus Dashboard Fabric Controller	Between nodes	50 ms
	To sites	50 ms
Nexus Dashboard Data Broker	Between nodes	150 ms
	To sites	500 ms

Nexus Dashboard Internal Networks

Two additional internal networks are required for communication between the containers used by the Nexus Dashboard:

- Application overlay is used for applications internally within Nexus Dashboard

Application overlay must be a /16 network and a default value is pre-populated during deployment.

- Service overlay is used internally by the Nexus Dashboard.

Service overlay must be a /16 network and a default value is pre-populated during deployment.

If you are planning to deploy multiple Nexus Dashboard clusters, they can use the same Application and Service subnets.



Communications between containers deployed in different Nexus Dashboard nodes is VXLAN-encapsulated and uses the data interfaces IP addresses as source and destination. This means that the Application Overlay and Service overlay addresses are never exposed outside the data network and any traffic on these subnets is routed internally and does not leave the cluster nodes.

For example, if you had another service (such as DNS) on the same subnet as one of the Overlay networks, you would not be able to access it from your Nexus Dashboard as the traffic on that subnet would never be routed outside the cluster. As such, when configuring these networks, ensure that they are unique and do not overlap with any existing networks or services you may need to access from the Nexus Dashboard cluster nodes.

For the same reason, we recommend not using **169.254.0.0/16** (the Kubernetes br1 subnet) for the App or Service subnets.

BGP Configuration and Persistent IPs

Previous releases of Nexus Dashboard allowed you to configure one or more persistent IP addresses for services (such as Nexus Dashboard Insights) that require retaining the same IP addresses even in case they are relocated to a different Nexus Dashboard node. However, in those releases, the persistent IPs had to be part of the management and data subnets and the feature could be enabled only if all nodes in the cluster were part of the same Layer 3 network. Here the services used Layer 2 mechanisms like Gratuitous ARP or Neighbor Discovery to advertise the persistent IPs within it's Layer 3 network

Beginning with Release 2.2(1), the Persistent IPs feature is supported even if you deploy the cluster nodes in different Layer 3 networks. In this case, the persistent IPs are advertised out of each node's

data links via BGP, which we refer to as "Layer 3 mode". The IPs must also be part of a subnet that is not overlapping with any of the nodes' management or data subnets. If the persistent IPs are outside the data and management networks, this feature will operate in Layer 3 mode by default; if the IPs are part of those networks, the feature will operate in Layer 2 mode.

BGP can be enabled during cluster deployment or from the Nexus Dashboard GUI after the cluster is up and running.

If you plan to enable BGP and use the persistent IP functionality, you must:

- Ensure that the peer routers exchange the advertised persistent IPs between the nodes' Layer 3 networks.
- Choose to enable BGP at the time of the cluster deployment as described in the subsequent sections or enable it afterwards in the Nexus Dashboard GUI as described in [Persistent IP Addresses](#).
- Ensure that the persistent IP addresses you allocate do not overlap with any of the nodes' management or data subnets.

Communication Ports: Nexus Dashboard

The following ports are required by the Nexus Dashboard cluster.



All services use TLS or mTLS with encryption to protect data privacy and integrity over the wire.

Table 2. Nexus Dashboard Communication Ports (Management Network)

Service	Port	Protocol	Direction	Connection
ICMP	ICMP	ICMP	In/Out	Other cluster nodes, CIMC, default gateway
SSH	22	TCP	In/Out	CLI and CIMC of the cluster nodes
TACACS	49	TCP	Out	TACACS server
DNS	53	TCP/UDP	Out	DNS server
HTTP	80	TCP	Out	Internet/proxy
NTP	123	UDP	Out	NTP server
HTTPS	443	TCP	In/Out	UI, other clusters (for multi-cluster connectivity), fabrics, Internet/proxy
LDAP	389 636	TCP	Out	LDAP server
Radius	1812	TCP	Out	Radius server

Service	Port	Protocol	Direction	Connection
KMS	9880	TCP	In/Out	Other cluster nodes and ACI fabrics
Infra-Service	30012 30021 30500-30600	TCP/UDP	In/Out	Other cluster nodes

Table 3. Nexus Dashboard Communication Ports (Data Network)

Service	Port	Protocol	Direction	Connection
SSH	22	TCP	Out	In-band of switches and APIC
HTTPS	443	TCP	Out	In-band of switches and APIC/NDFC
VXLAN	4789	TCP	In/Out	Other cluster nodes
KMS	9880	TCP	In/Out	Other cluster nodes and ACI fabrics
Infra-Service	3379 3380 8989 9090 9969 9979 9989 15233 30002-30006 30009-30010 30012 30014-30015 30018-30019 30025 30027	TCP	In/Out	Other cluster nodes
Kafka	30001	TCP	In/Out	In-band of switches and APIC/NDFC
Infra-Service	30016 30017	TCP/UDP	In/Out	Other cluster nodes
Infra-Service	30500-30600	TCP/UDP	In/Out	Other cluster nodes

Communication Ports: Nexus Dashboard Insights

In addition to the ports required by the Nexus Dashboard cluster nodes, which are listed above, the following ports are required by the Nexus Dashboard Insights service.

Table 4. Nexus Dashboard Insights Communication Ports (Data Network)

Service	Port	Protocol	Direction	Connection
Show Techcollection	2022	TCP	In/Out	In-band of switches and APIC/NDFC
Flow Telemetry	5640-5671	UDP	In	In-band of switches
TAC Assist	8884	TCP	In/Out	Other cluster nodes
KMS	9989	TCP	In/Out	Other cluster nodes and ACI fabrics
SW Telemetry	5695 30000 30570 57500	TCP	In/Out	Other cluster nodes

Communication Ports: Nexus Dashboard Fabric Controller

In addition to the ports required by the Nexus Dashboard (ND) cluster nodes, the following ports are required by the Nexus Dashboard Fabric Controller (NDFC) service.



The following ports apply to the Nexus Dashboard management network and/or data network interfaces depending on which interface provides IP reachability from the NDFC service to the switches.

Table 5. Nexus Dashboard Fabric Controller Communication Ports

Service	Port	Protocol	Direction	Connection
SSH	22	TCP	Out	SSH is a basic mechanism for accessing devices.
SCP	22	TCP	Out	SCP clients archiving NDFC backup files to remote server.
SMTP	25	TCP	Out	SMTP port is configurable through NDFC's Server Settings menu. This is an optional feature.

Service	Port	Protocol	Direction	Connection
DHCP	67	UDP	In	If NDFC local DHCP server is configured for Bootstrap/POAP purposes. This applies to LAN deployments only. NOTE: When using NDFC as a local DHCP server for POAP purposes, all ND master node IPs must be configured as DHCP relays. Whether the ND nodes' management or data IPs are bound to the DHCP server is determined by the LAN Device Management Connectivity in the NDFC Server Settings.
DHCP	68	UDP	Out	
SNMP	161	TCP/UDP	Out	SNMP traffic from NDFC to devices.
HTTPS/HTTP (NX-API)	443/80	TCP	Out	NX-API HTTPS/HTTP client connects to device NX-API server on port 443/80, which is also configurable. NX-API is an optional feature, used by limited set of NDFC functions. This applies to LAN deployments only.
HTTPS (vCenter, Kubernetes, OpenStack, Discovery)	443	TCP	Out	NDFC provides an integrated host and physical network topology view by correlating the information obtained from registered VMM domains, such as VMware vCenter or OpenStack, as well as container orchestrators, such as Kubernetes. This is an optional feature

The following ports apply to the External Service IPs, also known as persistent IPs, used by some of the NDFC services. These External Service IPs may come from the Nexus Dashboard management subnet pool or the data subnet pool depending on the configured settings.

Table 6. Nexus Dashboard Fabric Controller Persistent IP Ports

Service	Port	Protocol	Direction	Connection
SCP	22	TCP	In	SCP is used by various features to transfer files between devices and the NDFC service. The NDFC SCP service serves as the SCP server for both downloads and uploads. SCP is also used by the POAP client on the devices to download POAP-related files. The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the LAN Device Management Connectivity setting in the NDFC Server Settings.

Service	Port	Protocol	Direction	Connection
TFTP (POAP)	69	TCP	In	<p>Only used for device zero-touch provisioning via POAP, where devices can send (limited jailed write-only access to NDFC) basic inventory information to NDFC to start secure POAP communication. NDFC Bootstrap or POAP can be configured for TFTP or HTTP/HTTPS.</p> <p>The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the LAN Device Management Connectivity setting in the NDFC Server Settings.</p> <p>This applies to LAN deployments only.</p>
HTTP (POAP)	80	TCP	In	<p>Only used for device zero-touch provisioning via POAP, where devices can send (limited jailed write-only access to NDFC) basic inventory information to NDFC to start secure POAP communication. NDFC Bootstrap or POAP can be configured for TFTP or HTTP/HTTPS.</p> <p>The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the LAN Device Management Connectivity setting in the NDFC Server Settings.</p> <p>This applies to LAN deployments only.</p>
BGP	179	TCP	In/Out	<p>For Endpoint Locator, per fabric where it is enabled, an EPL service is spawned with its own persistent IP. This service is always associated with the Nexus Dashboard data interface. NDFC EPL service peers with the appropriate BGP entity (typically BGP Route-Reflectors) on the fabric to get BGP updates needed to track endpoint information.</p> <p>This feature is only applicable for VXLAN BGP EVPN fabric deployments.</p> <p>This applies to LAN deployments only.</p>

Service	Port	Protocol	Direction	Connection
HTTPS (POAP)	443	TCP	In	Secure POAP is accomplished via the NDFC HTTPS Server on port 443. The HTTPS server is bound to the SCP-POAP service and uses the same persistent IP assigned to that pod. The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the LAN Device Management Connectivity setting in the NDFC Server Settings. This applies to LAN deployments only.
Syslog	514	UDP	In	When NDFC is configured as a Syslog server, Syslogs from the devices are sent out toward the persistent IP associated with the SNMP-Trap/Syslog service pod The SNMP-Trap-Syslog service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the LAN Device Management Connectivity setting in the NDFC Server Settings
SCP	2022	TCP	Out	Transport tech-support file from persistent IP of NDFC POAP-SCP pod to a separate ND cluster running Nexus Dashboard Insights. The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the LAN Device Management Connectivity setting in the NDFC Server Settings
SNMP Trap	2162	UDP	In	SNMP traps from devices to NDFC are sent out toward the persistent IP associated with the SNMP-Trap/Syslog service pod. The SNMP-Trap-Syslog service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the LAN Device Management Connectivity setting in the NDFC Server Settings

Service	Port	Protocol	Direction	Connection
GRPC (Telemetry)	33000	TCP	In	SAN Insights Telemetry Server which receives SAN data (such as storage, hosts, flows, and so on) over GRPC transport tied to NDFC Persistent IP. This is enabled on SAN deployments only.
GRPC (Telemetry)	50051	TCP	In	Information related to multicast flows for IP Fabric for Media deployments as well as PTP for general LAN deployments is streamed out via software telemetry to a persistent IP associated with a NDFC GRPC receiver service pod. This is enabled on LAN and Media deployments only.

Communication Ports: Nexus Dashboard Fabric Controller for SAN Deployments

Nexus Dashboard Fabric Controller can be deployed on a single-node or 3-node Nexus Dashboard cluster. The following ports are required for NDFC SAN deployments on single-node clusters.

Table 7. Nexus Dashboard Fabric Controller Ports for SAN Deployments on Single-Node Clusters

Service	Port	Protocol	Direction	Connection
SSH	22	TCP	Out	SSH is a basic mechanism for accessing devices.
SCP	22	TCP	Out	SCP clients archiving NDFC backup files to remote server.
SMTP	25	TCP	Out	SMTP port is configurable through NDFC's Server Settings menu. This is an optional feature.
SNMP	161	TCP/UDP	Out	SNMP traffic from NDFC to devices.
HTTPS (vCenter, Kubernetes, OpenStack, Discovery)	443	TCP	Out	NDFC provides an integrated host and physical network topology view by correlating the information obtained from registered VMM domains, such as VMware vCenter or OpenStack, as well as container orchestrators, such as Kubernetes. This is an optional feature.

The following ports apply to the External Service IPs, also known as Persistent IPs, used by some of the NDFC services. These External Service IPs may come from the Nexus Dashboard management subnet pool or the data subnet pool depending on the configured settings.

Table 8. Nexus Dashboard Fabric Controller Persistent IP Ports for SAN Deployments on Single-Node Clusters

Service	Port	Protocol	Direction	Connection
SCP	22	TCP	In	SCP is used by various features to transfer files between devices and the NDFC service. The NDFC SCP service functions for both downloads and uploads.
Syslog	514	UDP	In	When NDFC is configured as a Syslog server, syslogs from the devices are sent out towards the persistent IP associated with the SNMP-Trap/Syslog service pod. The SNMP-Trap-Syslog service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the LAN Device Management Connectivity setting in the NDFC Server Settings.
SNMP Trap	2162	UDP	In	SNMP traps from devices to NDFC are sent out toward the persistent IP associated with the SNMP-Trap/Syslog service pod. The SNMP-Trap-Syslog service in NDFC has a persistent IP that is associated with either the management or data subnet.
GRPC (Telemetry)	33000	TCP	In	SAN Insights Telemetry Server which receives SAN data (such as storage, hosts, flows, and so on) over GRPC transport tied to NDFC Persistent IP. This is enabled on SAN deployments only.

Fabric Connectivity

You can connect the Nexus Dashboard cluster to your fabrics in two ways:

- The Nexus Dashboard cluster connected to the fabric via a Layer 3 network.
- The Nexus Dashboard nodes connected to the leaf switches as typical hosts.

For Cisco Cloud Network Controller fabrics, you will need to connect via a Layer 3 network.

Physical Node Cabling

If you deployed a virtual or cloud form factor cluster, you can skip this section.

The following figure shows the Nexus Dashboard physical node interfaces:

- **eth1-1** and **eth1-2** must be connected to the Management network
- **eth2-1** and **eth2-2** must be connected to the Data network




Figure 1. Node Connectivity

The interfaces are configured as Linux bonds (one for the data interfaces and one for the management interfaces) running in active-standby mode. All interfaces must be connected to individual host ports, PortChannel or vPC are not supported.

When Nexus Dashboard nodes are connected to Cisco Catalyst switches, packets are tagged with vlan0 if no VLAN is specified. In this case, you must add `switchport voice vlan dot1p` command to the switch interfaces where the nodes are connected to ensure reachability over the data network.

Connecting via External Layer 3 Network

We recommend connecting the Nexus Dashboard cluster to the fabrics via an external Layer 3 network as it does not tie the cluster to any one fabric and the same communication paths can be established to all sites. Specific connectivity depends on the type of applications deployed in the Nexus Dashboard:

- If you are deploying Nexus Dashboard Orchestrator to manage Cisco ACI fabrics only, you can establish connectivity from the data interface to either the in-band or out-of-band (OOB) interface of each site's APIC or both.
- If you are deploying Nexus Dashboard Orchestrator to manage Cisco NDFC fabrics, you must establish connectivity from the data interface to the in-band interface of each site's NDFC.
- If you are deploying Day-2 Operations applications, such as Nexus Dashboard Insights, you must establish connectivity from the data interface to the in-band network of each fabric and of the APIC.

If you plan to connect the cluster across an external Layer 3 network, keep the following in mind:

- For ACI fabrics, you must configure an L3Out and the external EPG for Cisco Nexus Dashboard data network connectivity in the management tenant.

Configuring external connectivity in an ACI fabric is described in [Cisco APIC Layer 3 Networking Configuration Guide](#).

- For NDFC fabrics, if the data interface and NDFC's in-band interface are in different subnets, you must add a route to the Nexus Dashboard's data network on NDFC.

You can add the route from the NDFC UI by navigating to **Administration > Customization > Network Preference > In-Band (eth2)**, then adding the route and saving.

- If you specify a VLAN ID for your data interface during setup of the cluster, the host port must be configured as trunk allowing that VLAN.

However, in most common deployments, you can leave the VLAN ID empty and configure the host port in access mode.

The following two figures show two distinct network connectivity scenarios when connecting the Nexus Dashboard cluster to the fabrics via an external Layer 3 network. The primary purpose of each depends on the type of application you may be running in your Nexus Dashboard.

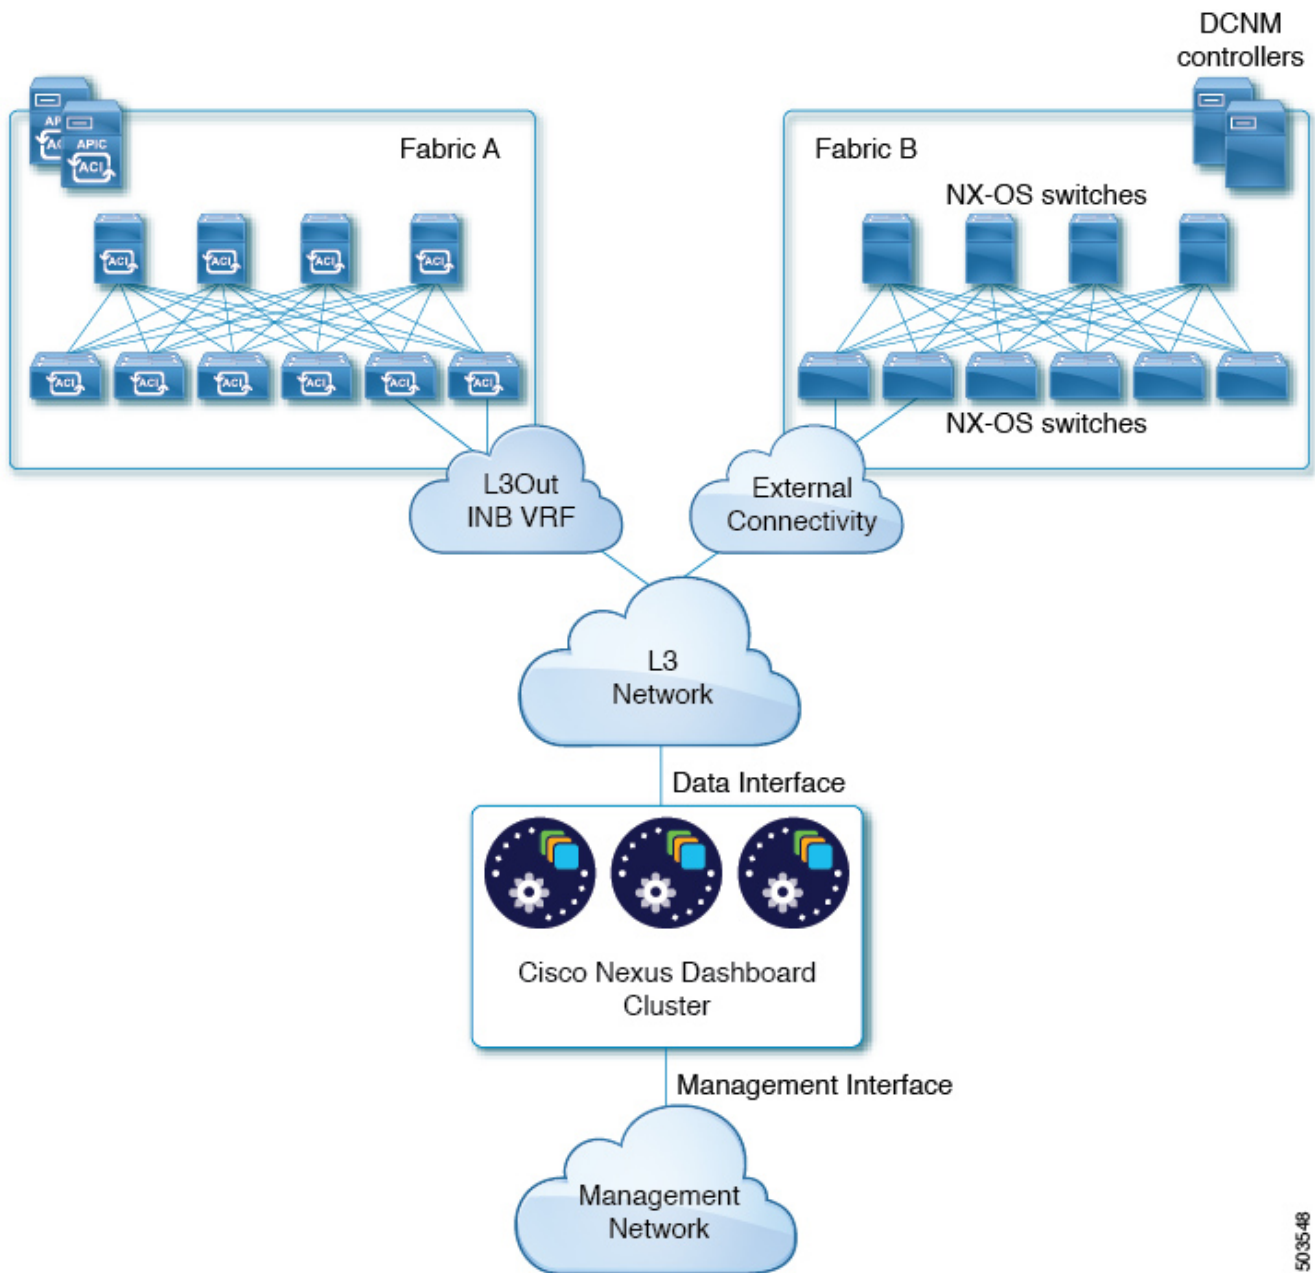
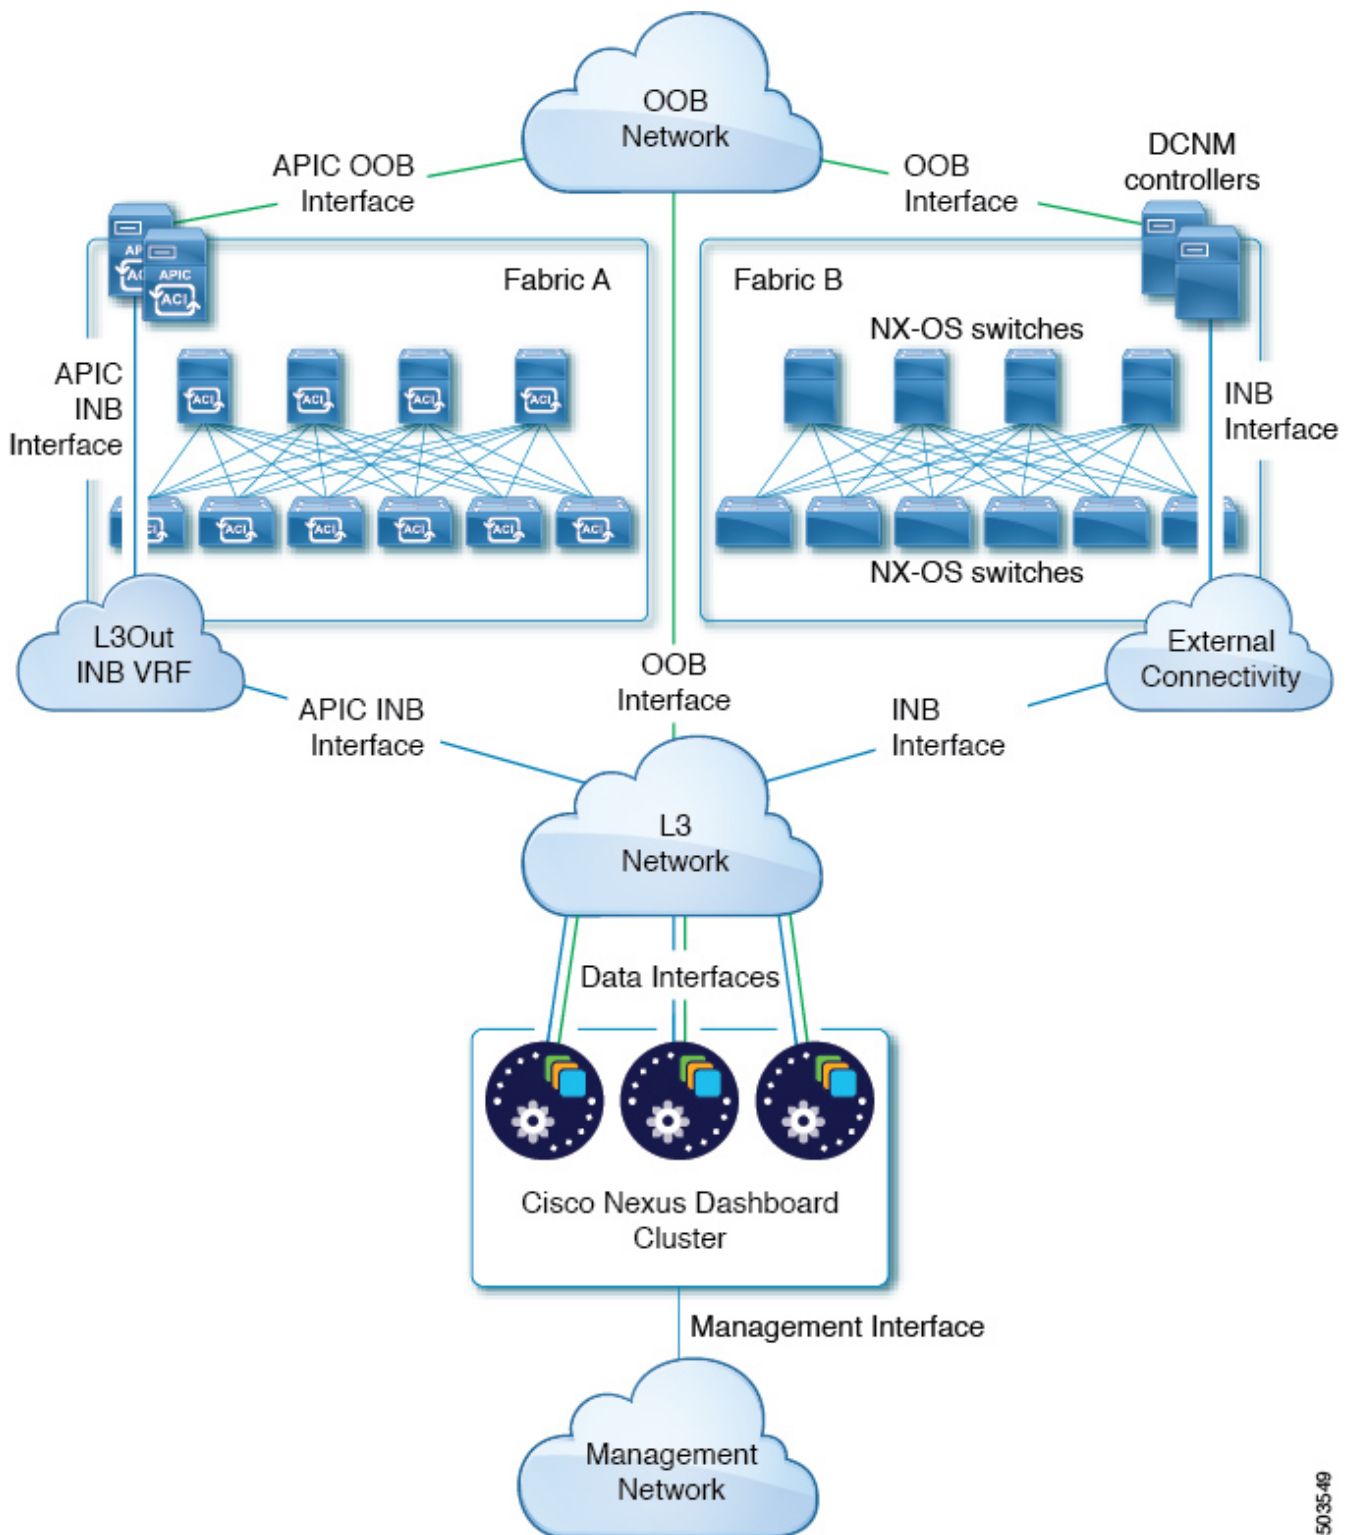


Figure 2. Connecting via External Layer 3 Network, Day-2 Operations Services

503548



503549

Figure 3. Connecting via External Layer 3 Network, Nexus Dashboard Orchestrator

Connecting the Nodes Directly to Leaf Switches

You can also connect the Nexus Dashboard cluster directly to one of the fabrics. This provides easy connectivity between the cluster and in-band management of the fabric, but ties the cluster to the specific fabric and requires reachability to other fabrics to be established through external connectivity. This also makes the cluster dependent on the specific fabric so issues within the fabric may impact Nexus Dashboard connectivity. Like in the previous example, connectivity depends on the type of applications deployed in the Nexus Dashboard:

- If you are deploying Nexus Dashboard Orchestrator to manage Cisco ACI fabrics only, you can

establish connectivity from the data interface to either the in-band or out-of-band (OOB) interface of each site's APIC.

- If you are deploying Nexus Dashboard Insights or Network Assurance Engine, you must establish connectivity from the data interface to the in-band interface of each fabric.

For ACI fabrics, the data interface IP subnet connects to an EPG/BD in the fabric and must have a contract established to the local in-band EPG in the management tenant. We recommend deploying the Nexus Dashboard in the management tenant and in-band VRF. Connectivity to other fabrics is established via an L3Out.

- If you are deploying Nexus Dashboard Insights with ACI fabrics, the data interface IP address and the ACI fabric's in-band IP address must be in different subnets.

If you plan to connect the cluster directly to the leaf switches, keep the following in mind:

- If deploying in VMware ESX or Linux KVM, the host must be connected to the fabric via trunk port.
- If you specify a VLAN ID for your data network during setup of the cluster, the Nexus Dashboard interface and the port on the connected network device must be configured as trunk


However, in most cases we recommend not assigning a VLAN to the data network, in which case you must configure the ports in access mode.

- For ACI fabrics:
 - We recommend configuring the bridge domain (BD), subnet, and endpoint group (EPG) for Cisco Nexus Dashboard connectivity in management tenant.

Because the Nexus Dashboard requires connectivity to the in-band EPG in the in-band VRF, creating the EPG in the management tenant means no route leaking is required.

- You must create a contract between the fabric's in-band management EPG and Cisco Nexus Dashboard EPG.
- If several fabrics are monitored with apps on the Services Engine cluster, L3Out with default route or specific route to other ACI fabric in-band EPG must be provisioned and a contract must be established between the cluster EPG and the L3Out's external EPG.

The following two figures show two distinct network connectivity scenarios when connecting the Nexus Dashboard cluster directly to the fabrics' leaf switches. The primary purpose of each depends on the type of application you may be running in your Nexus Dashboard.



503551

Figure 4. Connecting via an EPG/BD, Day-2 Operations Services




Figure 5. Connecting via an EPG/BD, Nexus Dashboard Orchestrator

503552

GUI Overview

After you have deployed the Nexus Dashboard cluster, you can perform all remaining actions using its GUI. To access Cisco Nexus Dashboard GUI, simply browse to any one of the nodes' management IP addresses:

```
https://<node-mgmt-ip>
```



Depending on the permissions of the user logged in to the Nexus Dashboard GUI, the UI will display only the objects and settings the user is allowed to access. The following sections describe all GUI elements as visible by an **admin** user. For more information on user configuration and permissions, see [Users](#).

Navigation Bar and User Preferences

As you navigate through the Nexus Dashboard UI and any installed services, the top of the screen will always display the common navigation bar:

- The **Home** button will return you to the **One View** page (described in the following section) from any page or service you are currently viewing.
- The **Feedback** button allows you to send feedback and suggestions or report issues as you are using the software.
- The **Help** menu provides access to the version information, new features in the current release, and the documentation for Nexus Dashboard as well as any services you have installed.
- The **user** menu allows you to log out, change the password for the currently logged in user, and configure one or more user-specific preferences:
 - **Show Welcome Screen On Login** toggles whether the new features screen is shown every time the current user logs in.
 - **Time Zone Preference** allows you to specify the time zone for the currently logged in user allowing multiple users across different geographical locations to more conveniently view any time-specific information in the UI.

When set to **Automatic**, your local browser time zone is used. This is the default and the same behavior as in previous releases of Nexus Dashboard.

When set to **Manual**, you can pick the geographic location from the map and the closest time zone will be set according to that.

The time zone conversion is done in the UI only, the backend and the APIs continue to return timestamps in the format in which they are saved, which is typically UTC.

This release supports the global time zone configuration for Nexus Dashboard and Insights service only, other services may continue to use automatic or internally configured time zone settings. The time zone setting for the Nexus Dashboard Insights service is absolute. In other words, if you have multiple sites across different geographical regions, all source time zones are mapped to the configured time zone.

One View Page

The first page you will see when you log into your Nexus Dashboard cluster is the **One View**. This page provides information about the current Nexus Dashboard cluster's status, sites, services, and resources usage:




Figure 6. One View

Here you can have a single place for status overview of the entire cluster (or all clusters in case of multi-cluster connectivity) for Dashboard Users. Note that the **Nexus Snapshot** information is available for Nexus Dashboard Insights service only.

You can always access the One View page by clicking the **Home** icon in the top left corner of the UI.

Admin Console Page

You can navigate to your Nexus Dashboard cluster's **Admin Console** by clicking **Admin Console** in the One View page after you log in. The **Overview** page of the admin console provides information about the current Nexus Dashboard cluster's status, sites, services, and resources usage.




Figure 7. System Overview

1. The global **Home** icon provides a convenient way of returning to the Nexus Dashboard’s home screen that allows you to navigate between different components:

- o **One View** page (described above), which provides a Single Pane of Glass (SPOG) view into all your connected cluster, sites, and services.

In multi-cluster deployments, **One View** will show all available resources and services across all of your clusters. For more information, see [Multi-Cluster Connectivity](#).

- o **Admin Console** (pictured above), where you can configure and administer your Nexus Dashboard cluster.
- o **Services** for one-click access to every service available in your cluster.

In multi-cluster deployments, **Services** will include all services across all of your clusters.

- o **Sites** for one-click access to the controller UI of any site onboarded to your cluster.

In multi-cluster deployments, **Sites** will include all sites across all of your clusters.

2. **Current Cluster** displays the name of the currently viewed cluster.

In multi-cluster deployments, you can click the name of the cluster to quickly switch to another connected cluster.

3. **Intent** icon allows you to access the most common tasks, such as adding sites or nodes, upgrading the cluster, creating users, and so on.

4. **Overview** tile displays System Status, Cluster Health, and Cisco Intersight Status.

You can click on the **Cluster Health** status to view specific details of any issues in your cluster.

5. **Sites and Services** tile displays the **Sites** by connectivity, as well as **Services** and **Infra Services** by status.

Connectivity indicates whether the sites are up (**Up**) or down (**Down**).

Status is displayed in number of services that are **healthy**, have **minor** faults, or have **critical** faults.

6. **Inventory** tile provides details of the **nodes**, **Pods**, **Deployments**, and other statistics about the currently selected cluster.



You can click different areas in the **System Overview** tab to open the corresponding GUI screens where you can see additional details or make configuration changes.

Sites Page

The **Sites** page in the left navigation pane allows you to onboard sites from a single location and then use those sites from any service deployed in your cluster.

Any site that is already onboarded is listed on this page, including the following:

- **Health Score**—current health status of the site as reported by the site’s controller.
- **Name** —the name of the site as you provided it during onboarding.
- **Connectivity Status**—indicates whether the site’s connectivity is established (**Up**) or not (**Down**).
- **Firmware Version**—the version of the controller software currently running in the site.
- **Services Used**—list of services currently using the specific site.

For additional information on onboarding sites, see [Site Management](#).

Services Page

The **Services** page in the left navigation pane allows you to access and manage services in your Nexus Dashboard.

Any service that is already installed and enabled is listed under the **Installed Services** tab, while the **App Store** tab provides an easy way to deploy additional services directly from the Cisco’s Data Center App Center page.

For additional information on managing services, see [Services Management](#).

System Resources Pages

The **System Resources** category in the left navigation pane displays the cluster resources, such as the nodes that make up your cluster and the Kubernetes API objects utilized by the cluster.

The category contains the following subcategories:

- **Nodes**—provides information about all **master**, **worker**, and **standby** nodes in your cluster along with their networking configuration and CPU/Memory utilization.
- **Pods**—provides information about pods, which are a fundamental unit of compute.

A pod is a group of containers that are scheduled together and are generally static. When a service requires a change in how it is deployed, new pods are created with the new configuration and the old pods are destroyed instead of changing the configuration of the existing pods in place.

- **Namespaces**—provides information about Kubernetes namespaces used to organize groups of other API objects.

Namespaces can be used to operate on all objects in the namespace at once or to restrict access to particular users or roles.

- **Services**—provides information about services (or set of dynamically-changing pods and containers) running in the cluster.

Each service consists of multiple pods and containers, which may be created, destroyed, or changed during cluster scaling or recovery, but the service's names provide a static way of accessing the specific service irrespective of its underlying configuration.

- Deployments, StatefulSets and DaemonSets provide the service developers with ways to describe how and where to deploy sets of Pods.
 - **Deployments**—Deployments are the most general of these objects and simply define a set of pods with the ability to set constraints about how many copies of the pod are deployed and on what type of node.
 - **DaemonSets**—define a Pod that runs on every host in the Kubernetes cluster and is automatically created whenever a Node is added to the cluster.
 - **StatefulSets**—define Pods that need to be run on a predictable host with a specific storage volume. If these pods go down, they are recreated in the same place with the same persistent identifier, so that they can use the same storage volume as their previous incarnation.

Operations Pages

The **Operations** category in the left navigation pane displays the actions that can be performed on Nexus Dashboard.

The category contains the following subcategories:

- **Firmware Management**—Firmware Management is used to perform cluster (firmware) upgrade or downgrade.
- **Tech Support**—An administrator can perform technical support collections.
- **Audit Logs**—Audit Logs are user triggered configuration changes.
- **Backup and Restore**—Backup and Restore displays the backed up and restored configuration.

Infrastructure Pages

The **Infrastructure** category in the left navigation pane allows you to management the Nexus Dashboard cluster, Cisco Intersight connector, and application Infra services.

- **Cluster Configuration**—provides cluster details (such as name, app subnet, and service subnet), allows you to configure cluster-wide settings (such as DNS and NTP servers, persistent IP addresses, and routes), and displays any current issues in the cluster.
- **Resource Utilization**—provides real-time information about the resource utilization of your Nexus Dashboard cluster.
- **Intersight**—provides access to Cisco Intersight device connector configuration.

The Cisco NI service depends on the Intersight Device Connector for service to be configured and available on the service node.

- **App Infra Services**—provides information about the infra services running on your Nexus Dashboard and allows you to restart individual microservices if needed.

Administrative Pages

The **Administrative** category in the left navigation pane allows you to manage authentication and users.

- **Authentication**—allows you to configure remote authentication domains as described in [Remote Authentication](#).
- **Security**—allows you to view and edit the security configurations, such keys and certificates.
- **Users**—allows you to create and update local Nexus Dashboard users as described in [Users](#) or view the users configured on any remote authentication servers you added to the Nexus Dashboard.

Site Management

With Cisco Nexus Dashboard, you can on-board multiple Cisco ACI, Cisco Cloud Network Controller, and Cisco NDFC fabrics as individual sites to the same cluster. Once the fabrics are on-boarded, they can be used by the applications running on the same Cisco Nexus Dashboard cluster.

To add a site, you need its controller's in-band or out-of-band IP address and credentials. The type of the IP address you will use for site onboarding depends on the Nexus Dashboard services that will use the site and is described in detail in the following sections. Sites added to the Cisco Nexus Dashboard cluster are not enabled in the services by default, so you will need to explicitly enable them directly from each service's own GUI.

After you on-board one or more sites to your Nexus Dashboard, you can view them in the Nexus Dashboard GUI by selecting **Sites** from the left navigation sidebar. You can also use the **Sites** page to launch directly into any of the site's GUIs by clicking the **Open** link next to the site's name.

If you are using remote authentication to login to your Nexus Dashboard and you have the same login domain and user configured in the site you are launching, you will be able to login to the site's GUI automatically without having to re-authenticate yourself.

Adding Sites

Before you begin

- Fabric connectivity must be already configured.
- If adding a Cisco APIC or Cloud Network Controller site, the site must be running Release 4.2(4) or later.
- If adding a Cisco APIC site, EPG/L3Out for Cisco Nexus Dashboard data network IP connectivity must be pre-configured.

Refer to [Fabric Connectivity](#) for more information.

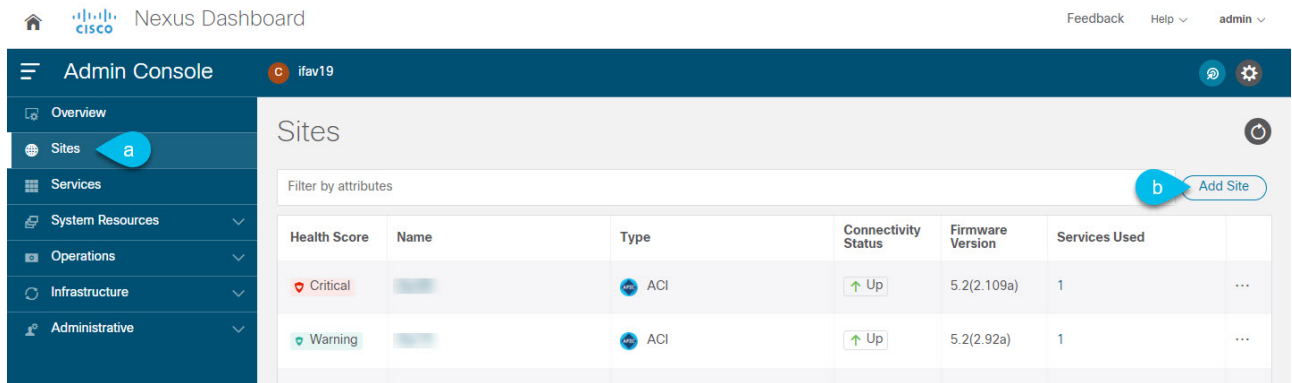
- If adding a Cisco APIC site and planning to deploy Cisco NIR application:
 - IP connectivity from Cisco Nexus Dashboard to Cisco APIC Inband IP over data network must be configured.
 - IP connectivity from Cisco Nexus Dashboard to the leaf nodes and spine nodes in-band IPs must be configured.
- If adding a Cisco NDFC site:
 - The site must be running Release 11.5(1) or later.
 - You must configure Layer 3 connectivity to the fabric and switches.
 - If your cluster is deployed in AWS or Azure, you must configure inbound rules on the data interface.

This is typically done during initial cluster deployment and described in detail in the [Cisco Nexus Dashboard Deployment Guide](#).

To add a site:

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

2. From the main navigation menu, select **Sites**.
3. Add a site.



- a. From the main navigation menu, select **Sites**.
- b. In top right of the main pane, click **Add Site**.

The **Add Site** screen opens.

4. Select the type of site you want to add.



While Cisco Nexus Dashboard supports on-boarding all three types of fabrics, for specific fabric types and versions compatible with your services, see the [Services Compatibility Matrix](#).

- **ACI**—for on-premises ACI sites managed by Cisco APIC
- **Cloud Network Controller**—for cloud sites managed by Cisco Cloud Network Controller
- **NDFC**—for on-premises sites managed by Cisco NDFC

5. Provide the site's information.

- a. If adding an **ACI** site, provide the following:

- **Site Name**—used throughout the Nexus Dashboard GUI when referring to this site.
- **Host Name/IP Address**—used to communicate with the Cisco APIC.

If you will use the site with Nexus Dashboard Orchestrator service only, you can provide either the in-band or out-of-band IP address of the APIC. If you will use the site with Nexus Dashboard Insights as well, you must provide the in-band IP address.



When providing the address, do not include the protocol (**http://** or **https://**) as part of the URL string or site addition will fail.

- **User Name** and **Password**—login credentials for a user with **admin** privileges on the site you are adding.
- (Optional) **Login Domain**—if you leave this field empty, the site's local login is used.
- (Optional) **Validate Peer Certificate**—allows Nexus Dashboard to verify that the certificates of hosts to which it connects (such as site controllers) are valid and are signed by a trusted Certificate Authority (CA).



You must have the certificate for this site already imported into your Nexus Dashboard before you can add a site using this option. If you have not yet added the certificates, cancel the **Add Site** wizard and follow the instructions described in [Validating Peer Certificates](#) first; then after you have imported the certificates, add the site as described here.

If you enable the **Verify Peer Certificate** option but don't import the valid certificate, site onboarding will fail.

- (Optional) **In-Band EPG**—required when connecting to an ACI fabric via an EPG and bridge domain. For more information on fabric connectivity, see [Fabric Connectivity](#).

If you plan to use this site with the Nexus Dashboard Insights service, you must provide the node management In-Band EPG.

b. If adding a **Cloud Network Controller** site, provide the following:

- **Site Name**—used throughout the Nexus Dashboard GUI when referring to this site.
- **Host Name/IP Address**—used to communicate with the Cloud Network Controller.



When providing the address, do not include the protocol ([http://](#) or [https://](#)) as part of the URL string or site addition will fail.

- **User Name** and **Password**—login credentials for a user with **admin** privileges on the site you are adding.
- (Optional) **Login Domain**—if you leave this field empty, the site's local login is used.
- (Optional) **Validate Peer Certificate**—allows Nexus Dashboard to verify that the certificates of hosts to which it connects (such as site controllers) are valid and are signed by a trusted Certificate Authority (CA).



You must have the certificate for this site already imported into your Nexus Dashboard before you can add a site using this option. If you have not yet added the certificates, cancel the **Add Site** wizard and follow the instructions described in [Validating Peer Certificates](#) first; then after you have imported the certificates, add the site as described here.

If you enable the **Verify Peer Certificate** option but don't import the valid certificate, site onboarding will fail.

- (Optional) **Enable Proxy**—enable this setting if your cloud site is reachable via a proxy.



Proxy must be already configured in your Nexus Dashboard's cluster settings. If the proxy is reachable via management network, a static management network route must also be added for the proxy IP address. For more information about proxy and route configuration, see [Cluster Configuration](#).

c. If adding an **NDFC** site, provide the following:

- **Host Name/IP Address**—used to communicate with the Cisco NDFC.

This must be the in-band IP address of NDFC.



When providing the address, do not include the protocol ([http://](#) or [https://](#)) as part of the URL string or site addition will fail.

- **User Name** and **Password**—login credentials for a user with **admin** privileges on the site you are adding.
- (Optional) **Login Domain**—if you leave this field empty, the site’s local login is used.
- (Optional) **Validate Peer Certificate**—allows Nexus Dashboard to verify that the certificates of hosts to which it connects (such as site controllers) are valid and are signed by a trusted Certificate Authority (CA).



You must have the certificate for this site already imported into your Nexus Dashboard before you can add a site using this option. If you have not yet added the certificates, cancel the **Add Site** wizard and follow the instructions described in [Validating Peer Certificates](#) first; then after you have imported the certificates, add the site as described here.

If you enable the **Verify Peer Certificate** option but don’t import the valid certificate, site onboarding will fail.

- **Sites**—click **Select Sites** to select the NDFC fabrics managed by the controller you provided.

6. Click **Add** to finish adding the site.
7. (Optional) Click on the **Geographical Location** map to specify where the site is located.
8. (Optional) Repeat these steps for any additional sites.

Editing Sites

To edit a site:

1. Navigate to your Nexus Dashboard’s **Admin Console**.
2. From the main navigation menu, select **Admin Console**.
3. From the main navigation menu, select **Sites**.
4. From the **Actions (...)** menu for the site you want to edit, select **Edit Site**.

The **Edit Site** screen opens.

5. In the **Edit Site** screen, make the required changes.
 - To remove a security domain, click the **Delete** icon next to an existing domain.
 - To add one or more security domains, click **+Add Security Domain**.
 - To re-provision the site, check the **Re-register Site** checkbox and provide the required information.

Re-registering a site may be required for Cloud Network Controller sites used with Nexus Dashboard Orchestrator in case the Cloud Network Controller’s public IP address changes.

You can also use this option if you changed the IP address information for a NDFC fabric managed by the Orchestrator service.



Re-registering a site is not supported for the Nexus Dashboard Insights service.

6. Click **Save** to save the changes

Deleting Sites

Before you begin

- Ensure that the site is not used by any applications installed in your Nexus Dashboard.

Deleting a site will cause an interruption to all applications using this site.

- When a Cisco ACI fabric is added as a site to Nexus Dashboard, some policies may be created in the Cisco APIC. If the Nexus Dashboard is clean rebooted without deleting the on-boarded site, the policies created on Cisco APIC will not be deleted. To clean up these policies on Cisco APIC, the site should be re-added and deleted.

To remove one or more sites:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the main navigation menu, select **Admin Console**.
3. From the main navigation menu, select **Sites**.
4. From the **Actions (...)** menu for the site you want to remove, select **Remove Site**.
5. In the **Confirm Delete** window, provide the login information for the site
6. Click **OK** to remove the site.

Services Management

With Cisco Nexus Dashboard, you can manage all of your services including their entire lifecycle from the **Services** GUI page. This page also allows you to explore the Cisco DC App Center and discover all the services that are available for the Nexus Dashboard.

Installing Services Using App Store

The App Store screen allows you to deploy services directly from the Cisco DC App Center.

Before you begin

- You must have administrative privileges to install services.
- The Cisco DC App Center must be reachable from the Nexus Dashboard via the Management Network directly or using a proxy configuration.

Setting up a proxy is described in [Cluster Configuration](#).

- Keep in mind that only the latest versions of services are available for installation using the App Store.

If you want to install a version of a service prior to the latest available in the App Store, you can follow the manual installation procedures as described in [Installing Services Manually](#).

- Ensure that the cluster is healthy before installing a service.

To install a service from the App Store:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. Install a service from the **App Store**.




Figure 8. Install Service from App Store

- a. From the main navigation menu, select **Services**.
- b. In the main pane, select the **App Store** tab.
- c. In the tile of the service you want to install, click **Install**.

Nexus Dashboard will download the service directly from the DC App Center and install it. After the process completes, the service will become available in the **Services** page.

This may take up to 20 minutes depending on the service.

3. Start the service.

By default, after the service is installed, it remains in the **disabled** state. Follow the steps described in [Enabling Services](#) to enable it.

This may take up to 20 minutes depending on the service.

Installing Services Manually

Alternatively, you may choose to manually download the services from the DC App Center and then upload them to the Nexus Dashboard to install.

Before you begin

- You must have administrative privileges to install services.
- Ensure that the cluster is healthy before installing a service.

To install a service manually:

1. Download the service's image.
 - a. Browse to the [Cisco DC App Center](#).
 - b. In the **Search for apps...** field, enter the name of the service you want to download and press Enter.

For example, **network insights**.
 - c. On the search results page, click the service.
 - d. On the service page, click **Download**.
 - e. In the **License Agreement** window, click **Agree and download**.

This will download the service's image file to your system.

2. Log in to your Nexus Dashboard GUI.
3. From the main navigation menu, select **Admin Console**.
4. Upload the service image.
 - a. From the main navigation menu, select **Services**.
 - b. In the top right of the main pane, click the **Actions** menu and select **Upload App**.
 - c. Choose the image file you downloaded.

You can choose to upload the service from an **http** service or from your local machine.

To upload a local image, select **Local** and click **Choose File** to select the service image you downloaded to your local system.

To use a remote server, select **Remote** and provide the URL to the image file.



If you are providing an **http** URL to the image, your web server must be configured to not interpret **.nap** files and serve them as-is. Typically, this means including the extension in the following line in the web server's

`httpd.conf` configuration file: `AddType application/x-gzip .gz .tgz .nap`

d. Click **Upload** to upload the app.

This may take up to 20 minutes depending on the service.

5. Wait for the upload and initialization process to finish.
6. Start the service.

By default, after the service is installed, it remains in the **disabled** state. Follow the steps described in [Enabling Services](#) to enable it.

This may take up to 20 minutes depending on the service.

Enabling Services

By default, after a service is installed, it remains in the **disabled** state. This section describes how to enable it.

Before you begin

- You must have already installed the service as described in [Installing Services Using App Store](#) or [Installing Services Manually](#).
- You must have configured the Network Scale parameters appropriate for your use case, as described in [Cluster Configuration](#).
- Ensure that the cluster is healthy before enabling a service.

To enable a service:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the main navigation menu, select **Services**.
3. In the service's tile, click **Start**.

In releases prior to Release 2.2(1), when you installed and enabled a service in your Nexus Dashboard cluster, you had to choose a service deployment profile, which defined the cluster resources (in number of CPUs and the amount of memory and storage) required for that specific service.

Beginning with Release 2.2(1), resource profile selection has been reduced to a number of more intuitive parameters directly related to your deployment use case. These parameters, such as number of switches or flows, describe the fabric size and use case intent and allow the cluster to intelligently determine the resources needed for the service. The parameters are categorized as "Network Scale" and must be provided prior to service deployment in the [Cluster Configuration](#) screen.

In addition, if the service requires a particular **App Infra Services** profile, the cluster will automatically update and restart that Infra service to satisfy the requirement before starting your application.

If the cluster does not contain the resources required to run the service, the service may provide a reduced capacity profile, which you can choose if you want to run the service in a reduced

capacity mode.

However, if the service you are trying to start is incompatible with the Nexus Dashboard version or the cluster size is insufficient to run the service even in reduced capacity mode, the cluster will return an error and you will not be able to start that service. If the service cannot be enabled due to cluster capacity, you may need to deploy additional worker node before you can start that service.

Updating Services

The process for updating services is similar to first deploying it, as described in [Installing Services Using App Store](#) or [Installing Services Manually](#).

When you upload a new version of an existing service, you will be able to select one of the available versions from the (...) menu on the service's tile in the **Services** screen.

To update an existing service:

1. Deploy the new version as described in [Installing Services Using App Store](#) or [Installing Services Manually](#).
2. Navigate to the **Services** screen in the Nexus Dashboard GUI.
3. Click the (...) menu on the service's tile and select **Available Version**.

Alternatively, you can click on the version count in the service tile to open the same menu.

4. In the available versions window that opens, click **Activate** next to the new version.

Disabling Services

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the main navigation menu, select **Services**.

All services installed in your Nexus Dashboard are displayed here.

3. Click the (...) menu on the service's tile and select **Disable** to disable the service.

Restarting Services

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the main navigation menu, select **Services**.

All services installed in your Nexus Dashboard are displayed here.

3. Click the (...) menu on the service's tile and select **Restart** to restart the service.

Uninstalling Services

Before you begin

You must disable the service before you can delete it.

1. Log in to your Nexus Dashboard GUI.
2. From the main navigation menu, select **Services**.

All services installed in your Nexus Dashboard are displayed here.

3. Click the (...) menu on the service's tile and select **Delete** to remove the service.

Operations

Firmware Management (Cluster Upgrades)

This section describes how to manage different firmware versions and perform cluster upgrades.

The upgrade process involves uploading a new image and then deploying it. As such, the same workflow can be used for cluster firmware downgrades as well.



The following sections provide reference information for firmware upgrades. For the latest upgrade process information, see the online version of the [Nexus Dashboard Deployment Guide](#).

Prerequisites and Guidelines

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.

The upgrade process is the same for all Nexus Dashboard form factors. However, if your existing cluster is deployed using physical servers, VMware ESX, Linux KVM, Azure, or AWS, you will use the target release's ISO image (nd-dk9.<version>.iso) to upgrade; if your existing cluster is deployed in Red Hat Enterprise Linux, you will use the RHEL-specific image (nd-rhel-<version>.tar).

- Ensure that you have read the Release Notes and upgrade guides 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 can find the service-specific documents at the following links:

- [Nexus Dashboard Fabric Controller Release Notes](#)
- [Nexus Dashboard Fabric Controller Upgrade Guide](#)
- [Nexus Dashboard Insights Release Notes](#)
- [Nexus Dashboard Insights Upgrade Guide](#)
- [Nexus Dashboard Orchestrator Release Notes](#)
- [Nexus Dashboard Orchestrator Upgrade Guide](#)
- If you are upgrading a physical Nexus Dashboard cluster, ensure that the nodes have a supported CIMC version for the target Nexus Dashboard release.

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

CIMC upgrade is described in detail in [Upgrading CIMC](#).

- You must perform configuration backups of your Nexus Dashboard and services before the upgrade to safeguard data and minimize any potential risk before proceeding with the upgrade.
- You must disable all services running in the cluster before upgrading to a 2.3.x release.

- 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 **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**.

- Ensure that no configuration changes are made to the cluster, such as adding worker or standby nodes, while the upgrade is in progress.
- After upgrading to this release, we recommend upgrading all the services to their latest versions. For a complete list of Nexus Dashboard and services interoperability support, see the [Nexus Dashboard and Services Compatibility Matrix](#).
- Nexus Dashboard does not support platform downgrades.

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

Adding Images

Before you can upgrade your Nexus Dashboard cluster, you need to make the upgrade image available by adding it using the GUI.

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`).

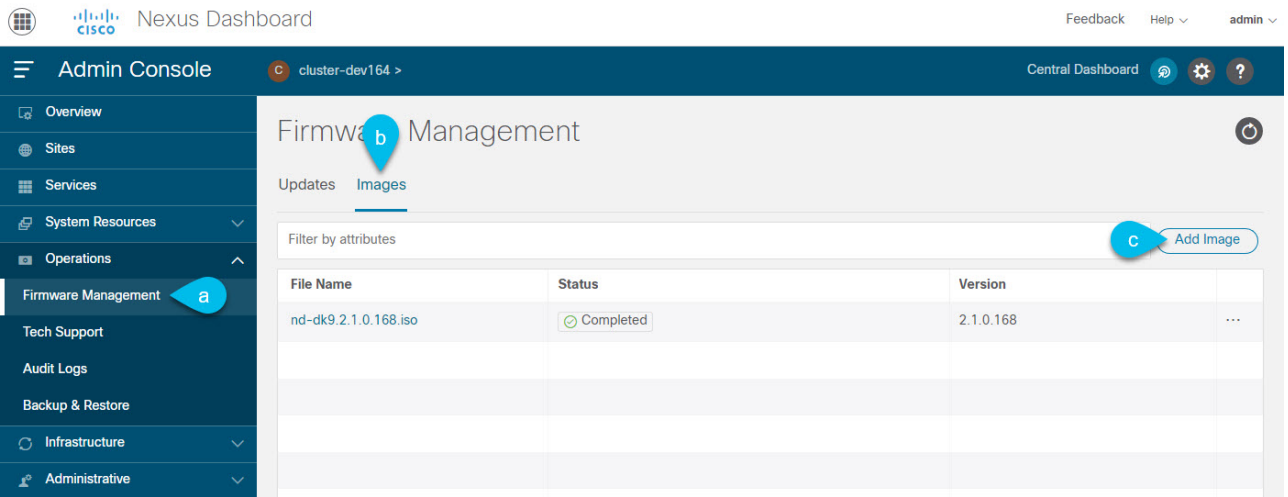


You must download the `.iso` image for all upgrades, even if you used the VMware ESX `.ova`, Linux KVM `.qcow2`, or a cloud provider's marketplace for initial cluster deployment.

- 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.

2. Navigate to your Nexus Dashboard's **Admin Console**.
3. Add an image.



- a. From the main navigation menu, select **Operations > Firmware Management**.
- b. In the main pane, select the **Images** tab.

The page will list any previously added images.

- c. In top right of the main pane, click the **Actions** menu and select **Add Image**.
4. In the **Add Firmware Image** window that opens, choose whether your image is stored on a remote server or local system.
 - a. If specifying a remote image, provide the full **URL** to the image.
 - b. If uploading a local image, click **Choose File** and select the image file from your local system.



If uploading from a local machine, slow upload speeds may cause the session to timeout which can interrupt the transfer. We recommend at least 40Mbps upload speed and increasing the session timeout to 1800 seconds (from the default 1200). You can change session timeout in the **Administrative > Security** page in your Nexus Dashboard GUI.

5. Click **Upload** to upload the image.

The **Images** tab will show the image upload progress, wait for it to finish before proceeding to the next section.

Upgrading the Cluster

Before you Begin

You must have the upgrade image already added to the Nexus Dashboard cluster as described in [Adding Images](#).

To upgrade your cluster:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. Initiate an update.
 - a. From the main navigation menu, select **Operations > Firmware Management**.
 - b. In the main pane, select the **Updates** tab.

c. Click **Set up Update** or **Modify Details**.

If this is the first time you are upgrading your cluster, simply click the **Setup Update** button in the middle of the page.


If you have previously upgraded the cluster, the last upgrade's details will be displayed in this page instead of the **Setup Update** button. In this case, click the **Modify Details** button at the top right of the screen.

3. In the **Setup/Version Selection** screen, select the target version and click **Next** to proceed.

If you uploaded multiple images to your Nexus Dashboard, they will be listed here.

4. Review the validation report and click **Install** to proceed with the upgrade.

Before the upgrade is triggered, the system will perform a number of validation checks and show the report:




5. In the **Setup/Confirmation** screen, review the details of the update and click **Begin Install** to proceed.

The screen will proceed to the **Install** tab and you will be able to see the progress of each node.

The process can take up to 20 minutes and you can navigate away from this screen in the meantime.

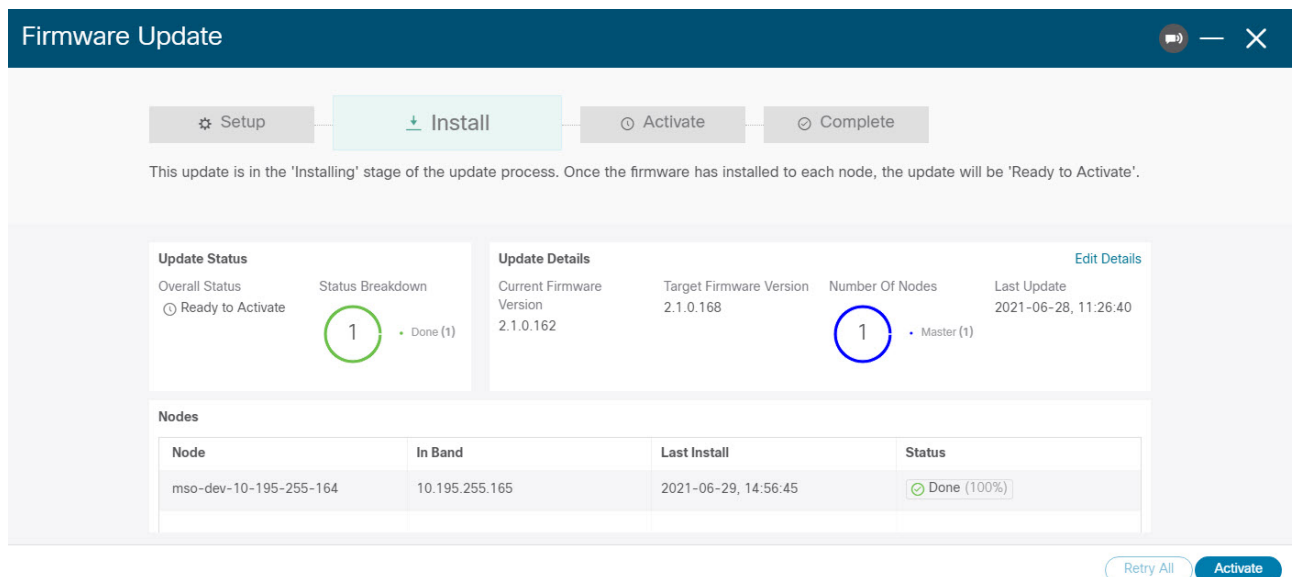
6. Wait for the image installation to complete.

You can check the installation status by navigating back to **Operations > Firmware Management** screen and clicking **View Details** link in the **Last Status** tile.



7. Click **Activate**.

If you navigated away from the installation screen, navigate back to **Operations > Firmware Management** screen and click **View Details** link in the **Last Status** tile.



It may take up to 20 additional minute for all the cluster services to start and the GUI may become unavailable during this process. The page will automatically reload when the process is completed. You can track the activation process in the **Activate** screen as shown below.

Deleting Images

Nexus Dashboard will retain any firmware images that you upload to it. If at any time you want to remove any of the images (for example, from older upgrades), you can use the following steps:

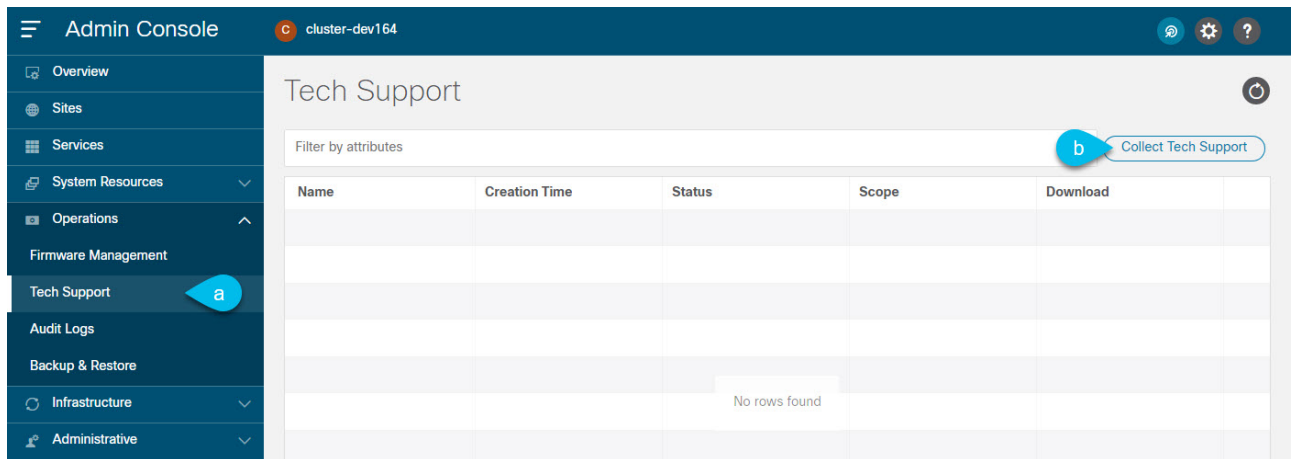
1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the main navigation menu, select **Operations > Firmware Management**.
3. In the main pane, select the **Images** tab.
4. Click the **Actions (...)** menu next to the image you want to delete and select **Delete Image**.
5. In top right of the main pane, click the **Actions** menu and select **Delete Image**.
6. In the **Confirm Delete** prompt, click **OK** to confirm.

Tech Support

Tech support enables you to collect logs and activities in the system for further troubleshooting by Cisco TAC. Cisco Nexus Dashboard provides best-effort tech support collection and gives ability to download tech support for individual nodes, the whole cluster, or applications. Tech support files are hosted on the Cisco Nexus Dashboard and can be downloaded at any time.

To collect Tech Support information:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. Collect Tech Support.



- a. From the main navigation menu, select **Operations > Tech Support**.
 - b. In top right of the main pane, click the **Actions** menu and select **Collect Tech Support**.
3. In the **Collect Tech Support** window that opens, provide a description.
 4. From the **Scope** dropdown, select the category for which you want to collect tech support information.
 - o **System** collects Infra tech support information.
 - o **App Store** collects App Store tech support information.
 - o Service-specific selections collects tech support information for that specific service.
 5. Click **Collect**.

After you begin Tech Support collection, you can see the progress in the same screen.

If for any reason the tech support collection process fails, you can also obtain the same information by logging into each node as the **rescue-user** and running one of the **acs techsupport collect** commands. For more information about specific **techsupport collect** command options, see [Useful Commands](#).

6. Download the Tech Support archive.

After the collection is finished, you can download the archive by clicking **Download** next it:


If you want to delete an existing Tech Support package, simply select it in the **Tech Support** screen and choose **Delete Tech Support** from the **Actions** menu.

Backup and Restore

This section describes how to back up or restore Nexus Dashboard cluster configuration.

Creating Configuration Backups

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. Start a back up.



- a. From the main navigation menu, select **Operations > Backup & Restore**.
 - b. In the main pane, select the **Backup Jobs** tab.
 - c. In top right of the main pane, click **Backup Configuration**.
3. In the **Backup Configuration** window that opens, provide the **Encryption Key** and the **File Name**.
The encryption key is used to encrypt the archive and must be at least 8 characters long.
 4. Click **Download** to start the backup.



Cisco Nexus Dashboard does not store configuration backups or encryption keys, so you must download and maintain them outside the Nexus Dashboard cluster.

Restoring Configuration

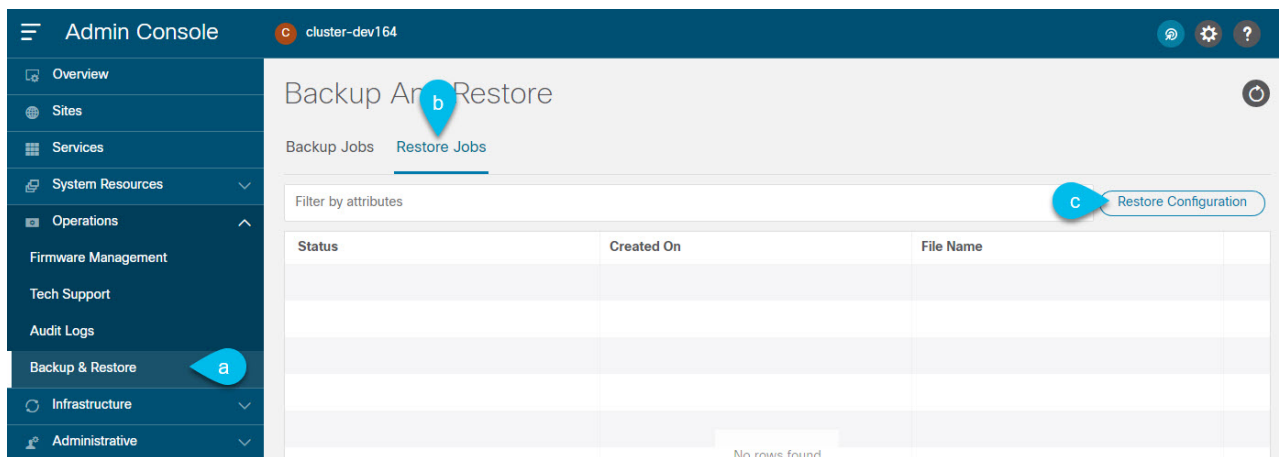
Before you begin

If your current configuration contains one or more of the following settings, you must remove them before restoring any backups:

- Persistent IPs, which are described in [Persistent IP Addresses](#).
- Syslog for streaming events, which is described in [Exporting Events](#).
- Static Routes, which are described in [Cluster Configuration](#).

To restore a configuration backup:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. Begin configuration restore.



- a. From the main navigation menu, select **Operations > Backup & Restore**.
- b. In the main pane, select the **Restore Jobs** tab.
- c. In top right of the main pane, click the **Restore Configuration**.

You do not need to select one of the listed backups. You will be asked to upload the configuration backup file in the next screen.

3. Provide the details.

- a. Provide the **Encryption Key**.

This must be the same encryption key that you used when creating the backup.

- b. Click **Choose File** and select the backup file.

Cisco Nexus Dashboard does not store configuration backups, so you must upload the backup file before restoring it

The file must be in **.tgz** or **tar.gz** format.

4. Click **Import** to start the restore process.

Event Analytics

The **Event Analytics** page in the **Operations** category allows you to see the system-wide list of events and alerts in your Nexus Dashboard cluster.

Events

The **Events** tab enables you to easily access your Nexus Dashboard's platform-level events and audit logs. The **Audit Logs** tab displays all events that occur during the cluster operation. In addition to viewing the events and logs directly in the Nexus Dashboard GUI, you can also configure the cluster to stream the events to an external syslog server, as described in [Cluster Configuration](#).

The **Events** tab includes high severity events that may require your attention to resolve:

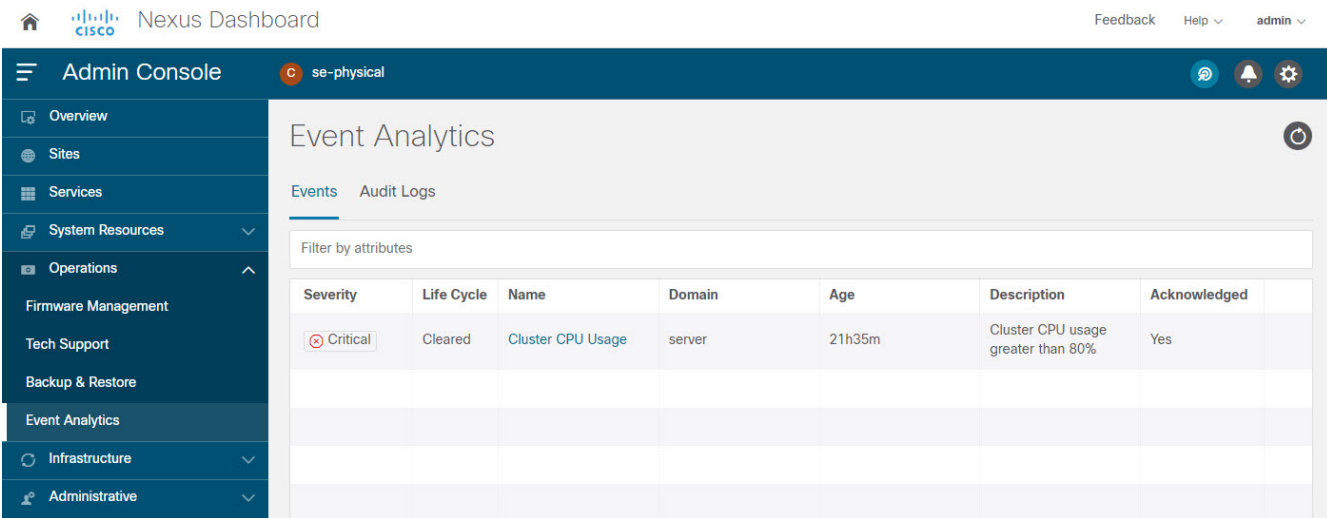


Figure 9. Event Analytics

You can view a summary of all critical events in the list or double-click any specific event for additional information about it. After you have viewed and analyzed an event, you can choose to acknowledge and clear it by clicking the **Actions (...)** menu next to the event in the list.

Audit Logs

Nexus Dashboard audit logging is automatically enabled when you first deploy the cluster and captures the operational changes made by the users in the environment.

You can view the audit logs directly in the GUI by selecting **Operations > Audit Logs** from the main navigation menu.

Note that the logs are not sorted by default; you can sort the list by clicking on any of the column headings.

You can choose to filter the list using the **Filter by attributes** field and providing a specific attribute and value pair.

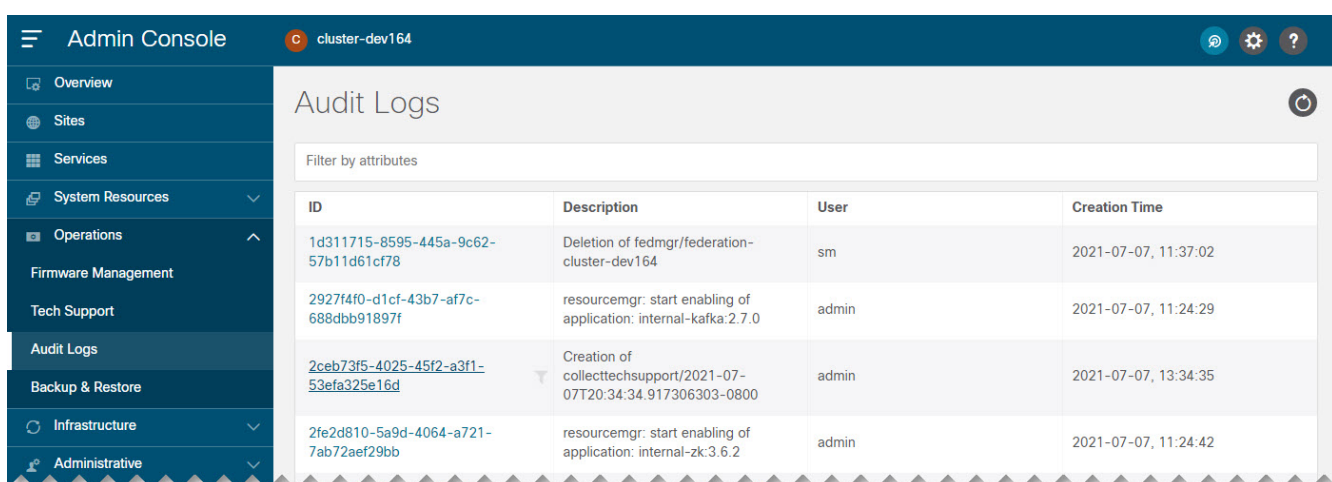


Figure 10. Audit Logs

Additionally, to view detailed information about a specific entry, simply click the entry in the list to open the **Details** tab.

Exporting Events

Nexus Dashboard can host multiple services one or more of which can generate various events, faults, and alerts. This information is published on and stored using Apache Kafka. While Release 2.1(2) allowed you to view and export cluster-level alerts to an external analyzer, there was no unified method to export all of this type of information to an external events monitoring service.

Beginning with Release 2.2(1), you can configure your cluster to export all platform-level, infrastructure-level, and service-level events to external monitoring and management systems. Each service running on Nexus Dashboard can define exactly which service-level events to aggregate and send to the cluster's Kafka service to export.

When configuring event streaming, the following restrictions apply:

- This release supports **syslog** event exporter only.
- Events are stored for up to 4 hours by default.

To configure event exporting:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the main navigation menu, select **Infrastructure > Cluster Configuration**.
3. Click the **Edit** icon in the **Syslog** tile.

In the **Syslog** dialog that opens, click **+Add Remote Destinations** to add a new server. Then provide the IP address, protocol, and port number for the server and choose whether you want to enable streaming to this syslog server at this time.

Infrastructure Management

Cluster Configuration

The cluster configuration GUI screen allows you to configure a number of options specific to the Nexus Dashboard cluster and its nodes. It will also display information about any issues that may be present in your Nexus Dashboard cluster.

Figure 11. Cluster Configuration



Only IPv4 addresses are supported for any of the following cluster configuration IP settings.

1. The **Multi-cluster Connectivity** tab allows you to connect multiple clusters together for a single pane of glass view and administration of the clusters and their sites, services, and configurations.

For more information, see [Multi-Cluster Connectivity](#).

2. The errors and warning tile will display the number of existing issues in your cluster. You can click **Expand** to see the full list of specific issues.
3. To configure a proxy for the Nexus Dashboard, click the **Edit** icon in the **Proxy Configuration** tile.

In certain deployment scenarios, such as with a combination of on-premises and cloud sites and the Nexus Dashboard cluster deployed inside a corporate network, you may have to access the Internet and the cloud sites through a proxy.



This release supports adding a single proxy server.

Note that Nexus Dashboard uses 2 main route tables—one for the Management network and one for the Data network—and by default, it will use the routing table of the originating IP address. In other words, Nexus Dashboard will attempt to reach the proxy from the routing table of the POD/Service that is trying to use the proxy.

For example, if you configure a proxy and establish Intersight connectivity from your Nexus Dashboard and then attempt to configure AppD integration from the Insights service running in the cluster, you may get an error stating that AppD host is not reachable. This happens because the proxy is only accessible from the management interface, so in such cases you also need to add a management network route for the proxy IP address, as described in "Management Network or Data Network routes" below.

To add a proxy server:

- a. Click **+Add Server** in the proxy configuration window.
- b. From the **Type** dropdown, select the type of traffic that you want to be proxied.
- c. In the **Server** field, provide the full address for the proxy server including the port if required.

For example <http://proxy.company.com:80>.

- d. If the server requires login credentials, provide the **Username** and **Password**.
- e. (Optional) Click **Add Ignore Host** to provide any hosts that will ignore the proxy.

You can add one or more hosts with which the cluster will communicate directly bypassing the proxy.

4. To add one or more Management Network or Data Network routes, click the **Edit** icon in the **Routes** tile.

Here you can define static routes for the management or data interfaces. For example, adding **10.195.216.0/21** as a Data Network route will cause all traffic destined to that subnet to transit out of the data network interface.

- o To add a management network route, click **Add Management Network Routes** and provide the destination subnet.
- o To add a data network route, click **Add Data Network Routes** and provide the destination subnet.

5. To add one or more External Service Pools, click the **Edit** icon in the **External Service Pools** tile.

This allows you to provide persistent IP addresses for services that require to retain the same IP addresses even in case they are relocated to a different Nexus Dashboard node.

For detailed information and configuration steps, see [Persistent IP Addresses](#).

6. To configure NTP settings, click the **Edit** icon in the **NTP** tile.

By default, the NTP server that you configured when deploying the Nexus Dashboard cluster is listed here.

You can provide additional NTP servers by clicking **+Add NTP Server**.

You can remove existing NTP server by clicking the **Delete** icon next to it. Keep in mind that at least one NTP server must be configured in your cluster.

7. To configure DNS settings, click the **Edit** icon in the **DNS** tile.

By default, the DNS server and search domain that you configured when deploying the Nexus Dashboard cluster are listed here.

You can provide additional DNS servers and search domains by clicking **+Add a Provider** or **+Add a Search Domain** respectively.

You can remove existing DNS server by clicking the **Delete** icon next to it.

8. To provide one or more **syslog** servers to stream event logs to, click the **Edit** icon in the **Syslog** tile.

In the **Syslog** dialog that opens, click **+Add Remote Destinations** to add a new server. Then provide the IP address, protocol, and port number for the server and choose whether you want to enable streaming to this syslog server at this time.

For more information, see [Event Analytics](#).

9. To configure **Network Scale**, click the **Edit** icon in the **Network Scale** tile.

In releases prior to Release 2.2(1), when you installed and enabled a service in your Nexus Dashboard cluster, you had to choose a service deployment profile, which defined the cluster resources (in number of CPUs and the amount of memory and storage) required for that specific service.

Beginning with Release 2.2(1), resource profile selection has been reduced to a number of more intuitive parameters directly related to your deployment use case. These parameters, such as number of switches or flows, describe the fabric size and use case intent and allow the cluster to intelligently determine the resources needed for the service. The parameters are categorized as "Network Scale".

- a. In the **Number of Sites** field, provide the target number of sites for your deployment that this Nexus Dashboard cluster will manage.
- b. In the **Number of Switches** field, provide the target number of switch nodes for your deployment.
- c. In the **Flows per second** field, provide the target number of flows for your Nexus Dashboard Insights service.

Persistent IP Addresses

You can provide persistent IP addresses for services that require to retain the same IP addresses even in case they are relocated to a different Nexus Dashboard node.

Nexus Dashboard Insights requires some services (such as SNMP trap, syslog, and others) to stream data from the switches in your fabrics to the service. An IP address is configured on the switches for this purpose. Typically, if the IP address changes when the service is relocated, the service will reconfigure the new IP address on the switches.

In order to avoid this IP reconfiguration impact on the fabric switches, the service can request that the services IP addresses are preserved, in which case you will need to define a set of IP addresses which can be assigned to the service for this purpose.

If a service requires persistent IP addresses, you will not be able to enable that service in the Nexus Dashboard until enough IP addresses are defined as described below.



This feature is supported for Nexus Dashboard Insights with NDFC fabrics only. In addition, if you are using Layer 2 functionality only (IPs configured as part of the management and data subnets) and your Nexus Dashboard is deployed in VMware ESX, you must enable promiscuous mode for both management and data network interface portgroups, as described in <https://kb.vmware.com/s/article/1004099>.

Prior to Release 2.2(1), this feature was supported only for clusters where all nodes were part of the same Layer 3 network and the persistent IPs were defined as part of the node's management or data networks. Here the application uses Layer 2 mechanisms like Gratuitous ARP or Neighbor Discovery to advertise the persistent IPs within its Layer 3 network.

Beginning with Release 2.2(1), the feature is supported even if you deploy the cluster nodes in different Layer 3 networks. In this case, the persistent IPs are advertised out of each node's data links via BGP, which we refer to as "Layer 3 mode". The IPs must not overlap with any of the nodes' management or data subnets. If the persistent IPs are outside the data and management networks, this feature will operate in Layer 3 mode by default; if the IPs are part of those networks, the feature will operate in Layer 2 mode.

Persistent IP Guidelines and Limitations

When configuring persistent IPs for your services:

- Ensure that you check the documentation for the services you plan to deploy as some services do not support this feature or require additional guidelines.

At this time, Persistent IPs are supported for Nexus Dashboard Insights and Nexus Dashboard Fabric Controller. You can find the service-specific documentation at the following links:

- [Nexus Dashboard Fabric Controller](#)
- [Nexus Dashboard Insights](#)

- You can choose which mode you want to operate in as long as the following conditions apply:
 - If you choose to operate in Layer 2 mode, the nodes must be part of the same data and management networks.
 - If you choose to operate in Layer 3 mode, all nodes must have BGP configuration provided

either during cluster deployment or after as described in [Enabling BGP On All Nodes](#).

- o You can switch between the two modes, in which case the existing services of a particular mode must be completely deleted and you will need to configure new persistent IPs corresponding to the new mode.
- If you configure one or more persistent IPs in Layer 3 mode and back up cluster configuration, the BGP settings required for this feature are not saved in the backup.

As such, you must ensure that you configure BGP for all cluster nodes before restoring any cluster configuration that contains Layer 3 persistent IPs in that cluster. If BGP is not configured prior to the configuration import, the import will fail.

Enabling BGP On All Nodes

If you want to operate in Layer 3 mode, you must enable and configure BGP for all nodes in your cluster. If you already configured BGP for each node during cluster deployment or if you want to operate in Layer 2 mode instead, you can skip this section and simply provide one or more persistent IPs from the nodes' management and data subnets, as described in [Configuring Persistent IPs](#). Note that if you choose to operate in Layer 2 mode, all nodes must be part of the same Layer 3 network. If you choose to operate in Layer 3 mode, at least one BGP peer must be configured on all cluster nodes to advertise the IPv4 or the IPv6 persistent IP addresses as described in this section.

Before you begin

- Ensure that the uplink peer routers are capable of exchanging the advertised persistent IPs across the Layer 3 networks of the cluster nodes.
- When a service requests a persistent IP address, the route advertised from the data links via BGP on the node where the service is running is maintained throughout the lifecycle of the service.

To configure BGP on the nodes:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the left navigation menu, select **System Resources > Nodes**.
3. Click the **Actions (...)** menu next to one of the nodes and choose **Edit**.
4. In the **Edit Node** screen, turn on **Enable BGP**.
5. In the **ASN** field, provide the autonomous system number for the node.
6. Click **+Add IPv4 BGP Peer** or **+Add IPv6 BGP Peer** to provide peer IP address information.
 - a. In the **Peer Address** field, provide the IPv4 or IPv6 address of the peer router for this node.

Multi-hop BGP peering is not supported, so you must ensure that the **Peer Address** is part of the node's data subnet.

- b. In the **Peer ASN** field, provide the autonomous system number of the peer router.

Only EBGP is supported, so you must ensure that the node ASN and Peer ASN are different.

- c. Click **Save** to save the changes.

7. Repeat these steps for every node in the cluster.

Every node in the cluster must have BGP configured.

You can configure the same ASN for all nodes or a different ASN per node

Configuring Persistent IPs

Before you begin

- For all persistent IPs, you must use either the Layer 2 or Layer 3 approach; a combination of the two is not supported.

If all nodes are in the same Layer 3 network, you can choose to use either the Layer 2 mode or Layer 3 mode for this feature. The two modes are described in [Persistent IP Addresses](#).

If the nodes are in different Layer 3 networks, you must configure the persistent IPs such that they don't overlap with either the management or the data subnets of the nodes.

- If the nodes in your cluster belong to different Layer 3 networks, you must have BGP enabled and configured as described in [Enabling BGP On All Nodes](#).
- There may be a momentary traffic interruption while a service using a persistent IP is relocated to a different node.

The interruption duration depends on the following factors:

- Time to detect the node failure
- Time for the service to get rescheduled to a different node
- Time for the service's external IP to get advertised from the scheduled node via GARP (IPv4) or neighbor discovery (IPv6) in case of Layer 2 mode
- Time for the service's external IP to get advertised from the scheduled node via BGP in case of layer 3 mode

To provide one or more persistent IP addresses:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the left navigation menu, select **Infrastructure > Cluster Configuration**.
3. In the **External Service Pools** tile, click the **Edit** icon.
4. In the **External Service Pools** screen that opens, click **+Add IP Address** to add one or more IP addresses for the management or data networks.

When editing persistent IPs, the following rules apply:

- If all nodes in your cluster are part of the same Layer 3 network, you can choose one of the following:
 - Layer 2 mode, in which case the IP addresses you add for management services must be part of the management subnet and the IP addresses for data services must be part of the data subnet.
 - Layer 3 mode, in which case the IP addresses you add must not overlap with the management or the data subnets of the nodes. In this case, adding IPs under "Management Service IPs" is not supported and you must add the IPs to the "Data Service IPs" category in the GUI.
- You must provide either IPv4 or IPv6 IP addresses, you cannot give both.

- You must add individual IP addresses one by one without any prefix; adding a range of IP addresses is not supported.
- You can remove any previously defined IPs, but you will not be able to remove any IPs that are currently in use by one or more services.

Multi-Cluster Connectivity

This release of Nexus Dashboard allows you to establish connectivity between multiple Nexus Dashboard clusters for a single pane of glass cluster administration, as well as access to any of the sites and services running on any of the connected clusters.

When you add a second cluster, a group of clusters is formed. The cluster from which you create the group becomes the "primary" cluster with a number of unique characteristics that do not apply to other clusters in the group:

- You must use the primary cluster to connect all additional clusters.
- You must use the primary cluster to remove any of the clusters from the group.

Establishing multi-cluster connectivity does not create any single databases with information from all clusters in the group. Every cluster continues to maintain its own configuration databases, while simultaneously being able to function as a proxy for all other clusters in the group regardless of which cluster an action or request is originated from or destined to.

Guidelines and Limitations

The following guidelines apply when configuring multi-cluster connectivity:

- This release supports multi-cluster connectivity between clusters deployed using physical or virtual (ESX) form factors only. In other words, you can join physical Nexus Dashboard clusters with virtual (ESX) clusters, but you cannot join virtual (KVM) or cloud clusters into the same group.
- For supported scale limits, such as number of clusters that can be connected together and number of sites across all clusters, see the [Nexus Dashboard Release Notes](#) for your release.
- Connectivity must be established between all nodes of all clusters, which will be connected via multi-cluster connectivity.
- The names of the sites onboarded in the clusters that you plan to connect together must be unique across those clusters.

Duplicate site names across different clusters may result in DNS resolution failures.

- The primary cluster, which you use to establish multi-cluster connectivity, must be running the same or later release of Nexus Dashboard as any other cluster in the group.

In other words, you cannot connect a Nexus Dashboard cluster running release 2.3.1 from a primary cluster that is running release 2.2.1.

- If you are upgrading multiple clusters that are connected together, you must upgrade the primary cluster first.
- From any cluster in the connected clusters group, you can view other clusters only if they are running the same or earlier version of Nexus Dashboard.

In other words, if **cluster1** is running release 2.3.1 and **cluster2** is running release 2.2.1, you can view **cluster2** from **cluster1** but not vice versa.

- Multi-Cluster connectivity and One View are supported for remote users only.

If you connect multiple clusters, but then login to one of the clusters as a local **admin** user, you will only be able to view and manage the local cluster into which you logged in.

To view and manage all clusters in the group, you must login as a remote user that is configured on all clusters.

- Nexus Dashboard Insights service in each cluster can view site groups from other Insights services across any cluster in the group.

However, when creating site groups, each Insights service can add sites which are onboarded in the same cluster where the service is installed only.

- Nexus Dashboard Orchestrator service supports managing only sites which are onboarded in the same cluster where the service is installed.

Connecting Multiple Clusters

Before you begin

- You must have familiarized yourself with the information provided in the [Guidelines and Limitations](#) section.
- You must have set up remote authentication and users on all clusters which you plan to connect.

Multi-Cluster connectivity and One View are supported for remote users only, so you must configure the same remote user with **admin** privileges for all clusters. For additional details, see [Remote Authentication](#).

To connect another cluster:

1. Log in to the Nexus Dashboard GUI of the cluster which you want to designate as the primary.
2. Add second cluster.

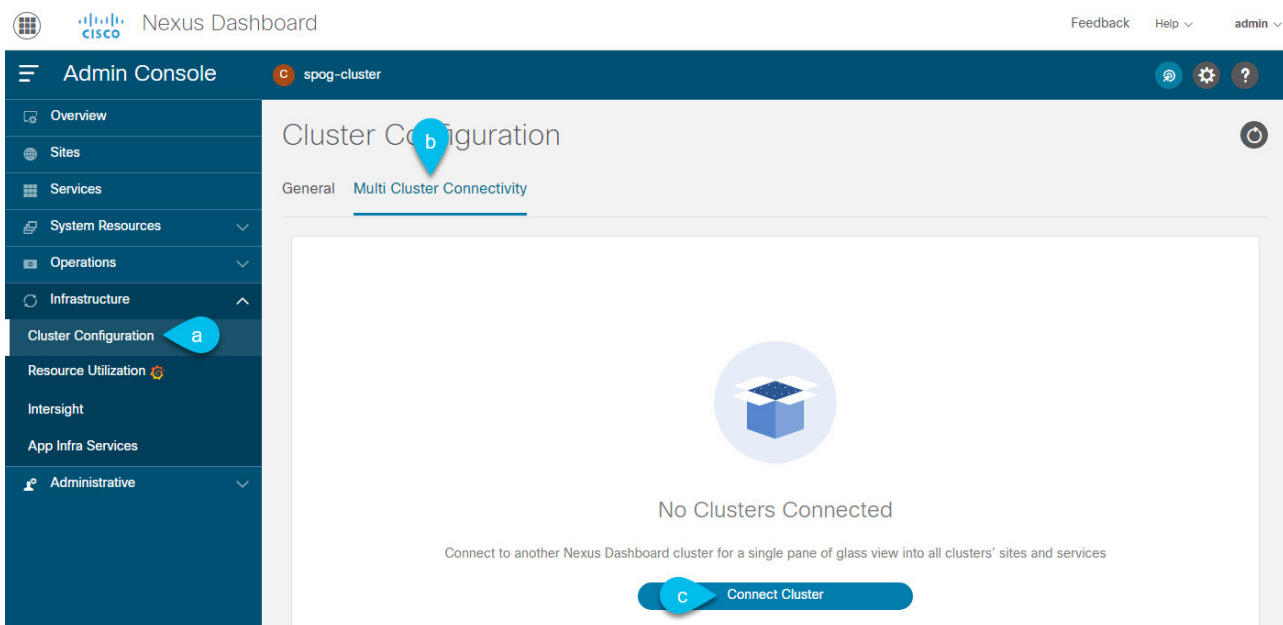


Figure 12. Adding Second Cluster

- a. From the main navigation menu, select **Infrastructure > Cluster Configuration**.
 - b. In the main pane, select the **Multi-Cluster Connectivity** tab.
 - c. Click **Connect Cluster**.
3. Provide cluster information.

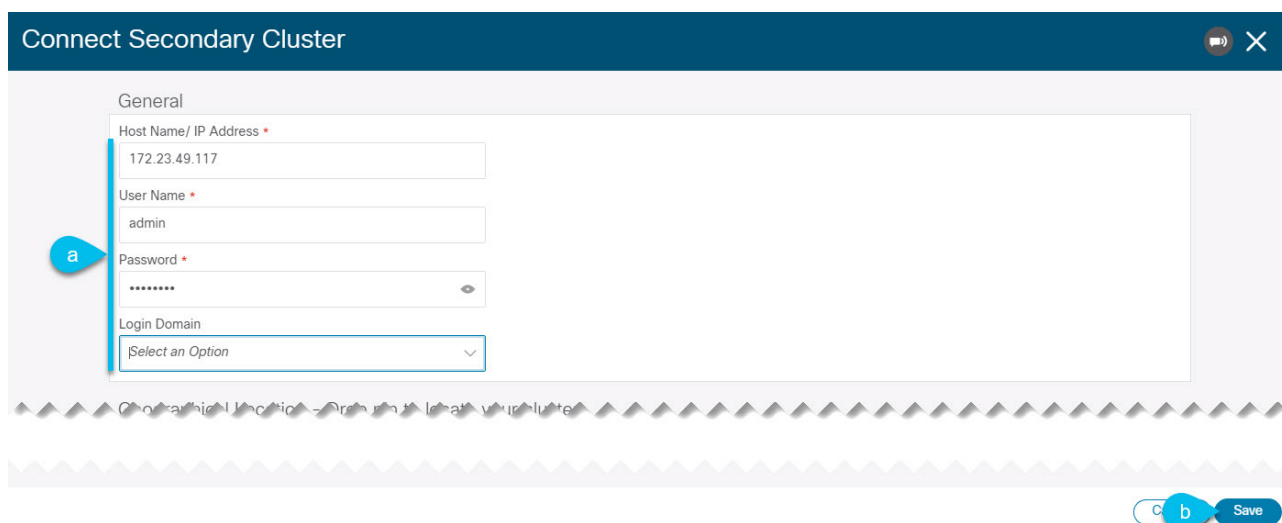


Figure 13. Providing Cluster Information

- a. In the information fields, provide the hostname or IP address and the authentication information for the cluster you are adding.

You only need to provide the management IP address of one of the nodes in the target cluster. Other nodes' information will be automatically synced after connectivity is established.

- b. Then click **Save**.

The user you provide must have administrative rights on the cluster you are adding. The user credentials are used once when you are first establishing connectivity to the additional cluster. After initial connectivity is established, all subsequent communication is done through secure keys. The secure keys are provisioned to each cluster while adding it to the group.

The cluster you are adding must not be part of an already existing group of clusters.

- Repeat the procedure for any additional Nexus Dashboard cluster which you want to add to the group.

After multiple clusters are added to the group, you can see their status in the **Cluster Configuration > Multi-Cluster Connectivity** page.

Note that while you can view and manage any cluster from any other cluster as long as they are part of the same multi-cluster group, you can only add and remove clusters from the group when viewing the **primary** cluster.

The **Multi-Cluster Connectivity** page will display all clusters that are part of the multi-cluster group. The **Actions** button will be displayed only when viewing the primary cluster. To modify the cluster group, you will need to navigate to the primary as described in [Navigating Between Clusters](#), at which point the **Actions** button will become available.

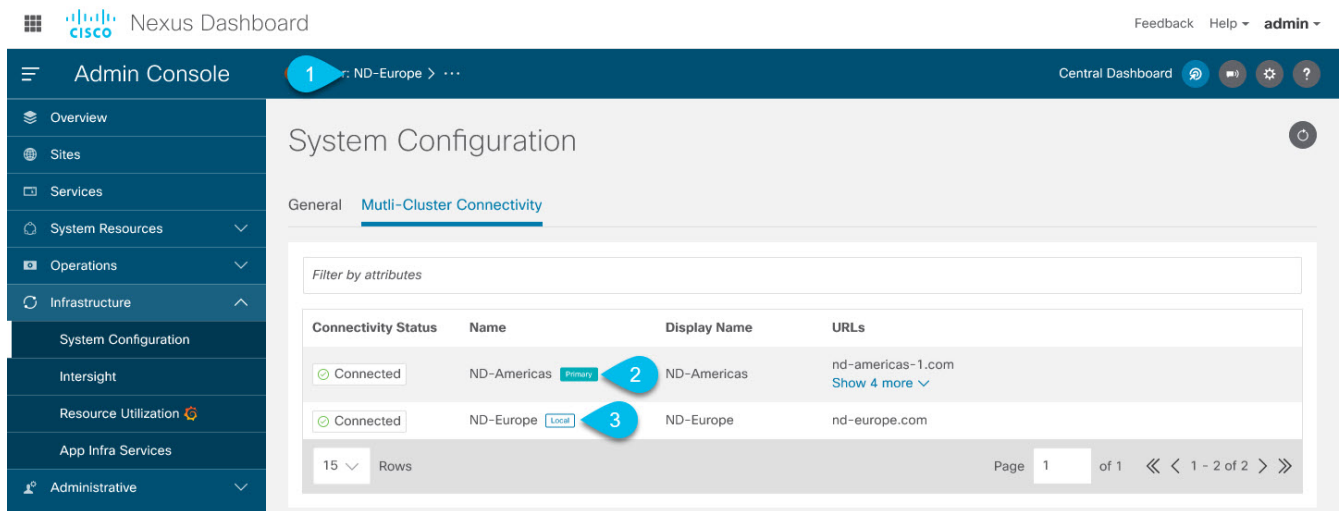


Figure 14. Primary vs Non-primary Clusters

- The **Cluster: <name>** area shows the cluster you are currently viewing.

When you first log into a cluster that is part of a cluster group, it will be displayed here. You can click on the name of the cluster to navigate to and manage a remote cluster that is part of the same group.

- The **Primary** label indicates the group's primary cluster.

You must be viewing this cluster to make any changes to the cluster group, such as adding or removing clusters.

- The **Local** label indicates the cluster you logged into.

This is the cluster whose address is displayed in the browser's URL field. If you navigate to a different cluster as mentioned above, the browser URL and the **Local** label will not change.

Central Dashboard

A central multi-cluster connectivity dashboard UI page becomes available if you connect multiple clusters together and can be accessed by clicking **Central Dashboard** in the top right of any Nexus Dashboard UI page. If you log in to a cluster that is not connected to any other clusters, this UI option

will not be visible.

This page provides an overview and status of the entire system with all clusters, sites, and services across the entire group of clusters you have created and allows you to quickly find obvious issues, such as connectivity loss to one of the clusters:

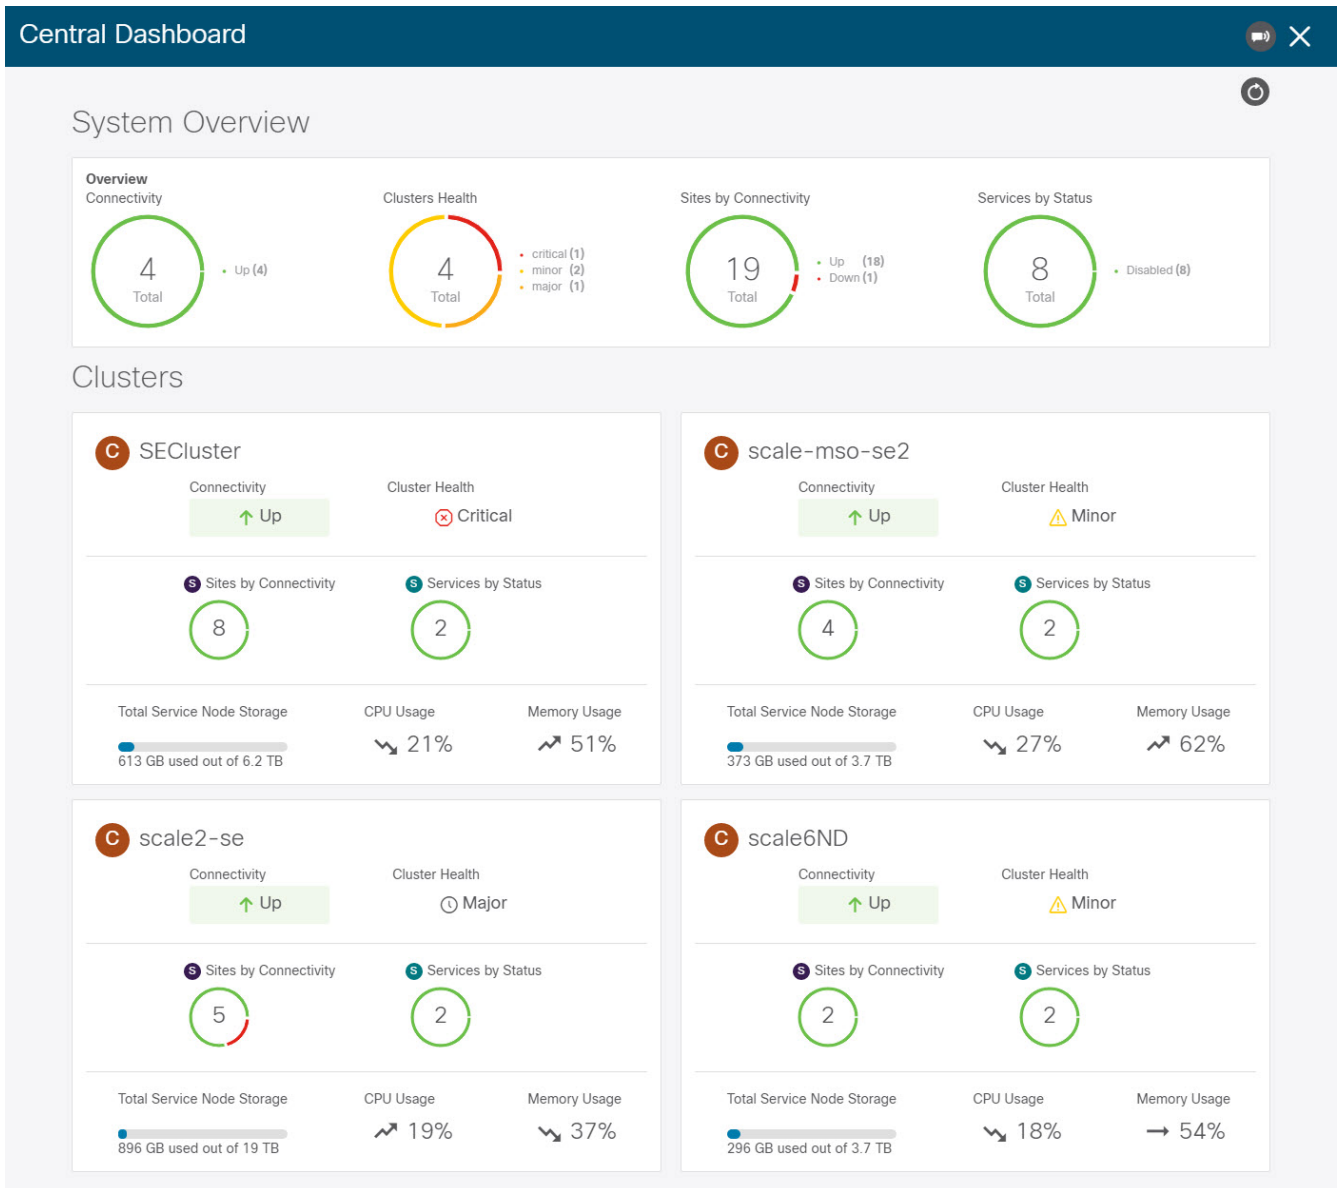


Figure 15. Central Dashboard



The cluster connectivity displayed in this screen indicates connectivity of each cluster to the **primary** cluster only. It does not cover full-mesh connectivity across all clusters in the group.

Navigating Between Clusters

When you connect two or more clusters together, you can view and manage any of the clusters and their sites and services directly from the cluster to which you are already logged in through a single pane of glass.

To change the currently viewed cluster, simply click the cluster name on any of the Nexus Dashboard pages:

The screenshot shows the 'Cluster Configuration' page in the Nexus Dashboard. The left sidebar contains navigation options: Overview, Sites, Services, System Resources, Operations, Infrastructure, Cluster Configuration, Resource Utilization, Intersight, and App Infra Services. The main content area is titled 'Cluster Configuration' and has a sub-tab 'Multi Cluster Connectivity'. Below this, there is a 'Filter by attributes' input field and a 'Connect Cluster' button. A table displays the following data:

Connectivity Status	Name	URL	
Up	se-cluster	172.23.49.118, 172.23.49.117	...
Up	tb100-cluster	172.31.200.113, 172.31.200.114	...

Figure 16. Navigating Between Clusters

After you change the current cluster, you may need to click the **Refresh** button in the top right of the current page to display the information from the cluster you just selected. From here on, you can perform any actions as if you were logged in directly into that cluster.

Disconnecting Clusters

To disconnect a cluster from an existing group:

1. Log in to the Nexus Dashboard GUI of the primary cluster.
 - Adding and removing clusters from the group must be done from the primary cluster.
2. From the main navigation menu, select **Infrastructure > Cluster Configuration**.
3. In the main pane, select the **Multi-Cluster Connectivity** tab.
4. From the **Actions (...)** menu for the cluster you want to remove, select **Disconnect Cluster**
5. In the confirmation window, click **Ok**.

Deploying Additional Physical Nodes

Initial cluster deployment is described in [Nexus Nexus Dashboard Deployment Guide](#). The following sections describe how to deploy an additional physical node so you can add it as a **worker** or **standby** node.



When adding nodes to an existing cluster, the additional nodes must be of the same form factor (such as physical or virtual) as the rest of the nodes in the cluster. This release does not support clusters with nodes of different form factors.

After you deploy an additional node, you can add it to the cluster based on its role:

- For more information about **worker** nodes, see [Managing Worker Nodes](#)
- For more information about **standby** nodes, see [Managing Standby Nodes](#)

Prerequisites and Guidelines for Physical Nodes

- Ensure that you have reviewed and completed the general prerequisites described in the [Platform](#)

[Overview](#), especially the network and fabric connectivity sections.

- Ensure that you have Reviewed and complete any additional prerequisites described in the *Release Notes* for the services you have deployed.

Some services may have additional caveats for **worker** and **standby** nodes. You can find the service-specific documents at the following links:

- [Nexus Dashboard Fabric Controller Release Notes](#)
 - [Nexus Dashboard Insights Release Notes](#)
 - [Nexus Dashboard Orchestrator Release Notes](#)
- For maximum number of **worker** and **standby** nodes in a single cluster, see the [Nexus Dashboard Release Notes](#) for your release.
 - Ensure you are using the supported hardware and the servers are racked and connected.

The physical appliance form factor is supported on the UCS-C220-M5 and UCS-C225-M6 original Nexus Dashboard platform hardware only. The following table lists the PIDs and specifications of the physical appliance servers:

Table 9. Supported UCS-C220-M5 Hardware

PID	Hardware
SE-NODE-G2	<ul style="list-style-type: none">- UCS C220 M5 Chassis- 2x 10 core 2.2G Intel Xeon Silver CPU- 256 GB of RAM- 4x 25G Virtual Interface Card 1455- 4x 2.4TB HDDs- 400GB SSD- 1.2TB NVMe drive- 1050W power supply

Table 10. Supported UCS-C225-M6 Hardware

PID	Hardware
ND-NODE-G4	<ul style="list-style-type: none"> - UCS C225 M6 Chassis - 2.8GHz AMD CPU - 256 GB of RAM - 4x 2.4TB HDDs - 960GB SSD - 1.6TB NVME drive - Intel X710T2LG 2x10 GbE (Copper) - Intel E810XXVDA2 2x25/10 GbE (Fiber Optic) - 1050W power supply



The above hardware supports Nexus Dashboard software only. If any other operating system is installed, the node can no longer be used as a Nexus Dashboard node.

- Ensure that you are running a supported version of Cisco Integrated Management Controller (CIMC).

The minimum supported and recommended versions of CIMC are listed in the "Compatibility" section of the [Release Notes](#) for your Nexus Dashboard release.

- Ensure the hardware is running the same Nexus Dashboard release as your existing cluster.

If the new node is running an earlier release, you must manually upgrade to the current release, as described in [Manual Cluster Upgrades](#).

If for any reason you are unable to run the manual upgrade, you can reinstall the software, as described in [Re-Imaging Nodes](#).

Deploying Physical Nodes

Once you have completed all prerequisites described above, simply connect the node and power it own.

Once the node is deployed, you can add it to the cluster:

- To add the node as a **worker** node, see [Managing Worker Nodes](#)
- To add the node as a **standby** nodes, see [Managing Standby Nodes](#)

Deploying Additional Virtual Nodes in VMware ESX

Initial cluster deployment is described in [Nexus Nexus Dashboard Deployment Guide](#). The following

sections describe how to deploy an additional node in VMware ESX so you can add it as a **worker** or **standby** node.



When adding nodes to an existing cluster, the additional nodes must be of the same form factor (physical or virtual) as the rest of the nodes in the cluster. This release does not support clusters with nodes of different form factors.

After you deploy an additional node, you can add it to the cluster based on its role:

- For more information about **worker** nodes, see [Managing Worker Nodes](#)
- For more information about **standby** nodes, see [Managing Standby Nodes](#)

Prerequisites and Guidelines for ESX Nodes

- Ensure that you reviewed and completed the general prerequisites described in the [Platform Overview](#), especially the network and fabric connectivity sections.
- When deploying in VMware ESX, you can choose to deploy using a vCenter or directly in the ESXi host.

For detailed information, see one of the following sections.

- When deploying in VMware ESX, you can deploy two types of nodes:
 - 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

Table 11. Supported Hardware

Nexus Dashboard Version	Data Node Requirements	App Node Requirements
Release 2.3.1	<p>VMware ESXi 7.0, 7.0.1, 7.0.2, 7.0.3</p> <p>VMware vCenter 7.0.1, 7.0.2 if deploying using vCenter</p> <p>Each VM requires the following:</p> <ul style="list-style-type: none"> ▪ 32 vCPUs with physical reservation of at least 2.2GHz ▪ 128GB of RAM with physical reservation ▪ 3TB SSD storage for the data volume and an additional 50GB for the system volume <p>Data nodes must be deployed on storage with the following minimum performance requirements:</p> <ul style="list-style-type: none"> ○ The SSD must be attached to the data store directly or in JBOD mode if using a RAID Host Bus Adapter (HBA) ○ The SSDs must be optimized for Mixed Use/Application (not Read-Optimized) ○ 4K Random Read IOPS: 93000 ○ 4K Random Write IOPS: 31000 <p>We recommend that each Nexus Dashboard node is deployed in a different ESXi server.</p>	<p>VMware ESXi 7.0, 7.0.1, 7.0.2, 7.0.3</p> <p>VMware vCenter 7.0.1, 7.0.2 if deploying using vCenter</p> <p>Each VM requires the following:</p> <ul style="list-style-type: none"> ▪ 16 vCPUs with physical reservation of at least 2.2GHz ▪ 64GB of RAM with physical reservation ▪ 500GB HDD or SSD storage for the data volume and an additional 50GB for the system volume <p>Some services require App nodes to be deployed on faster SSD storage while other services support HDD. Check the Nexus Dashboard Capacity Planning tool to ensure that you use the correct type of storage.</p> <p>We recommend that each Nexus Dashboard node is deployed in a different ESXi server.</p>

Deploying ESX Node Using vCenter

Before you begin

Ensure that you meet the requirements and guidelines described in [Prerequisites and Guidelines for ESX Nodes](#).

This section describes how to deploy an additional Cisco Nexus Dashboard node in VMware ESX using vCenter.

1. Obtain the Cisco Nexus Dashboard OVA 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. Click the **Download** icon next to the Nexus Dashboard OVA image (**nd-dk9.<version>.ova**).

2. Log in to your VMware vCenter.

Depending on the version of your vSphere client, the location and order of configuration screens may differ slightly. The following steps provide deployment details using VMware vSphere Client 6.7.

3. Start the new VM deployment.



- a. Right-click the ESX host where you want to deploy.
- b. Then select "Deploy OVF Template..." .

The **Deploy OVF Template** wizard appears.

4. In the **Select an OVF template** screen, provide the OVA image, then click **Next**.

Deploy OVF Template

1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select an OVF template

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

URL

http:// /nd-dk9.2.2.0.83.ova

Local file

Choose Files No file chosen

CANCEL

BACK **NEXT**

a. Provide the image.

If you hosted the image on a web server in your environment, select **URL** and provide the URL to the image.

If your image is local, select **Local file** and click **Choose Files** to select the OVA file you downloaded.

b. Click **Next** to continue.

5. In the **Select a name and folder** screen, provide a name and location for the VM.

Deploy OVF Template

1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select a name and folder

Specify a unique name and target location

Virtual machine **a** nd-cluster-vm1

Select a location for the virtual machine.

172.31.141.49

b Datacenter1

CANCEL

BACK **NEXT**

a. Provide the name for the virtual machine.

b. Select the location for the virtual machine.

c. Click **Next** to continue

6. In the **Select a compute resource** screen, select the ESX host.

Deploy OVF Template

✓ 1 Select an OVF template
✓ 2 Select a name and folder
3 Select a compute resource
4 Review details
5 Select storage
6 Ready to complete

Select a compute resource
Select the destination compute resource for this operation

▼ Datacenter1
 > 172.23.136.84
 > 172.23.136.86
 > 172.23.136.87
 > 172.23.136.88

Compatibility
✓ Compatibility checks succeeded.

CANCEL BACK NEXT

a. Select the vCenter datacenter and the ESX host for the virtual machine.

b. Click **Next** to continue

7. In the **Review details** screen, click **Next** to continue.

8. In the **Configuration** screen, select the node profile you want to deploy.

Deploy OVF Template

✓ 1 Select an OVF template
✓ 2 Select a name and folder
✓ 3 Select a compute resource
✓ 4 Review details
5 Configuration
6 Select storage
7 Select networks
8 Customize template
9 Ready to complete

Configuration
Select a deployment configuration

App
 Data

Description
Use this deployment profile to configure an App OVA with 16vCPUs, 64GB RAM, and 500GB Disk.

2 Items

CANCEL BACK NEXT

a. Select either **App** or **Data** node profile based on your use case requirements.

b. For more information about the node profiles, see [Prerequisites and Guidelines for ESX Nodes](#).

c. Click **Next** to continue

9. In the **Select storage** screen, provide the storage information.

Deploy OVF Template

1 Select an OVF template
2 Select a name and folder
3 Select a compute resource
4 Review details
5 Configuration
6 **Select storage**
7 Select networks
8 Customize template
9 Ready to complete

Select storage
Select the storage for the configuration and disk files

Encrypt this virtual machine (Requires Key Management Server)

Select virtual disk format: Thick Provision Lazy Zeroed

VM Storage Policy: Datastore Default

Name	Capacity	Provisioned	Free	Type	Cluster
datastore1 (3)	925.25 GB	225.74 GB	707.7 GB	VMFS 5	

Compatibility
✔ Compatibility checks succeeded.

CANCEL BACK NEXT

- From the **Select virtual disk format** drop-down, select **Thick Provision Lazy Zeroed**.
- Select the datastore for the virtual machine.

We recommend a unique datastore for each node.

- Click **Next** to continue

- In the **Select networks** screen, choose the VM network for the Nexus Dashboard's Management and Data networks and click **Next** to continue.

There are two networks required by the Nexus Dashboard cluster:

- o **fabric0** is used for the Nexus Dashboard cluster's Data Network
- o **mgmt0** is used for the Nexus Dashboard cluster's Management Network.

For more information about these networks, see "Network Connectivity" .

- In the **Customize template** screen, provide the required information.

Deploy OVF Template

✓ 1 Select an OVF template
✓ 2 Select a name and folder
✓ 3 Select a compute resource
✓ 4 Review details
✓ 5 Configuration
✓ 6 Select storage
✓ 7 Select networks
8 Customize template
9 Ready to complete

Customize template
Customize the deployment properties of this software solution.

✓ All properties have valid values

Resource Configuration	1 settings
1. Data Disk Size (GB)	Data disk size (min 500GB, max 1536GB (1.5TB)) 500
Node Configuration	3 settings
1. Password	Local "rescue-user" password Password Confirm Password
2. Management Network Address and subnet	Management network address. Enter IP/subnet 172.31.140.46/24
3. Management Gateway IP	Management network gateway IP address. Enter IP only 172.31.140.

CANCEL **BA e** NEXT

a. Provide the sizes for the node's data disks.

We recommend using the default values for the required data volume.

The default values will be pre-populated based on the type of node you are deploying, with App node having a single 500GB disk and Data node having a single 3TB disk.

Note that in addition to the data volume, a second 50GB system volume will also be configured but cannot be customized.

b. Provide and confirm the **Password**.

This password is used for the **rescue-user** account on each node. We recommend configuring the same password for all nodes, however you can choose to provide different passwords for the second and third node.

c. Provide the **Management Network** IP address, netmask.

d. Provide the **Management Network** IP gateway.

e. Click **Next** to continue.

12. In the **Ready to complete** screen, verify that all information is accurate and click **Finish** to begin deploying the node.

13. Once the VM deployment is finished, power on the VM.

14. Add the node as **master** or **standby**.

Once the node is deployed, you can add it to the cluster:

- o To add the node as a **worker** node, see [Managing Worker Nodes](#)
- o To add the node as a **standby** nodes, see [Managing Standby Nodes](#)

Deploying ESX Node Directly in ESXi

Before you begin

Ensure that you meet the requirements and guidelines described in [Prerequisites and Guidelines for ESX Nodes](#).

This section describes how to deploy an additional Cisco Nexus Dashboard node in VMware ESX using vCenter.

1. Obtain the Cisco Nexus Dashboard OVA 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. Click the **Download** icon next to the Nexus Dashboard OVA image (**nd-dk9.<version>.ova**).

2. Log in to your VMware ESXi.

Depending on the version of your ESXi server, the location and order of configuration screens may differ slightly. The following steps provide deployment details using VMware ESXi 6.7.

3. Right-click the host and select Create/Register VM.

4. In the Select creation type screen, choose Deploy a virtual machine from an OVF or OVA file, then click Next.

5. In the Select OVF and VMDK files screen, provide the virtual machine name (for example, **nd-node-worker1**) and the OVA image you downloaded in the first step, then click Next.

6. In the Select storage screen, choose the datastore for the VM, then click Next.

7. In the Select OVF and VMDK files screen, provide the virtual machine name (for example, **nd-node-worker1**) and the OVA image you downloaded in the first step, then click Next.

8. In the Deployment options screen, choose Disk Provisioning: Thick, uncheck the Power on automatically option, then click Next to continue.

There are two networks, fabric0 is used for the data network and mgmt0 is used for the management network.

9. In the Ready to complete screen, verify that all information is accurate and click Finish to begin deploying the first node.

10. Wait for the VM to finish deploying, ensure that the VMware Tools periodic time synchronization is disabled, then start the VM.

To disable time synchronization:

- a. Right-click the node's VM and select Edit Settings.

- b. In the Edit Settings window, select the VM Options tab.

- c. Expand the VMware Tools category and uncheck the Synchronize guest time with host option.

11. Open the node's console and configure the node's basic information.

- a. Begin initial setup.

You will be prompted to run the first-time setup utility:

```
[ OK ] Started atomix-boot-setup.  
Starting Initial cloud-init job (pre-networking)...  
Starting logrotate...  
Starting logwatch...  
Starting keyhole...  
[ OK ] Started keyhole.  
[ OK ] Started logrotate.  
[ OK ] Started logwatch.  
Press any key to run first-boot setup on this console...
```

b. Enter and confirm the admin password

This password will be used for the rescue-user SSH login and for adding this node to the cluster.

```
Admin Password:  
Reenter Admin Password:
```

c. Enter the management network information.

```
Management Network:  
IP Address/Mask: 192.168.9.172/24  
Gateway: 192.168.9.1
```

d. Review and confirm the entered information.

You will be asked if you want to change the entered information. If all the fields are correct, choose **n** to proceed.

If you want to change any of the entered information, enter **y** to re-start the basic configuration script.

```
Please review the config  
Management network:  
Gateway: 192.168.9.1  
IP Address/Mask: 192.168.9.172/24  
Re-enter config? (y/N): n
```

12. Add the node as **master** or **standby**.

Once the node is deployed, you can add it to the cluster:

- o To add the node as a **worker** node, see [Managing Worker Nodes](#)

- o To add the node as a **standby** nodes, see [Managing Standby Nodes](#)

Deploying Additional Virtual Nodes in Linux KVM

Initial cluster deployment is described in [Nexus Nexus Dashboard Deployment Guide](#). The following sections describe how to deploy an additional node in Linux KVM so you can add it as a **standby** node.



When adding nodes to an existing cluster, the additional nodes must be of the same form factor (physical or virtual) as the rest of the nodes in the cluster. This release does not support clusters with nodes of different form factors.

After you deploy an additional node, you can add it to the cluster as a **standby** node as described in [Managing Standby Nodes](#).

Prerequisites and Guidelines for KVM Nodes

- Ensure that you reviewed and complete the general prerequisites described in the [Platform Overview](#), especially the network and fabric connectivity sections.
- Ensure that your VM has sufficient resources:

Table 12. Supported Hardware

Nexus Dashboard Version	VM Requirements
Release 2.2.x	<ul style="list-style-type: none"> ▪ Supported Linux distribution: <ul style="list-style-type: none"> ◦ For Nexus Dashboard Orchestrator, you must deploy in CentOS Linux ◦ For Nexus Dashboard Fabric Controller, you must deploy in CentOS or Red Hat Enterprise Linux ▪ Supported versions of Kernel and KVM: <ul style="list-style-type: none"> ◦ Kernel <code>3.10.0-957.el7.x86_64</code> or later ◦ KVM <code>libvirt-4.5.0-23.el7_7.1.x86_64</code> or later ▪ 16 vCPUs ▪ 64 GB of RAM ▪ 500 GB disk <p>Each node requires a dedicated disk partition</p> <ul style="list-style-type: none"> ▪ The disk must have I/O latency of 20ms or less. You can verify the I/O latency using the following command: <pre style="margin-left: 20px;"> fiio --rw=write --ioengine=sync --fdatasync=1 --directory=test -data_with_se --size=22m --bs=2300 --name=mytest And confirm that the 99.00th=[<value>] in the fsync/fdatasync/sync_file_range section is under 20ms. </pre> ▪ We recommend that each Nexus Dashboard node is deployed in a different KVM server.

Deploying KVM Nodes

Before you begin

Ensure that you meet the requirements and guidelines described in [Prerequisites and Guidelines for KVM Nodes](#).

This section describes how to deploy an additional Cisco Nexus Dashboard node in Linux KVM.

1. Download the Cisco Nexus Dashboard image.
 - a. Browse to the **Software Download** page.

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

- b. From the left sidebar, choose the Nexus Dashboard version you want to download.

- c. Download the Cisco Nexus Dashboard image for Linux KVM (`nd-dk9.<version>.qcow2`).
2. Copy the image to the Linux KVM servers where you will host the nodes.

If you have already copied the image, for example when initially deploying the cluster, you can use the same base image and skip this step. The following steps assume you copied the image into the `/home/nd-base` directory.

You can use `scp` to copy the image, for example:

```
# scp nd-dk9.<version>.qcow2 root@<kvm-host-ip>:/home/nd-base
```

3. Create the required disk images for the node.

You will create a snapshot of the base `qcow2` image you downloaded and use the snapshots as the disk images for the node's VM. You will also need to create a second disk image for the node.

- a. Log in to your KVM host as the `root` user.
- b. Create a directory for the node's snapshot.

The following steps assume you create the snapshot in the `/home/nd-node1` directory.

```
# mkdir -p /home/nd-node1/  
# cd /home/nd-node1
```

- c. Create the snapshot.

In the following command, replace `/home/nd-base/nd-dk9.<version>.qcow2` with the location of the base image you created in the previous step.

```
# qemu-img create -f qcow2 -b /home/nd-base/nd-dk9.<version>.qcow2  
/home/<node-name>/nd-node1-disk1.qcow2
```

The following steps assume you are adding `nd-node4`.

- d. Create the additional disk image for the node.

Each node requires two disks: a snapshot of the base Nexus Dashboard `qcow2` image and a second 500GB disk.

```
# qemu-img create -f qcow2 /home/nd-node1/nd-node4-disk2.qcow2 500G
```

Before you proceed to the next step, you should have the following:

- `/home/nd-node4/nd-node4-disk1.qcow2`, which is a snapshot of the base `qcow2` image you downloaded in Step 1.
- `/home/nd-node4/nd-node4-disk2.qcow2`, which is a new 500GB disk you created.

4. Create the first node's VM.

You can use CLI or KVM GUI to create the VM with the following configuration:

- o 16 vCPUs
- o 64GB of RAM
- o Operating system type set to **linux2020**
- o Network device model set to **virtio**
- o Management interface mapped to bus **0x00** and slot **0x03** and Data interface mapped to bus **0x00** and slot **0x04**



Nexus Dashboard expects the Management interface to be connected to bus **0x00** and slot **0x03** and the Data interface to bus **0x00** and slot **0x04**. If this is not the case, the cluster will not have network connectivity.

For example, to create the VM using CLI:

```
# virt-install --name <node-name> \  
  --vcpus 16 --ram 64000 --osinfo linux2020 \  
  --disk path=/home/nd-node4/nd-node4-disk1.qcow2 \  
  --disk path=/home/nd-node4/nd-node4-disk2.qcow2 \  
  --network bridge:br-  
oob,model=virtio,address.type=pci,address.domain=0,address.bus=0,address.slot=3 \  
  --network bridge:br-  
vnd,model=virtio,address.type=pci,address.domain=0,address.bus=0,address.slot=4 \  
  --noautoconsole --import
```

5. Open the node's console and configure the node's basic information.

- a. Press any key to begin initial setup.

You will be prompted to run the first-time setup utility:

```
[ OK ] Started atomix-boot-setup.  
Starting Initial cloud-init job (pre-networking)...  
Starting logrotate...  
Starting logwatch...  
Starting keyhole...  
[ OK ] Started keyhole.  
[ OK ] Started logrotate.  
[ OK ] Started logwatch.  
Press any key to run first-boot setup on this console...
```

- b. Enter and confirm the **admin** password.

This password will be used for the **rescue-user** SSH login as well as the initial GUI password.

Admin Password:
Reenter Admin Password:

c. Enter the management network information.

Management Network:
IP Address/Mask: 192.168.9.172/24
Gateway: 192.168.9.1

d. When asked if the node is the "Cluster Leader", choose "no".

Since you are adding a **worker** or **standby** node, do not designate it as the cluster leader

Is cluster leader?: n

e. Review and confirm the entered information.

You will be asked if you want to change the entered information. If all the fields are correct, enter **n** to proceed. If you want to change any of the entered information, enter **y** to re-start the basic configuration script.

Please review the config
Management network:
Gateway: 192.168.9.1
IP Address/Mask: 192.168.9.172/24
Cluster leader: no

6. Wait for the initial bootstrap process to complete.

After you provide and confirm management network information, wait for the initial bootstrap process to complete:

Please wait for system to boot: [#####] 100%
System up, please wait for UI to be online.

System UI online, please login to https://192.168.9.172 to continue.

7. Add the node to the cluster as **master** or **standby**.

Once the bootstrap process is finished, you can add it to the cluster:

- o To add the node as a **worker** node, see [Managing Worker Nodes](#)
- o To add the node as a **standby** nodes, see [Managing Standby Nodes](#)

Managing Worker Nodes

You can add a number of worker nodes to an existing 3-node cluster for horizontal scaling to enable application co-hosting.

For additional information about application co-hosting and cluster sizing, see the [Platform Overview](#) section of this document.



Worker nodes are not supported for cloud form factors of Nexus Dashboard clusters deployed in AWS or Azure.

Adding Worker Nodes

This section describes how to add a worker node to your cluster to enable horizontal scaling

Before you begin

- Ensure that the existing master nodes and the cluster are healthy.
- Prepare and deploy the new node as described in [Deploying Additional Physical Nodes](#), [Deploying Additional Virtual Nodes in VMware ESX](#), [Deploying ESX Node Directly in ESXi](#), or [Deploying Additional Virtual Nodes in Linux KVM](#).
- Ensure that the node you are adding is powered on.
- If you are adding a physical node, ensure that you have the new node's CIMC IP address and login information.

You will need to use the CIMC information to add the new node using the Nexus Dashboard GUI.

- If you are adding a virtual node, ensure that you have the node's management IP address and login information.

To add a worker node:

1. Log in to the Cisco Nexus Dashboard GUI.
2. From the main navigation menu, select **System Resources > Nodes**.
3. In the main pane, click **Add Node**.

The **Add Node** screen opens.

4. In the **Add Node** screen, provide the node information.
 - a. Provide the **Name** of the node.
 - b. From the **Type** dropdown, select **Worker**.
 - c. Provide the **Credentials** information for the node, then click **Verify**.

For physical nodes, this is the IP address, username, and password of the server's CIMC. The CIMC will be used to configure the rest of the information on the node.

For virtual nodes, this is the IP address and **rescue-user** password you defined for the node when deploying it.

- d. Provide the **Management Network** information.

For virtual nodes, the management network information will be pre-populated with the information pulled from the node based on the IP address and credentials you provided in the previous sub-step.

For physical nodes, you must provide the management network IP address, netmask, and gateway now.

e. Provide the **Data Network** information.

You must provide the data network IP address, netmask, and gateway. Optionally, you can also provide the VLAN ID for the network. For most deployments, you can leave the VLAN ID field blank.

f. (Optional) Provide IPv6 information for the management and data networks.

Starting with release 2.1.1, Nexus Dashboard supports dual stack IPv4/IPv6 for the management and data networks.

If you want to provide IPv6 information, you must do it when adding the node.

All nodes in the cluster must be configured with either only IPv4 or dual IPv4/IPv6 stack.

5. Click **Save** to add the node.

The configuration will be pushed to the node and the node will be added to the list in the GUI.

1. If you are running Nexus Dashboard Insights application, disable and re-enable the application.

After you add the new worker node, you must disable and re-enable the application for its services to be properly distributed to the new node.

Deleting a Worker node

Before you begin

- Ensure that the master nodes and the cluster are healthy.

To delete an existing worker node:

1. Log in to the Cisco Nexus Dashboard GUI.
2. From the main navigation menu, select **System Resources > Nodes**.
3. Select the checkbox next to the worker node you want to delete.
4. From the **Actions** menu, choose **Delete** to delete the node.

Managing Standby Nodes

You can add up to two standby nodes, which you can use to quickly restore the cluster functionality in case one or more master nodes fail by replacing the failed master node with the standby node.

Standby nodes are similar to worker nodes in deployment, initial configuration, and upgrades. However, unlike worker nodes, the cluster will not use the standby nodes for any workloads.



Standby nodes are not supported for single-node clusters or clusters deployed in

AWS or Azure.

The following two cases are supported:

- Single master node failure

You can use the UI to convert the standby node into a new master node.

- Two master nodes failure

You will need to perform manual failover of one of the nodes to restore cluster functionality. Then fail over the second node using standard procedure.

Adding Standby Nodes

This section describes how to add a standby node to your cluster for easy cluster recover in case of a master node failure.

Before you begin

- Ensure that the existing master nodes and the cluster are healthy.
- Prepare and deploy the new node as described in [Deploying Additional Physical Nodes](#), [Deploying Additional Virtual Nodes in VMware ESX](#), [Deploying ESX Node Directly in ESXi](#), or [Deploying Additional Virtual Nodes in Linux KVM](#).

You can failover only between nodes of identical types (physical or virtual), so you must deploy the same type of node as the nodes in your cluster which you may need to replace. In case of virtual nodes deployed in VMware ESX, which have two node profiles (**OVA-app** and **OVA-data**), you can failover only between nodes of the same profile.

- Ensure that the node you are adding is powered on.
- If you are adding a physical node, ensure that you have the new node's CIMC IP address and login information.

You will need to use the CIMC information to add the new node using the Nexus Dashboard GUI.

- If you are adding a virtual node, ensure that you have the node's management IP address and login information.

To add a standby node:

1. Log in to the Cisco Nexus Dashboard GUI.
2. From the main navigation menu, select **System Resources > Nodes**.
3. In the main pane, click **Add Node**.

The **Add Node** screen opens.

4. In the **Add Node** screen, provide the node information.
 - a. Provide the **Name** of the node.
 - b. From the **Type** dropdown, select **Standby**.
 - c. Provide the **Credentials** information for the node, then click **Verify**.

For physical nodes, this is the IP address, username, and password of the server's CIMC. The CIMC will be used to configure the rest of the information on the node.

For virtual nodes, this is the IP address and **rescue-user** password you defined for the node when deploying it.

d. Provide the **Management Network** information.

For virtual nodes, the management network information will be pre-populated with the information pulled from the node based on the IP address and credentials you provided in the previous sub-step.

For physical nodes, you must provide the management network IP address, netmask, and gateway now.

e. Provide the **Data Network** information.

You must provide the data network IP address, netmask, and gateway. Optionally, you can also provide the VLAN ID for the network. For most deployments, you can leave the VLAN ID field blank.

f. (Optional) Provide IPv6 information for the management and data networks.

Starting with release 2.1.1, Nexus Dashboard supports dual stack IPv4/IPv6 for the management and data networks.

If you want to provide IPv6 information, you must do it when adding the node.

All nodes in the cluster must be configured with either only IPv4 or dual IPv4/IPv6 stack.

5. Click **Save** to add the node.

The configuration will be pushed to the node and the node will be added to the list in the GUI.

Replacing Single Master Node with Standby Node

This section describes failover using a pre-configured **standby** node. If your cluster does not have a standby node, follow the steps described in one of the sections in [Troubleshooting](#) instead.

Before you begin

- Ensure that at least 2 master nodes are healthy.

If two of the master nodes are unavailable, you will need to manually restore the cluster as described in [Replacing Two Master Nodes with Standby Nodes](#)

- Ensure that you have at least one **standby** node available in the cluster.

Setting up and configuring **standby** nodes is described in [Adding Standby Nodes](#).

- Ensure that the **master** node you want to replace is powered off.



You cannot re-add the **master** node you are replacing back to the cluster after the failover is complete. If the **master** node you replace is still functional and you want to re-add it to the cluster after the failover, you must factory reset or re-

image it as described in [Troubleshooting](#) and add it as a **standby** or **master** node only.

To failover a single master node:

1. Log in to the Cisco Nexus Dashboard GUI.
2. From the main navigation menu, select **System Resources > Nodes**.
3. Click the **Actions (...)** menu next to the **Inactive** master node that you want to replace.
4. Choose **Failover**.

Note that you must have a standby node already configured and added or the **Failover** menu option will not be available.

5. In the **Fail Over** window that opens, select a standby node from the dropdown.
6. Click **Save** to complete the failover.

The failed master node will be removed from the list and replaced by the standby node you selected. The status will remain **Inactive** while the services are being restored to the new master node.

It can take up to 10 minutes for all services to be restored, at which point the new master node's status will change to **Active**.

Replacing Two Master Nodes with Standby Nodes

This section describes failover using a pre-configured **standby** node. If your cluster does not have a standby node, follow the steps described in one of the sections in [Troubleshooting](#) instead.

If only one of your master nodes failed, you can use the GUI to replace it with a standby node as described in [Replacing Single Master Node with Standby Node](#).

However, when two master nodes are unavailable, the cluster goes offline. In this case, most operations including the UI are disabled and no configuration changes can be made to the cluster. You can still SSH into the remaining master node as the **rescue-user**, which is used to recover the cluster by manually failing over one of the failed master nodes to a standby node. Once two **master** nodes are available again, the cluster can resume normal operation, at which point you can recover the second master node using the normal procedure.

Before you begin

- Ensure that you have at least one **standby** node available in the cluster.

Setting up and configuring **standby** nodes is described in [Adding Standby Nodes](#).

- Ensure that the **master** nodes you want to replace are powered off.



You cannot re-add the **master** node you are replacing back to the cluster after the failover is complete. If the **master** node you replace is still functional and you want to re-add it to the cluster after the failover, you must factory reset or re-image it as described in [Troubleshooting](#) and add it as a **standby** or **master** node only.

- If you had installed the Nexus Dashboard Fabric Controller (NDFC) service in the cluster, you must have a configuration backup available to restore after you recover the cluster.

The Fabric Controller service cannot recover from a two **master** node failure of the Nexus Dashboard cluster where it is running. After you recover the cluster, you must re-install the NDFC service and restore its configuration from a backup.

To fail over two master nodes:

1. Log in to the remaining master node via CLI as **rescue-user**.
2. Execute the failover command.

In the following command, replace **<node1-data-ip>** and **<node2-data-ip>** with the data network IP addresses of the failed nodes:

```
# acs failover --failedIP <node1-data-ip> --failedIP <node2-data-ip>
```



Even though only the first node is failed over, the second failed node you provide is required internally to recover the cluster.

By default, the healthy master node will automatically pick an available standby node and fail over the first failed node you provide (**<node1-data-ip>**) to it.

If you would like to provide a specific standby node, you can add **<standby-node-data-ip>** to the above command:

```
# acs failover --failedIP <node1-data-ip> --failedIP <node2-data-ip> \  
--standbyIP <standby-node1-data-ip>
```

3. Confirm that you want to proceed.

Warning: Failover can be a disruptive operation and should only be performed as last resort option to recover cluster from disasters using standby where two master nodes have lost their state due to hardware faults.

Proceed? (y/n): y

The master node will copy the configuration state to the standby node and both nodes will restart. It may take up to 30 minutes for the nodes to come up and the cluster to be restored. You can check the progress by navigating to the master node's UI.

4. If necessary, re-install the NDFC service and restore NDFC configuration.

We recommend fully recovering the cluster by replacing the 3rd **master** node before any configuration changes. However, if you must recover your production NDFC configuration as soon as possible, you can do so now:

- a. Using a browser log in to one of the two active **master** nodes of the ND cluster.

- b. Disable the NDFC service.

This is described in [Disabling Services](#).

- c. Delete the NDFC service.

This is described in [Uninstalling Services](#).

- d. Re-install the NDFC service.

This is described in [Services Management](#).

This step assumes all installation pre-requisites are completed from the initial service deployment. For detailed information on all NDFC requirements, see the [Nexus Dashboard Fabric Controller Installation Guide](#) for your release.

- e. Restore NDFC configuration from a backup.

This is described in the "[Operations > Backup and Restore](#)" chapter of the [NDFC Fabric Controller Configuration Guide](#) for your release.

5. After the cluster is back up, fail over the second failed master node.

At this point, you can use the standard procedure described in [Replacing Single Master Node with Standby Node](#). If you do not have a second **standby** node, you can add it to the cluster while it has only 2 **master** nodes, as described in [Adding Standby Nodes](#).

Deleting Standby Nodes

Before you begin

- Ensure that the master nodes and the cluster are healthy.

To delete an existing standby node:

1. Log in to the Cisco Nexus Dashboard GUI.
2. From the main navigation menu, select **System Resources > Nodes**.
3. Select the checkbox next to the standby node you want to delete.
4. From the **Actions** menu, choose **Delete** to delete the node.

Administrative

You can choose how the users logging into the Nexus Dashboard GUI are authenticated. This release supports local authentication as well as LDAP, RADIUS, and TACACS remote authentication servers. User roles and permissions are described in this section, remote authentication configuration is described in [Remote Authentication](#), and local user configuration is described in [Users](#).

Roles and Permissions

Cisco Nexus Dashboard allows user access according to roles defined by role-based access control (RBAC). Roles are used in both local and external authentication and apply to either the Nexus Dashboard, the services running in it, or both. All roles can be assigned with **read-only** or **write** privileges. Read-only access allows the user to view objects and configurations, while write access allows them to make changes.

The following sections describes the user roles available in Nexus Dashboard and their associated permissions within the platform as well as the individual services.

The same roles can be configured on a remote authentication server and the server can be used to authenticate the Nexus Dashboard users. Additional details about remote authentication are available in the [Remote Authentication](#) section.

Nexus Dashboard and Orchestrator Roles

User Role	ND Platform	Orchestrator Service
Administrator	Provides full access to all settings, features, and tasks. The only role that allows adding and removing services.	Full access.
Approver	Same as Dashboard role.	Allows approval or denial of template configurations; does not allow editing or deploying templates.
Dashboard User	Allows access to the Dashboard view and launching applications; does not allow any changes to the Nexus Dashboard configurations.	No access.
Deployer	Same as Dashboard role.	Allows the user to deploy templates to sites; does not allow editing or approving templates.
Policy Manager	Same as Dashboard role.	No access.

User Role	ND Platform	Orchestrator Service
Site Administrator	Allows access to configurations related to the site on-boarding and configuration.	Allows changing the site status between managed and unmanaged , as well as fabric resource template, fabric policy template, and monitoring template (access SPAN) configurations.
Site Manager	Allows access to deployment of policies to the site.	Allows policy, schema, and monitoring template (tenant SPAN) configurations.
Tenant Manager	Same as Dashboard role.	Allows tenant policy, schema, and monitoring template (tenant SPAN) configurations.
User Manager	Allows access to users settings, such as creating users, changing permissions, and adding remote authentication providers.	No access.

Each role above is associated with a set of permissions, which in turn are used to show relevant and hide not relevant elements from the user's view.


Nexus Dashboard Insights

The Insights service does not support RBAC and any account that can log in to the Nexus Dashboard has full access to Insights.

Nexus Dashboard Fabric Controller Roles

User Role	Nexus Dashboard Fabric Controller
NDFC Access Admin	<p>Allows you to perform operations related to network interfaces in NDFC's Interface Manager screen.</p> <p>An Access Admin user can perform the following actions:</p> <ul style="list-style-type: none"> ▪ Add, edit, delete and deploy layer 2 port channels, and vPC. ▪ Edit host vPC, and ethernet interfaces. ▪ Save, preview, and deploy from management interfaces. ▪ Edit interfaces for LAN classic, and external fabrics if it isn't associated with policy. Except for nve, management, tunnel, subinterface, SVI, interface grouping, and loopback interfaces <p>However, an Access Admin user cannot perform the following actions:</p> <ul style="list-style-type: none"> ▪ Cannot edit layer 3 port channels, ST FEX, AA FEX, loopback interfaces, nve interfaces, and subinterfaces ▪ Cannot edit member interfaces and port channels of Layer 3, ST FEX, AA FEX ▪ Cannot edit interfaces with policy associated from underlay and link ▪ Cannot edit peer link port channel ▪ Cannot edit management interface ▪ Cannot edit tunnels
NDFC Device Upgrade Admin	Allows you to perform operations related to device upgrades in NDFC's Image Management screen.
NDFC Network Admin	Allows full administrative access.

User Role	Nexus Dashboard Fabric Controller
NDFC Network Operator	<p data-bbox="469 170 1161 203">Allows read-only access the following NDFC menus:</p> <ul data-bbox="496 241 683 443" style="list-style-type: none"><li data-bbox="496 241 663 275">▪ Dashboard<li data-bbox="496 297 644 331">▪ Topology<li data-bbox="496 353 624 387">▪ Monitor<li data-bbox="496 409 683 443">▪ Applications <p data-bbox="469 481 1107 515">A Network Operator user can view the following:</p> <ul data-bbox="496 553 818 801" style="list-style-type: none"><li data-bbox="496 553 699 586">▪ Fabric builder<li data-bbox="496 609 710 642">▪ Fabric settings<li data-bbox="496 665 818 698">▪ Preview configurations<li data-bbox="496 721 624 754">▪ Policies<li data-bbox="496 777 659 810">▪ Templates <p data-bbox="469 848 1426 882">However, a Network Operator user cannot perform the following actions:</p> <ul data-bbox="496 920 1461 1095" style="list-style-type: none"><li data-bbox="496 920 1461 954">▪ Cannot change expected configurations of any switch within any fabric<li data-bbox="496 976 1115 1010">▪ Cannot deploy any configurations to switches<li data-bbox="496 1032 1461 1095">▪ Cannot access the administration options like licensing, creating more users, and so on

User Role	Nexus Dashboard Fabric Controller
NDFC Network Stager	<p>Allows you to make configuration changes, but a Network Admin user will need to deploy the changes later.</p> <p>A Network Stager user can perform the following actions:</p> <ul style="list-style-type: none"> ▪ Edit interface configurations ▪ View or edit policies ▪ Create interfaces ▪ Change fabric settings ▪ Edit or create templates <p>However, a Network Stager user cannot perform the following actions:</p> <ul style="list-style-type: none"> ▪ Cannot make any configuration deployments to switches ▪ Cannot perform deployment-related actions from the DCNM Web UI or the REST APIs ▪ Cannot access the administration options like licensing, creating more users, and so on ▪ Cannot move switches in and out of maintenance mode ▪ Cannot move fabrics in and out of deployment-freeze mode ▪ Cannot install patches ▪ Cannot upgrade switches ▪ Cannot create or delete fabrics ▪ Cannot import or delete switches <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  <p>A Network Stager can only define intent for existing fabrics, but cannot deploy those configurations. A Network Admin can deploy the changes and edits that are staged by a user with the Network Stager role.</p> </div>

Choosing Default Authentication Domain

By default, the login screen will select the **local** domain for user authentication; you can manually change the domain at login time by selecting any of the available login domains from the dropdown menu.

Alternatively, you can set a different default login domain to the most commonly used as follows:



The domain must already exist before you can set it as the default domain. Adding remote authentication domains is described in [Remote Authentication](#).

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. Choose the default login domain.

- a. From the main navigation menu, select **Administrative > Authentication**.
- b. In top right of the **Default Authentication** tile, click the **Edit** icon.

The **Default Authentication** window opens.

3. In the **Default Authentication** that opens, choose the **Login Domain** from the dropdown.

Remote Authentication

Cisco Nexus Dashboard supports a number of remote authentication providers, including LDAP, TACACS, and Radius.

When configuring external authentication servers:

- You must configure each user on the remote authentication servers.
- All LDAP configurations are case sensitive.

For example, if you have **OU=Cisco Users** on the LDAP server and **OU=cisco users** on the Nexus Dashboard, the authentication will not work.

- For LDAP configurations, we recommend using **CiscoAVPair** as the attribute string. If, for any reason, you are unable to use an Object ID **1.3.6.1.4.1.9.22.1**, an additional Object IDs **1.3.6.1.4.1.9.2742.1-5** can also be used in the LDAP server.

Alternatively, instead of configuring the Cisco AVPair values for each user, you can create LDAP group maps in the Nexus Dashboard.

- Single sign-on (SSO) between the Nexus Dashboard, sites, and applications is available for remote users only.
- When using SSO to cross-launch into an APIC site from your Nexus Dashboard's **Sites** page, the AV pairs defined for the Nexus Dashboard user are also used when logging into the APIC.

For example, a user defined as **admin** for the Nexus Dashboard cluster will also have **admin** privileges in the APIC.

Configuring Remote Authentication Server

When configuring the remote authentication server for the Nexus Dashboard users, you must add a custom attribute-value (AV) pair, specifying the username and the roles assigned to them.

The user roles and their permissions are the same as for the local users you would configure directly in the Nexus Dashboard GUI as described in [Roles and Permissions](#).

The following tables list the Nexus Dashboard user roles and the AV pair you would use to define the roles on a remote authentication server, such as LDAP.

Table 13. Nexus Dashboard AV Pairs

User Role	AV Pair Value
Administrator	admin

User Role	AV Pair Value
Approver	approver
Dashboard User	app-user
Deployer	deployer
Policy Manager	config-manager
Site Administrator	site-admin
Site Manager	site-policy
Tenant Manager	tenant-policy
User Manager	aaa

Table 14. Nexus Dashboard Fabric Controller AV Pairs

User Role	AV Pair Value
NDFC Access Admin	access-admin
NDFC Device Upgrade Admin	device-upg-admin
NDFC Network Admin	network-admin
NDFC Network Operator	network-operator
NDFC Network Stager	network-stager

The AV pair string format differs when configuring a read-write role, read-only role, or a combination of read-write and read-only roles for a specific user. A typical string includes the domain, followed by the read-write roles separated from the read-only roles using the slash (/) character; individual roles are separated by the pipe (|) character:

```
shell:domains=<domain>/<writeRole1>|<writeRole2>/<readRole1>|<readRole2>
```

For example, the following string illustrates how to assign the **Tenant Manager** and **Policy Manager** roles to a user, while still allowing them to see objects visible to the **User Manager** users:

```
shell:domains=all/tenant-policy|site-policy/aaa
```

Note that if you want to configure only the read-only or only read-write permissions for a user, you must still include the slash (/) character. The following examples show how to set just the read-write or read-only access to the objects available to **Site Administrator** role:

- Read-only: `shell:domains=all//site-admin`
- Read-write: `shell:domains=all/site-admin/`

Adding LDAP as Remote Authentication Provider

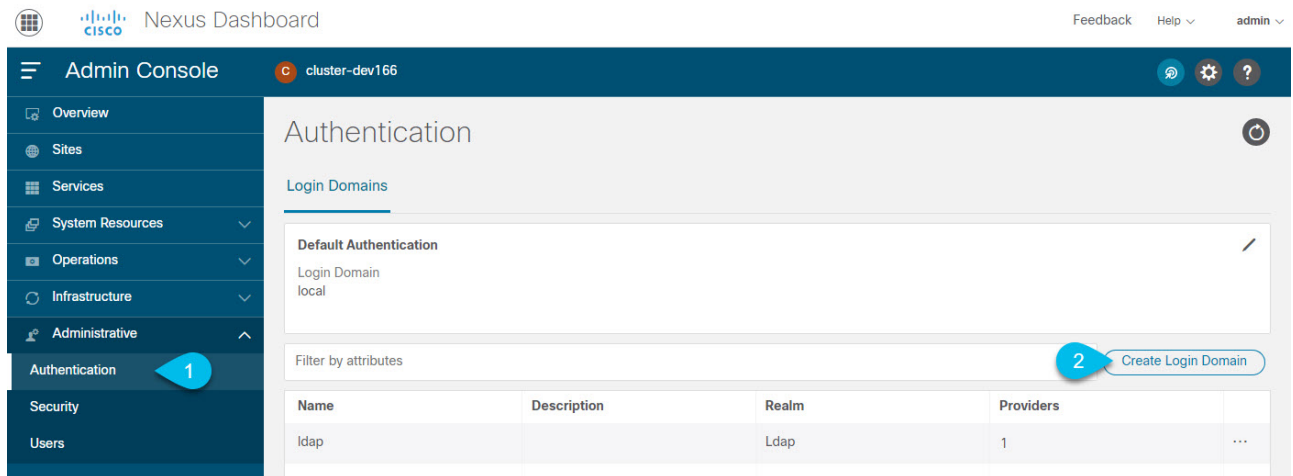
Before you begin

- You must have at least one user already configured on the LDAP server as described in [Configuring Remote Authentication Server](#).

You will need to use an existing user for end-to-end verification of LDAP configuration settings.

To add an LDAP remote authentication provider:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. Add an authentication domain.



- a. From the main navigation menu, select **Administrative > Authentication**.
 - b. In top right of the main pane, click the **Actions** menu and select **Create Login Domain**.
3. In the **Create Login Domain** screen that opens, provide domain details.
 - a. Provide the **Name** for the domain.
 - b. (Optional) Provide its **Description**.
 - c. From the **Realm** dropdown, select **Ldap**.
 - d. Then click **+Add Provider** to add a remote authentication server.

The **Add Provider** window opens.

4. Provide the remote authentication server details.
 - a. Provide the **Hostname** or **IP Address** of the server.
 - b. (Optional) Provide the **Description** of the server.
 - c. Provide the **Port** number.

The default port is **389** for LDAP.

- d. Provide the **Base DN** and **Bind DN**.

The Base DN and Bind DN depend on how your LDAP server is configured. You can get the Base DN and Bind DN values from the distinguished name of the user created on the LDAP server.

Base DN is the point from which the server will search for users. For example, **DC=nd,DC=local**.

Bind DN is the credentials used to authenticate against the server. For example, **CN=admin,CN=Users,DC=nd,DC=local**.

- e. Provide and confirm the **Key**.

This is the password for your Bind DN user. Anonymous bind is not supported, so you must provide a valid value in these fields.

- f. Specify the **Timeout** and number of **Retries** for connecting to the authentication server.

- g. Provide the **LDAP Attribute** field for determining group membership and roles.

The following two options are supported:

- **ciscoAVPair** (default)—used for LDAP servers configured with Cisco AVPair attributes for user roles.
- **memberOf**—used for LDAP servers configured with LDAP group maps. Adding a group map is described in a following step.

- h. (Optional) Enable **SSL** for LDAP communication.

If you enable SSL, you must also provide the **SSL Certificate** and the **SSL Certificate Validation** type:

- **Permissive**: Accept a certificate signed by any certificate authority (CA) and use it for encryption.
- **Strict**: Verify the entire certificate chain before using it.

- i. (Optional) Enable **Server Monitoring**.

If you choose to enable monitoring, you must also provide the **Username** and **Password** for it.

- j. In the **Validation** fields, provide a **Username** and **Password** of a user already configured on the LDAP server you are adding.

Nexus Dashboard will use this user to verify the end-to-end authentication to ensure that the settings you provided are valid.

- k. Click **Save** to complete provider configuration.

- l. Repeat this step for any additional LDAP authentication servers you want to use with this domain.

5. (Optional) Enable and configure **LDAP Group Map Rules**.

If you want to authenticate your LDAP users using Cisco AV pair strings, skip this step.

- a. In the **LDAP Auth Choice**, select **LDAP Group Map Rules**.
- b. Click **Add LDAP Group Map Rule**.

The **Add LDAP Group Map Rule** window opens.

- c. Provide the **Group DN** for the group.
 - d. Select one or more **Roles** for the group.
 - e. Click **Save** to save the group configuration.
 - f. Repeat this step for any additional LDAP groups.
6. Click **Create** to finish adding the domain.

Adding Radius or TACACS as Remote Authentication Provider


Before you begin

- You must have at least one user already configured on the remote authentication server as described in [Configuring Remote Authentication Server](#).

You will need to use an existing user for end-to-end verification of the provider configuration settings.

To add a Radius or TACACS remote authentication provider:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. Add an authentication domain.



- a. From the main navigation menu, select **Administrative > Authentication**.
 - b. In top right of the main pane, click the **Actions** menu and select **Create Login Domain**.
3. In the **Create Login Domain** screen that opens, provide domain details.
 - a. Provide the **Name** for the domain.
 - b. (Optional) Provide its **Description**.
 - c. From the **Realm** dropdown, select **Radius** or **Tacacs**.
 - d. Then click **+Add Provider** to add a remote authentication server.

The **Add Provider** window opens.

4. Provide the remote authentication server details.
 - a. Provide the **Hostame** or **IP Address** of the server.
 - b. (Optional) Provide the **Description** of the server.

c. Choose **Authorization Protocol** used by the server.

You can choose **PAP**, **CHAP**, or **MS-CHAP**.

d. Provide the **Port** number.

The default port is **1812** for RADIUS and **49** for TACACS

e. Provide and confirm the **Key**.

This is the password used for connecting to the provider server.

f. (Optional) Choose whether you want to enable **Server Monitoring**.

If you choose to enable monitoring, you must also provide the **Username** and **Password** for it.

g. In the **Validation** fields, provide a **Username** and **Password** of a user already configured on the remote server you are adding.

Nexus Dashboard will use this user to verify the end-to-end authentication to ensure that the settings you provided are valid.

h. Click **Save** to complete provider configuration.

i. Repeat this step for any additional remote authentication servers.

5. Click **Create** to finish adding the domain.

Validating Remote User Logins

Nexus Dashboard provides a way to validate reachability of the remote authentication provider by performing a login attempt using a specific user's credentials.

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

2. Navigate to the domain you want to test.

The screenshot shows the Nexus Dashboard Admin Console interface. The main navigation menu on the left includes Overview, Sites, Services, System Resources, Operations, Infrastructure, Administrative, Authentication (highlighted with callout 'a'), Security, and Users. The main content area is titled 'Authentication' and shows 'Login Domains'. A table lists the domains:

Name	Description	Realm
local		
IdapTest		Ldap
radiusTest		Radius

The right sidebar shows the configuration for the selected 'IdapTest' domain. It includes sections for General (Description, Realm) and Settings (Providers, Host Name). A blue callout 'c' points to the details icon in the top right of the sidebar.

a. From the main navigation menu, select **Administrative > Authentication**.

b. Click on a specific domain.

c. In the right properties sidebar, click the details icon.

The domain's **Overview** page opens.

3. In the **Overview** page, click **Validate** next to the provider you want to test.
4. In the **Validate Provider** window, enter the **Username** and **Password** of a user defined in this authentication provider and click **Validate**

You will see a message indicating whether authentication was successful or not.

If authentication failure message is displayed, ensure that the authentication provider server is reachable and the user credentials you used to test are valid and configured on the provider.

Editing Remote Authentication Domains

If you want to make changes to a domain you have created:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the main navigation menu, select **Administrative > Authentication**.
3. From the **Actions** menu for the domain, select **Edit Login Domain**.



You cannot change the name and the type of the authentication domain, but you can make changes to the description and provider configuration.



If you make any changes to the login domain, including simply updating the description, you must re-enter the **key** for all existing providers.

Deleting Remote Authentication Domains

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the main navigation menu, select **Administrative > Authentication**.
3. From the **Actions** menu for the domain, select **Delete Login Domain**.



4. In the **Confirm Delete** prompt, click **OK** to confirm.

Multi-Factor Authentication

Starting with Release 2.1.2, you can configure your Nexus Dashboard to use multi-factor authentication (MFA) for user login.

When configuring multi-factor authentication:

- You will configure each user in your MFA provider, as described in [Configuring Okta Account as MFA Provider](#)

This release supports only Okta as MFA provider.

- You will establish MFA provider and client integration, as described in [Configuring MFA Client](#).

This release supports only Duo as MFA client.

- You will add the MFA provider as an external authentication domain in Nexus Dashboard, as described in [Adding Okta as Remote Authentication Provider](#).

Configuring Okta Account as MFA Provider

The following steps provide basic configuration required to enable MFA for Nexus Dashboard using Okta as a provider. Detailed Okta configurations are outside the scope of this document, see Okta documentation for all available options.

To configure Okta for Nexus Dashboard MFA:

1. Log in to your Okta account.

To create an account, browse to <https://developer.okta.com>.

2. Create a new app integration.
 - a. From the left navigation menu, select **Applications > Applications**.
 - b. Click **Create App Integration**.
 - c. For **Sign-in method**, select **OIDC - OpenID Connect**.
 - d. For **Application Type**, select **Web Application**.
 - e. Click **Next**.
 - f. Provide **App integration name**, for example, **nd-mfa**.

The following steps assume you used **nd-mfa** as the app integration name. If you choose a different name, replace **nd-mfa** where appropriate.

- g. For **Sign-in redirect URIs**, enter <https://<nd-node1-ip>/oidccallback>

Then click **+Add URI** to provide the URIs for all nodes in the cluster.

- h. For **Controlled Access**, choose **Skip group assignment for now**.

- i. Leave other fields at their default values and click **Save**.

3. Add the required attributes to the default user.

- a. From the left navigation menu, select **Directory > Profile Editor**.
 - b. Click the **Okta User (default)** profile.
 - c. Click **+Add Attribute**.
 - d. For **Data type**, choose **string**.
 - e. For **Display name**, **Variable name**, and **Description**, enter **CiscoAVPair**.
 - f. Ensure that **Attribute required** is **unchecked**.
 - g. Leave other fields at default values and click **Save and Add Another**.
 - h. For **Data type**, choose **string**.
 - i. For **Display name**, **Variable name**, and **Description**, enter **nduser**.
 - j. Ensure that **Attribute required** is **unchecked**.
 - k. Leave other fields at default values and click **Save**.
4. Add the required attributes to the **nd-mfa** user you created.
- a. From the left navigation menu, select **Directory > Profile Editor**.
 - b. Click the **nd-mfa User (default)** profile.
 - c. Click **+Add Attribute**.
 - d. For **Data type**, choose **string**.
 - e. For **Display name**, **Variable name**, and **Description**, enter **CiscoAVPair**.
 - f. Ensure that **Attribute required** is **checked**.
 - g. Leave other fields at default values and click **Save and Add Another**.
 - h. For **Data type**, choose **string**.
 - i. For **Display name**, **Variable name**, and **Description**, enter **nduser**.
 - j. Ensure that **Attribute required** is **checked**.
 - k. Leave other fields at default values and click **Save**.
5. Map the attributes.
- a. From the left navigation menu, select **Directory > Profile Editor**.
 - b. Click the **nd-mfa User** profile.
 - c. In the **Attributes** area of the main window, click **Mappings**.

The **nd-mfa User Profile Mappings** window opens.

nd-mfa User Profile Mappings

×

nd-mfa User Profile Mappings

nd-mfa to Okta User Okta User to nd-mfa

nd-mfa User Profile appuser

Okta User User Profile user

appuser.userName login string

Use default username setting for Okta user

appuser.given_name firstName string

appuser.family_name lastName string

Choose an attribute or enter an expression... middleName string

appuser.CiscoAVPair CiscoAVPair string

appuser.nduser nduser string

Preview Enter an Okta user to preview their mappi Save Mappings Cancel

- d. At the top of the **nd-mfa User Profile Mappings** window, click **nd-mfa to Okta User**.
 - e. Select **app.CiscoAVPair** from the dropdown menu next to **CiscoAVPair**.
 - f. Select **app.nduser** from the dropdown menu next to **nduser**.
 - g. Click **Save Mappings**.
 - h. Click **Apply updates now**.
6. Create users.
- a. From the left navigation menu, select **Directory > People**.
 - b. Click **+Add person**.
 - c. Provide the user information.
 - d. Click **Save and Add Another** to add another user or click **Save** to finish.

You must add all users that you want to be able to log in to your Nexus Dashboard.

7. Assign users to the app.
- a. From the left navigation menu, select **Applications > Application**.
 - b. Click the application you created (**nd-mfa**).
 - c. Select the **Assignments** tab.
 - d. Choose **Assign > Assign to People**.

The **Assign nd-mfa to People** window opens.

- e. In the **Assign nd-mfa to People** window, click **Assign** next to the user you want to be able to log in to your Nexus Dashboard.
- f. In the user details window that opens, provide a value for **CiscoAVPair** and **nduser** fields.

The **CiscoAVPair** values are described in the [Configuring Remote Authentication Server](#), for example `shell:domains=all/admin/`.

The **nduser** value will be used as the username for this user when logging in to your Nexus Dashboard.

- g. Click **Save and Go Back**.
- h. Assign another user or click **Done** to finish.

You must add all users that you created in a previous step.

8. Configure **Claims** for the app.

- a. From the left navigation menu, select **Security > API**.
- b. Click the **default** name.
- c. Select the **Claims** tab.
- d. Click **+Add Claim** to add the **CiscoAVPair** claim.
- e. In the **Name** field, enter **CiscoAVPair**.
- f. From the **Include in token type** dropdown, select **ID Token**.

We recommend using **ID Token**, however **Access Token** is also supported.

- g. In the **Value** field, enter `appuser.CiscoAVPair`.
- h. Click **Save**.
- i. Click **+Add Claim** to add the **nduser** claim.
- j. In the **Name** field, enter **nduser**.
- k. From the **Include in token type** dropdown, select **ID Token**.

You must create both claims in the same token, mixing **ID Token** and **Access Token** is not supported.

- l. In the **Value** field, enter `appuser.nduser`.
- m. Click **Save**.

9. Gather the required Okta account information for adding it as authentication provider for your Nexus Dashboard.

- a. From the left navigation menu, select **Security > API**.
- b. Click the **default** name.
- c. Note down the **Issuer** value.

Settings Edit

Name	default
Audience	api://default
Description	Default Authorization Server for your Applications
Issuer	https://dev- .okta.com/oauth2/default
Metadata URI	https://dev- .okta.com/oauth2/default/well-known/oauth-authorization-server
Signing Key Rotation ?	Automatic
Last Rotation	15 Nov 2021

- d. From the left navigation menu, select **Application > Applications**.
- e. Click the application you created (**nd-mfa**).
- f. Note down the **Client ID** and **Client Secret** values.

Client Credentials Edit

Client ID	<input type="text" value=""/> Copy
	Public identifier for the client that is required for all OAuth flows.
Client secret	<input type="password" value="....."/> Copy
	Secret used by the client to exchange an authorization code for a token. This must be kept confidential! Do not include it in apps which cannot keep it secret, such as those running on a client.

Configuring MFA Client

This release supports only Cisco Duo as MFA client.

The following steps provide basic configuration required to enable using Cisco Duo for Nexus Dashboard MFA. Detailed Duo configurations are outside the scope of this document, see Cisco Duo documentation for all available options.

To configure Duo:

1. Log in to your Okta account.
2. Add DUO as an MFA type.
 - a. From the left navigation menu, select **Security > Multifactor**.
 - b. In the **Factor Types** tab, select **Duo Security**.

If you do not have the **Duo Security** option, you will need to open a support case with Okta from <https://support.okta.com/help/s/opencase>.

- c. In the **Duo Security** window, provide the required information.

For more information on how to obtain integration key, secret key, and API hostname, see <https://duo.com/docs/okta>.

Ensure that **Duo Username Format** is set to **Email**.

- d. Click **Save**.

3. Create a Duo rule.

- a. From the left navigation menu, select **Applications > Application**.
- b. Click the application you created (**nd-mfa**).
- c. Select the **Sign On** tab.
- d. In the **Sign On Policy** area, click **+Add Rule**.
- e. Provide the name for the rule.
- f. In the **Access** area, enable **Prompt for factor** and select **Every sign on**.
- g. Specify other options as required by your use case.
- h. Click **Save**.

4. Configure Okta and Duo integration.

There are two ways you can allow the users you configured in Okta to use the Duo app for MFA—have the Duo admin add all the same users in Duo dashboard or have each individual user log in to Okta and self-enroll.

To configure users in Duo dashboard:

- a. Log in to your Duo dashboard as admin user.
- b. From the left navigation menu, select **Users**.
- c. Click **Add User** and provide the details that match the user's information in Okta.
- d. Repeat this step for all users you added in Okta.

To self-enroll:

- a. Instruct every user you created in [Configuring Okta Account as MFA Provider](#) to log in to Okta on their own using your specific Okta domain.

You can determine the Okta domain to use by navigating to **Application > Application**, then clicking the **nd-mfa** application you created and copying the **Okta domain** URL:

← Back to Applications

nd-mfa

Active ▾ View Logs

General Sign On Assignments Okta API Scopes

General Settings Edit

Okta domain .okta.com

APPLICATION

- b. Once they're logged in, they can navigate to the **Settings** page from the top right user menu.
- c. Choose **Duo Security Setup** and follow the instructions on the screen.

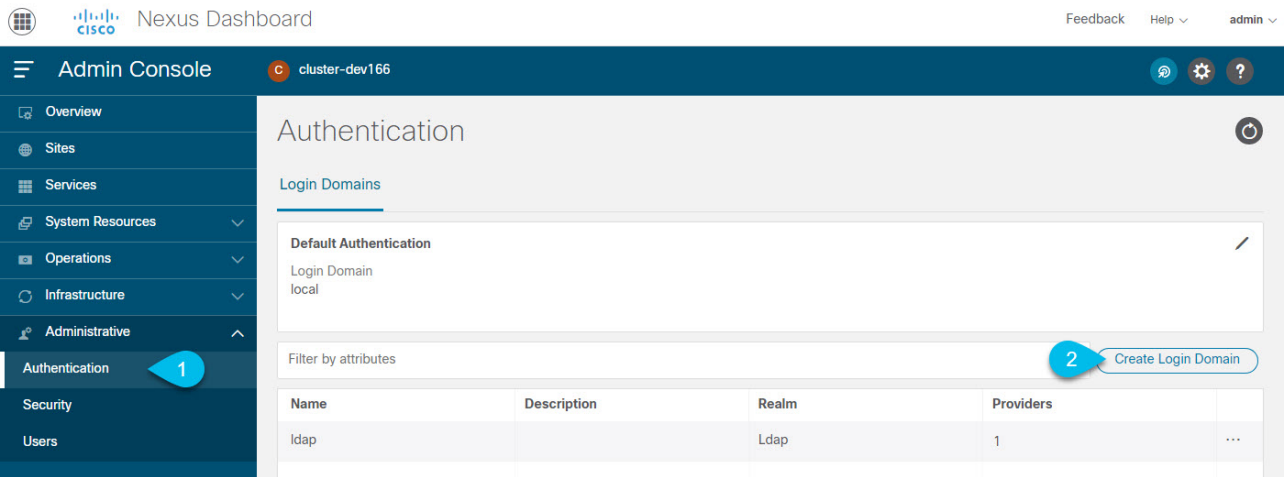
Adding Okta as Remote Authentication Provider

Before you begin

- You must have at least one user already configured in Okta as described in [Configuring Okta Account as MFA Provider](#).
- You must have the **Client ID**, **Client Secret**, and **Issuer** information from your Okta account available, which is described in the last step of [Configuring Okta Account as MFA Provider](#).
- If you want to use a proxy to connect to your Okta account, the proxy must already be configured as described in [Cluster Configuration](#).

To add Okta as remote authentication provider:

1. Log in to your Nexus Dashboard as an **admin** user.
2. Navigate to the **Admin Console**.
3. Add an authentication domain.



- a. From the main navigation menu, select **Administrative > Authentication**.
 - b. In top right of the main pane, click the **Actions** menu and select **Create Login Domain**.
4. In the **Create Login Domain** screen that opens, provide domain details.
- a. Provide the **Name** for the domain.
 - b. (Optional) Provide its **Description**.
 - c. From the **Realm** dropdown, select **OIDC**.
 - d. In the **Client ID** field, enter the client ID you obtained from your Okta account.
 - e. In the **Client Secret** field, enter the client secret you obtained from your Okta account.
 - f. In the **Issuer** field, enter the URI you obtained from your Okta account.
 - g. (Optional) Check the **User Proxy** option if you want to connect to Okta over a proxy.
 - h. Leave the **Scopes** options unchecked.

This release supports the **openid** scope only.

5. Click **Create** to finish adding the domain.

Logging In To Nexus Dashboard Using MFA

1. Navigate to one of your Nexus Dashboard IPs as your typically would.
2. From the **Login Domain** dropdown, select the OIDC domain you created in [Adding Okta as Remote Authentication Provider](#).

The **Username** and **Password** fields will not be displayed.

3. Click **Login**.

You will be redirected to the Okta login page.

4. Log in using a user that was configured in Okta as described in [Configuring Okta Account as MFA Provider](#).

A push notification will be sent to your Duo client.

5. Approve the login using Duo.

You will be redirected back to the Nexus Dashboard UI and logged in using the Okta user.

Users

The **Users** GUI page allows you to view and manage all users that have access to the Nexus Dashboard.

The **Local** tab displays all local users while the **Remote** tab displays users that are configured on the remote authentication servers you have added as described in the [Remote Authentication](#) section.

Note that Single sign-on (SSO) between the Nexus Dashboard, sites, and applications is available for remote users only. For more information on configuring remote users, see [Remote Authentication](#).

Adding Local Users

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. Create a new local user.
 - a. From the main navigation menu, select **Administrative > Users**.
 - b. In top right of the main pane, click **Create Local User**.
3. In the **Create Local User** screen that opens, provide user details.
 - a. Provide the **User ID** that will be used for login in.
 - b. Provide and confirm the initial **Password**.
 - c. Provide the **First Name**, **Last Name**, and **Email Address** for the user.
 - d. Choose the user's **Roles** and **Privileges**.

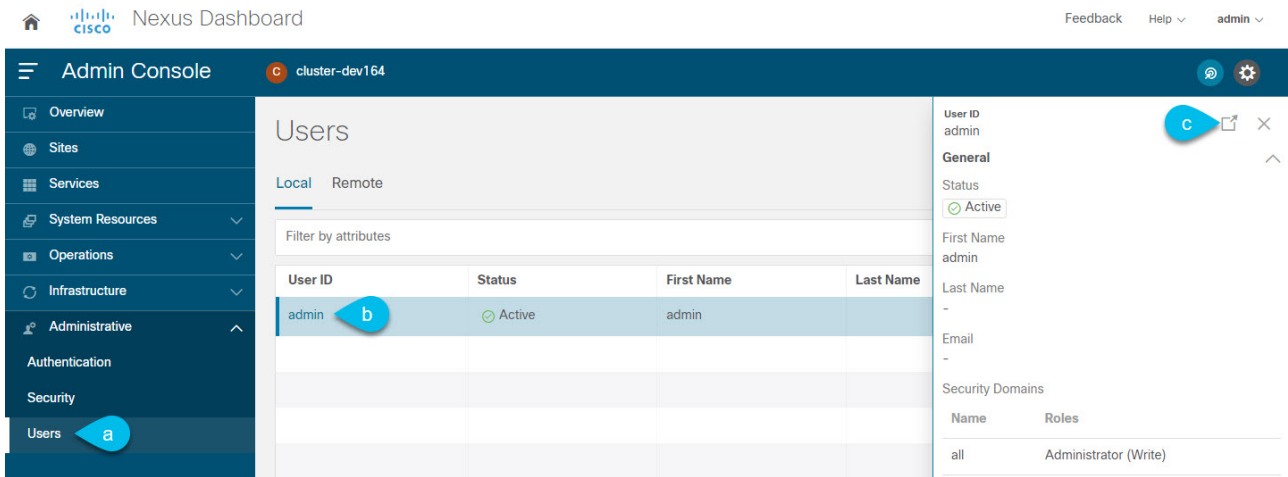
You can select one or more roles for each user. The available roles and their permissions are described in [Roles and Permissions](#).

For all of the user roles you select, you can choose to enable read-only or read-write access. In case of read-only access, the user will be able to view the objects and settings allowed by their user **Role** but unable to make any changes to them.

- e. Click **Create** to save the user.

Editing Local Users

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. Open user details screen.



- a. From the main navigation menu, select **Administrative > Users**.
 - b. In the main pane, click on the user's name.
 - c. In the details pane that opens, click the **Details** icon.
3. In the **<user-name>** details screen that opens, click the **Edit** icon.
 4. In the **Edit User** screen that opens, update the settings as necessary.

Security

The **Security** GUI page allows you to view and manage certificates used by the Nexus Dashboard.

Security Configuration

The **Administrative > Security Configuration** page allows you to configure authentication session timeouts and security certificates used by your Nexus Dashboard cluster.

Before you begin

- You must have the keys and certificates you plan to use with Nexus Dashboard already generated.

Typically, this includes the following files:

- Private key (**nd.key**)
- Certificate Authority's (CA) public certificate (**ca.crt**)
- CA-signed certificate (**nd.crt**)

Generating these files for self-signed certificates is described in [Generating Private Key and Self-Signed Certificate](#).

- We recommend creating a configuration backup of your Nexus Dashboard cluster before making changes to the security configurations.

For more information about backups, see [Backup and Restore](#).

To edit security configuration:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. Edit security configuration.

- a. From the main navigation menu, select **Administrative > Security**.
 - b. In the main pane, select the **Security Configuration** tab.
 - c. In top right of the main pane, click the **Edit** icon.
3. In the **Security Configuration** screen that opens, update one or more fields as required:

Note that uploading the keys and certificate files is not supported and you will need to paste the information in the following fields.

- a. Update the **Session Timeout**.

This field defines the duration of the API tokens with the default duration set to 20 minutes.

- b. Update the **Idle Timeout**.

This field defines the duration of the UI session.

- c. In the **Domain Name** field, provide your domain.
- d. Click the **SSL Ciphers** field and select any additional cipher suites you want to enable from the dropdown or click the **x** icon on an existing cipher suite to remove it.

Cipher suites define algorithms (such as key exchange, bulk encryption, and message authentication code) used to secure a network connection. This field allows you to customize which cipher suites your Nexus Dashboard cluster will use for network communication and disable any undesired suites, such as the less secure TLS1.2 and TLS1.3.

- e. In the **Key** field, provide your private key.
- f. In the **RSA Certificate** field, provide the CA-signed or self-signed certificate.
- g. In the **Root Certificate** field, provide the CA's public certificate.
- h. (Optional) If your CA provided an Intermediate Certificate, provide it in the **Intermediate Certificate** field.
- i. Click **Save** to save the changes.

After you save your changes, the GUI will reload using the new settings.

Security Domains

A restricted security domain allows an administrator to prevent a group of users from viewing or modifying any objects created by a group of users in a different security domain, even when users in both groups have the same assigned privileges.

For example, an administrator in restricted security domain (**domain1**) will not be able to see sites, services, cluster or user configurations in another security domain (**domain2**).

Note that a user will always have read-only visibility to system-created configurations for which the user has proper privileges. A user in a restricted security domain can be given a broad level of privileges within that domain without the concern that the user could inadvertently affect another group's physical environment.

To create a security domain:

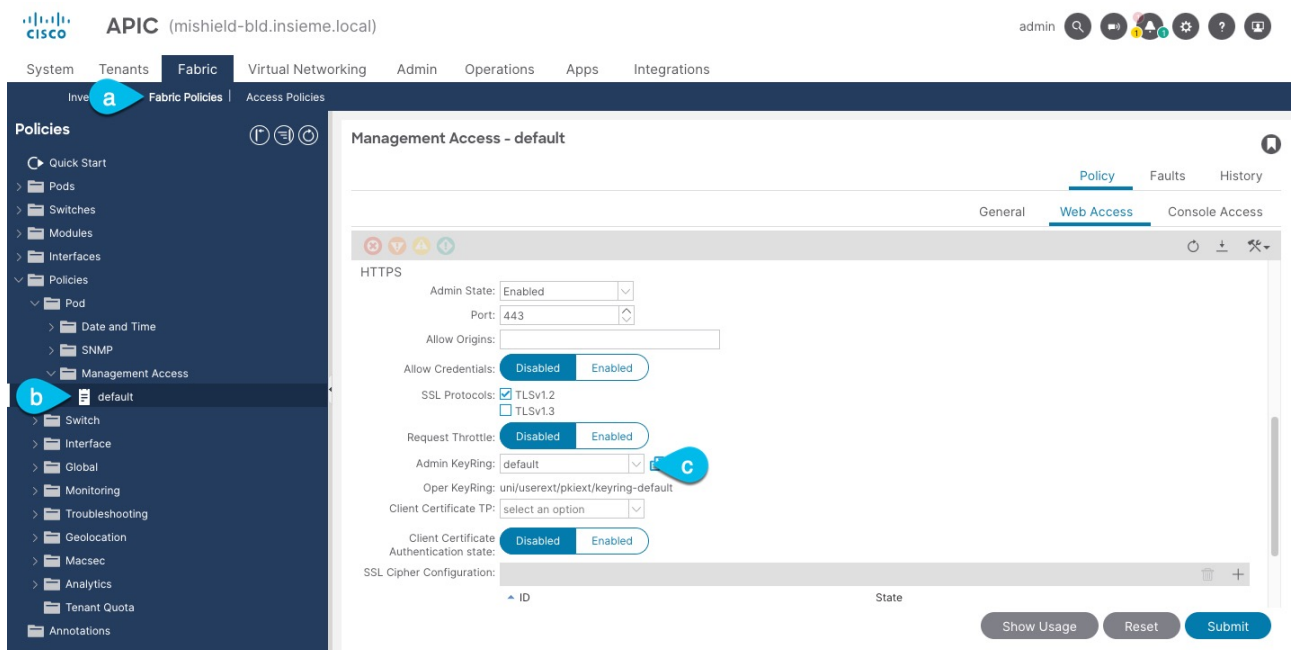
1. Navigate to your Nexus Dashboard's **Admin Console**.
2. Create a new security domain.
 - a. From the main navigation menu, select **Administrative > Security**.
 - b. In the main pane, select the **Security Domains** tab.
 - c. In top right of the main pane, click **Create Security Domain**.
3. In the **Create Security Domain** screen that opens, provide the domain details.
 - a. Provide the **Name** for the domain.
 - b. (Optional) Provide a description for the domain.
 - c. Click **Create** to save the domain.

Validating Peer Certificates

Beginning with release 2.3.1, you can import a site controller's Certificate Authority (CA) root certificate chain into Nexus Dashboard. This allows you to verify that the certificates of hosts to which your Nexus Dashboard connects (such as site controllers) are valid and are signed by a trusted Certificate Authority (CA) when you add the sites.

Exporting Certificate Chain From Cisco APIC

1. Log in to your Cisco APIC.
2. Check which key ring is being used for management access:

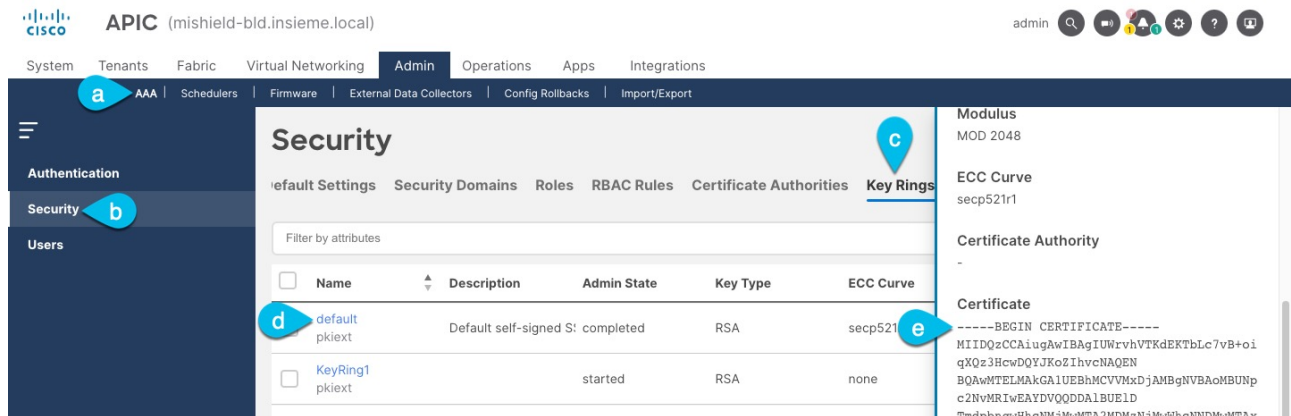


- a. In the top navigation bar, choose **Fabric > Fabric Policies**.
- b. In the left navigation menu, choose **Policies > Pod > Management Access**.
- c. In the main pane, note the name in the **Admin KeyRing** field.

In the above example, the **default** key ring is being used. However, if you created a custom key ring with a custom certificate chain, the name of that key ring would be listed in the **Admin KeyRing** field.

Custom security configuration for Cisco APIC is described in detail in [Cisco APIC Security Configuration Guide](#) for your release.

3. Export the certificate used by the key ring:



- In the top navigation bar, choose **Admin > AAA**.
- In the left navigation menu, choose **Security**.
- In the main pane, choose the **Key Rings** tab.
- Click the name of the key ring you found in the previous step and copy the **Certificate**.

The above example shows the **default** key ring from the previous step. However, if you had a custom key ring configured, choose the CA certificate chain used to create the key ring.

You must include the **-----BEGIN CERTIFICATE-----** and **-----END CERTIFICATE-----** in the text you copy, for example:

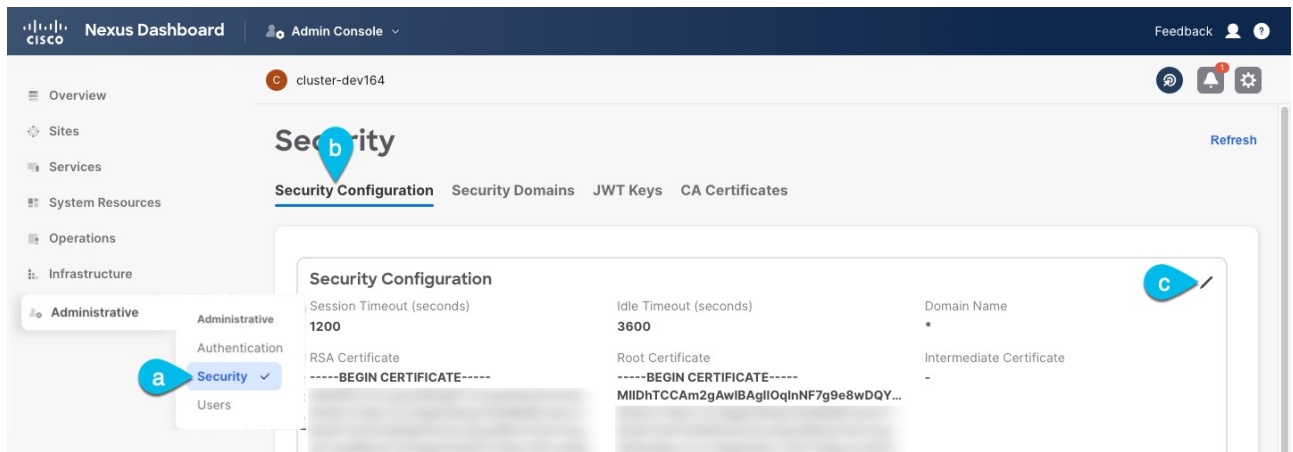
```
-----BEGIN CERTIFICATE-----  
MIIDQzCCAIugAwIBAgIUWrvhVTKdEKTbLc7vB+oiqXQz3HcwDQYJKoZIhvcNAQEN  
[...]  
-----END CERTIFICATE-----
```

Exporting Certificate Chain From Cisco NDFC

- Log in to the Nexus Dashboard that's hosting the service.

In case of NDFC, there is no separate certificate from the service so you use the Nexus Dashboard host's certificate.

- Export the certificate.



- In the main navigation menu, choose **Administrative > Security**.
- In the main pane, choose the **Security Configuration** tab
- In the **Security Configuration** page, click the **Edit** icon.
- Copy the **Root Certificate**.



We recommend copying the certificate chain from the **Edit** page instead of directly from the **Security Configuration** page to ensure that there are no spaces or new line characters (`\n`) in the string you copy.

You must include the `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` in the text you copy, for example:

```
-----BEGIN CERTIFICATE-----
MIIDhTCCAm2gAwIBAgIIQlnNF7g9e8wDQYJKoZIhvcNAQELBQAwSzELMAkGA1UE
[...]
-----END CERTIFICATE-----
```

Exporting Certificate Chain From Cisco DCNM

- SSH in to your Cisco DCNM as the `sysadmin` user.

Unlike the other site controllers, DCNM certificates are not available in the UI, so you must use the CLI.

```
# ssh -l sysadmin <dcnm-ip-address>
```

For additional information about certificate management in DCNM, see the "[Certificate Management](#)" chapter of the *Cisco DCNM Installation and Upgrade Guide for LAN Fabric Deployment*.

- Change into the `/var/lib/dcnm/afw/apigateway/` directory.

The certificate (`dcnmweb.crt`) file is located in this directory.

```

dcnm# cd /var/lib/dcnm/afw/apigateway/
dcnm# ls -ltr /* View the contents of the folder
total 128
-rw----- 1 root root 1844 Nov 18 13:14 dcnmweb.key.2019-11-20T132939-08:00
-rw-r--r-- 1 root root 1532 Nov 18 13:14 dcnmweb.crt.2019-11-20T132939-08:00
-rw----- 1 root root 1844 Nov 20 10:15 dcnmweb.key.2019-11-20T132950-08:00
-rw-r--r-- 1 root root 1532 Nov 20 10:15 dcnmweb.crt.2019-11-20T132950-08:00
-rw----- 1 root root 1844 Dec 22 13:59 dcnmweb.key
-rw-r--r-- 1 root root 1532 Dec 22 13:59 dcnmweb.crt

```

3. Check for root certificate.

Depending on which Certificate Authority you used to sign your certificate, the root certificate may be included in the `dcnmweb.crt` file or may be provided as a separate file.

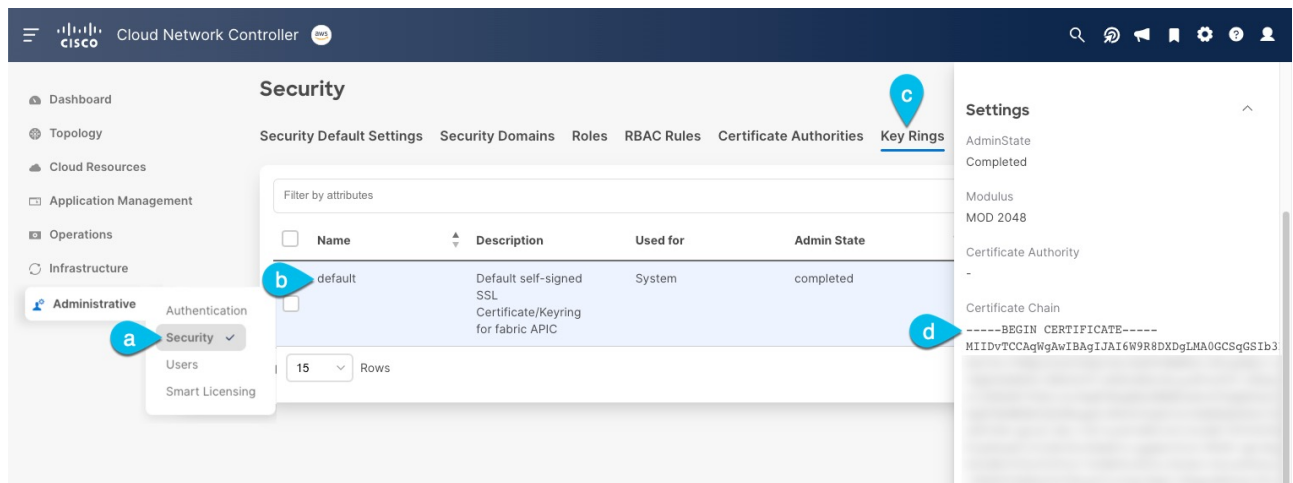
To check whether the root certificate is included:

```
dcnm# openssl x509 -text -noout -in dcnmweb.crt
```

If the file includes the root certificate, copy it. Otherwise, use the root certificate file you should have obtained when signing your certificate.

Exporting Certificate Chain From Cisco Cloud Network Controller

1. Log in to your Cisco Cloud Network Controller.
2. Export the certificate.



- a. In the main navigation menu, choose **Administrative > Security**.
- b. In the main pane, choose the **Key Rings** tab.
- c. Click the name of the certificate you want to import into your Nexus Dashboard and copy the **Certificate Authorities**.

The above example shows the `default` key ring. However, if you had a custom key ring configured, choose the CA certificate chain used to create the key ring.

You must include the `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` in the

text you copy, for example:

```
-----BEGIN CERTIFICATE-----  
MIIDvTCCAqWgAwIBAgIJAI6W9R8DXDgLMA0GCSqGSIb3DQEBDQUAMEAxGzAJBg  
NV  
[...]  
-----END CERTIFICATE-----
```

Importing Certificates Into Nexus Dashboard

1. Log in to your Nexus Dashboard where you plan to onboard the sites.
2. Import the certificate into Nexus Dashboard.
 - a. Log in to your Nexus Dashboard where you will onboard the sites.
 - b. In the left navigation menu, choose **Administrative > Security**.
 - c. In the main pane, select the **CA Certificates** tab.
 - d. Click **Add CA certificate**, provide a unique name for the certificate, and paste the certificate chain you copied from your site's controller.
3. Proceed with adding the site as you typically would, but enable the **Verify Peer Certificate** option.

Note that if you enable the **Verify Peer Certificate** option but don't import the valid certificate, site onboarding will fail.

Adding sites is described in [Adding Sites](#).

Cisco Intersight

Cisco Intersight is a Software-as-a-Service (SaaS) infrastructure management platform that is augmented by other intelligent systems. It provides global management of the Cisco Unified Computing System (Cisco UCS) and Cisco HyperFlex hyperconverged infrastructure, Cisco APIC, and other platforms including Nexus Dashboard.

Data center apps, such as Cisco Nexus Dashboard Insights, connect to the Cisco Intersight portal through a Device Connector that is embedded in the management controller of each system, in this case your Nexus Dashboard platform. Device Connector provides a secure way for the connected devices to send information and receive control instructions from the Cisco Intersight portal, using a secure internet connection.

When an Intersight-enabled device or application starts, the Device Connector starts at boot by default and attempts to connect to the cloud service. If the **Auto Update** option is enabled, the Device Connector is automatically updated to the latest version through a refresh by the Intersight service when you connect to Cisco Intersight. For more information on the **Auto Update** option, see [Configuring Device Connector Settings](#).

For additional information on Cisco Intersight, see https://www.intersight.com/help/saas/getting_started/overview.



If you upgraded from Application Services Engine and your Intersight device connector is claimed with a proxy configured, you will need to re-configure the proxy in the **Cluster Configuration** screen. For more information, see [Cluster Configuration](#).

Configuring Device Connector Settings

Devices are connected to the Cisco Intersight portal through a Device Connector, which provides a secure way for the connected devices to send information and receive control instructions from the Cisco Intersight portal.

All device connectors must properly resolve svc.intersight.com and allow outbound initiated HTTPS connections on port **443**. If a proxy is required for an HTTPS connection, you must configure the proxy settings in your Nexus Dashboard.

This section describes how to configure the basic Device Connector settings.

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the main navigation menu, select **Infrastructure > Intersight**.
3. In the top right of the main pane, click **Settings**.
4. Click the **General** tab to configure basic options.
 - a. Use the **Device Connector** knob to enable or disable the Device Connector.

This enables you to claim the device and leverage the capabilities of Intersight. If it is disabled, no communication is allowed to Cisco Intersight.

- b. In the **Access Mode** area, choose whether to allow Intersight the capability to make changes

to this device.

- **Allow Control** (default) – enables you to perform full read or write operations from the cloud based on the features available in Cisco Intersight.
- **Read-only** – ensures that no changes are made to this device from Cisco Intersight.

For example, actions such as upgrading firmware or a profile deployment will not be allowed in read-only mode. However, the actions depend on the features available for a particular system.

- c. Use the **Auto Update** knob to enable automatic Device Connector updates.

We recommend that you enable automatic updates so that the system automatically updates the Device Connector software. When enabled, the Device Connector will automatically upgrade its image whenever there is any upgrade push from Intersight.

If you disable the automatic updates, you will be asked to manually update the software when new releases become available. Note that if the Device Connector is out-of-date, it may be unable to connect to Cisco Intersight.

5. Click **Save** to save the changes.

6. Click the **Certificate Manager** tab if you want to import additional certificates.

By default, the device connector trusts only the built-in certificate. If the device connector establishes a TLS connection and a server sends a certificate that does not match the built-in certificate, the device connector terminates TLS connections because it cannot determine if the server is a trusted device.

You can choose to upload additional certificates by clicking the **Import** button in this screen. The imported certificates must be in the **.pem** (base64 encoded) format. After a certificate is successfully imported, it is listed in the list of **Trusted Certificates** and if the certificate is correct, it is shown in the **In-Use** column.

You can click the **View** icon at the end of the certificate's row to view its details such as name, issue and expiration dates.

Target Claim

This section describes how to claim the Nexus Dashboard platform as a device for Cisco Intersight.

Before you begin

You must have configured the Intersight Device Connector as described in [Configuring Device Connector Settings](#).

To claim the device:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the main navigation menu, select **Infrastructure > Intersight**.
3. Check whether the Device Connector is already configured.
 - o If you see a green dotted line connecting **Internet** to **Intersight** in the **Device Connector** page

and the text **Claimed**, then your Intersight Device Connector is already configured and connected to the Intersight cloud service, and the device is claimed. In this case, you can skip the rest of this section.

- o If you see a red dotted line connecting to **Internet** in the **Device Connector** page, you must configure a proxy for your Nexus Dashboard cluster to be able to access the Internet, as described in [Cluster Configuration](#) before continuing with the rest of this section.
- o If you see a yellow dotted line and a caution icon connecting **Internet** to **Intersight** in the **Device Connector** page and the text **Not Claimed**, then your Intersight Device Connector is not yet configured and connected to the Intersight service, and the device is not yet claimed. Follow these procedures to configure the Intersight Device Connector and connect to the Intersight cloud service, and claim the device. In this case, proceed with the rest of the steps to configure the device.

4. If necessary, update the device connector software.

If there is a new Device Connector software version available and you do not have the **Auto Update** option enabled, you will see a message at the top of the screen informing you that Device Connector has important updates available. Enabling the auto-update feature is described in [Configuring Device Connector Settings](#).

To manually update the Device Connector, click the **Update Now** link.

5. Note down the **Device ID** and **Claim Code** listed on the Nexus Dashboard's **Intersight** page.

The screenshot displays the 'Intersight Device Connector' page in the Nexus Dashboard. On the left is a navigation menu with options like Dashboard, System Overview, Sites, Service Catalog, System Resources, Operations, Infrastructure, Cluster Configuration, Resource Utilization, Intersight, App Infra Services, and Administrative. The main content area shows the 'Device Connector' status as 'Not Claimed' with a yellow warning banner. A blue box highlights the 'Device ID' and 'Claim Code' fields. The Device ID is 'WZP22470ZM2&WZP23150D47&WZP23150D4W&WZP23310KD6&WZP242017UP' and the Claim Code is 'FDE7B01E1FC4'. A diagram below shows the connection flow: Device Connector (computer icon) -> Internet (globe icon) -> Intersight (cloud icon). A 'Not Claimed' message states: 'The connection to the Cisco Intersight Portal is successful, but device is still not claimed. To claim the device open Cisco Intersight, create a new account and follow the guidance or go to the Devices page and click Claim a New Device for existing account. Open Intersight'. The version number '1.0.9-731' is visible at the bottom left of the main content area.

6. Log into the Cisco Intersight cloud site at <https://www.intersight.com>.

7. Follow the instructions described in the **Target Claim** section of the Intersight documentation to claim the device.

After the device is claimed in Intersight, you should see green dotted lines connecting **Internet** to **Intersight** in your Nexus Dashboard's **Device Connector** page along with the text **Claimed**.



You may need to click **Refresh** in top right of the page to update the latest status.

Unclaiming the Device

To unclaim the Nexus Dashboard as a device from Intersight:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the main navigation menu, select **Infrastructure > Intersight**.
3. In the main pane, click **Unclaim**.

Troubleshooting

Useful Commands

You can log in to any of the cluster nodes as **rescue-user** for a limited access to system data. You can use the following commands to perform various operations in Cisco Nexus Dashboard.

Cluster Troubleshooting:

- **acs health** – displays cluster health information and any existing issues.
- **acs show cluster** – displays cluster configuration.
- **acs show nodes** – displays information about all nodes in the cluster.
- **acs show masters** – displays information about **master** nodes in the cluster.
- **acs show workers** – displays information about **worker** nodes in the cluster.
- **acs show standbys** – displays information about **standby** nodes in the cluster.
- **acs ntp show** – displays NTP information.
- **acs techsupport collect -s system** – collects Infra tech support information.
- **acs techsupport collect -s cisco-mso** – collects Nexus Dashboard Orchestrator service tech support information.
- **acs techsupport collect -s cisco-nir** – collects Nexus Dashboard Insights service tech support information.
- **acs techsupport collect -s cisco-appcenter** – collects App Store tech support information.
- **acs version** – returns the Nexus Dashboard version.

Resetting Devices:

- **acs reboot** – reboots the node with all services and configurations intact.
- **acs reboot clean** – removes all data for Nexus Dashboard and applications, but preserves the Nexus Dashboard bootstrap configuration and pod images.

When you first bring up your Nexus Dashboard cluster, initial deployment process installs all required pod images. Retaining pod images will speed up cluster bring up after reboot.

If you plan to re-install all the nodes in the cluster, you must clean up the site and app information first. In this case, ensure that the sites are disabled in all applications and removed from the ND cluster.

- **acs reboot clean-wipe** – removes all data for Nexus Dashboard and applications including application images, but preserves the Nexus Dashboard bootstrap configuration.

When the cluster boots up again, pod images will be re-installed.

If you plan to re-install all the nodes in the cluster, you must clean up the site and app information first. In this case, ensure that the sites are disabled in all applications and removed from the ND

cluster.

- **acs reboot factory-reset**—removes all data for Nexus Dashboard and applications including cluster bootstrap configuration, but preserves application images.

When you first bring up your Nexus Dashboard cluster, initial deployment process installs all required pod images. Retaining pod images will speed up cluster bring up.

If you plan to re-install all the nodes in the cluster, you must clean up the site and app information first. In this case, ensure that the sites are disabled in all applications and removed from the ND cluster.

- **acs reboot factory-wipe**—removes all data for Nexus Dashboard and applications, including application images and cluster bootstrap configuration.

When the cluster boots up again, the pod images will be re-installed.

If you plan to re-install all the nodes in the cluster, you must clean up the site and app information first. In this case, ensure that the sites are disabled in all applications and removed from the ND cluster.

System and Connectivity Troubleshooting:

- The **/logs** directory is mounted into the **rescue-user** container and can be inspected with standard tools.
- **ping** command is supported with most options.
- **ip** command supports a read-only subset of commands, including **ip addr show** and **ip route show**.
- **kubectl** command supports read-only Kubernetes commands.

For example, you can use it to get a list of all pods running in the system:

```
$ kubectl get pods -A
NAMESPACE      NAME                                READY STATUS RESTARTS  AGE
aaamgr          aaamgr-54494fdb8-q8rc4             2/2   Running 0         3d3h
authy-oidc      authy-oidc-75fdf44b57-x48xr        1/1   Running 3 (3d3h ago) 3d4h
authy           authy-857fbb7fdc-7cwgg             3/3   Running 0         3d4h
cisco-appcenter apiserver-686655896d-kmqhq         1/1   Running 0         3d3h
[...]
```

- **acs elasticsearch** command invokes a custom utility that allows you to get debug information about the services.

```
$ acs elasticsearch --name cisco-ndfc-controller-elasticsearch health
{
  "cluster_name" : "cisco-ndfc-controller-elasticsearch",
  "status" : "green",
  "timed_out" : false,
```

```

"number_of_nodes" : 3,
"number_of_data_nodes" : 3,
"discovered_master" : true,
"active_primary_shards" : 10,
"active_shards" : 21,
"relocating_shards" : 0,
"initializing_shards" : 0,
"unassigned_shards" : 0,
"delayed_unassigned_shards" : 0,
"number_of_pending_tasks" : 0,
"number_of_in_flight_fetch" : 0,
"task_max_waiting_in_queue_millis" : 0,
"active_shards_percent_as_number" : 100.0
}

```

You can obtain the list of the service-specific pod names using the `kubectl` command, for example:

```

$ kubectl get pods -A | grep elasticsearch
cisco-ndfc-controller-elasticsearch es-data-0 2/2 Running 0 109m
cisco-ndfc-controller-elasticsearch es-data-1 2/2 Running 0 163m
cisco-ndfc-controller-elasticsearch es-data-2 2/2 Running 0 104m

```

Application Information:

- `acs apps instances` command displays all applications running on the cluster.
- `acs apps actions` command displays the history operations done on the applications, such as installations, upgrades, or deletions.

Upgrading CIMC

When you upgrade Nexus Dashboard software, you may also have to upgrade the version of Cisco Integrated Management Controller (CIMC) that is running in your Nexus Dashboard nodes.

Supported CIMC versions for each Nexus Dashboard release are listed in the [Release Notes](#) specific to that release.

The following steps describe how to upgrade the Nexus Dashboard CIMC using the Cisco Host Upgrade Utility (HUU). Additional details about the Host Upgrade Utility are available at [Upgrading the Firmware on a Cisco UCS C-Series Server Using the HUU](#).

Before You Begin

- Check the [Release Notes](#) for your Nexus Dashboard release to confirm the CIMC versions supported by that release.
- Allow for the appropriate amount of time for the upgrade.

The time required for the upgrade process depends on a number of factors, such as the speed of

the link between the local machine and the UCS-C chassis, the source and target software images, as well as other internal component versions.

- If you're upgrading a single node that is running an older firmware to add it to an existing cluster, you will perform the following steps on that node only and not on all nodes in the cluster.
- Updating CIMC may also require updating your browser and/or Java software version to run the vKVM used to upgrade the CIMC.



Upgrading the CIMC version does not affect your production network as the Nexus Dashboard nodes are not in the data path of the traffic.

To upgrade the Nexus Dashboard CIMC software:

1. Open your browser, navigate to the CIMC IP address, and log in using the CIMC credentials.

Note that the CIMC credentials may be different from the Nexus Dashboard GUI credentials.

2. Determine the model of UCS platform for your Nexus Dashboard by locating the first part of the BIOS version under **Server > Summary**.

Nexus Dashboard supports the UCS-C220-M5 and UCS-C225-M6 servers.

Cisco Integrated Management Controller (Cisco IMC) Information

Server Properties	Cisco IMC Information
Product Name: SE-NODE-G2	Hostname: C220-WMP250600S0
Serial Number: WMP250600S0	IP Address: 172.28.185.116
PID: SE-NODE-G2	MAC Address: 48:8B:0A:45:EC:D0
UUID: 09A2D89E-A6C0-4F6D-9C91-2665E18FF8DC	Firmware Version: 4.1(2a)
BIOS Version: C220M5.4.1.2a.0.0624200115	Current Time (UTC): Tue Mar 21 21:07:09 2023
Description: <input type="text"/>	Local Time: Tue Mar 21 21:07:09 2023 UTC +0000
Asset Tag: <input type="text"/>	Timezone: UTC

3. If necessary, clear the existing logs.

System Event Log

Time	Severity	Description
2023-03-21 17:57:28 UTC	Normal	BIOS_POST_CMPLT: Presence sensor, Device Inserted / Device Pres
2023-03-21 17:57:27 UTC	Normal	System Software event: System Event sensor, OEM System Boot Eve

- a. Click the hamburger menu to show the available options.
- b. Choose **Faults and Logs**.
- c. In the main pane, choose the **System Event Log** tab and wait for the logs to load.
- d. If the log is full, click **Clear Log**.

4. Download the appropriate HUU ISO image.

Software Download

Downloads Home / Servers - Unified Computing / UCS C-Series Rack-Mount Standalone Server Software / UCS C220 M5 Rack Server Software / Unified Computing System (UCS) Server Firmware- 4.2(2g)

Search...

Expand All Collapse All

Suggested Release >

Latest Release >

All Release >

4.2 >

4.2(3d)

4.2(3b)

4.2(2g)

4.2(2f)

UCS C220 M5 Rack Server Software

Release 4.2(2g)

My Notifications

Related Links and Documentation
Release Note for 4.2(2g)

File Information	Release Date	Size
Cisco UCS Host Upgrade Utility ucs-c220m5-huu-4.2.2g.iso	23-Nov-2022	914.78 MB

Download icon (C)

a. Navigate to the software download page for your server model.

For UCS-C220-M5, browse to <https://software.cisco.com/download/home/286318809/type/283850974>.

For UCS-C225-M6, browse to <https://software.cisco.com/download/home/286329390/type/283850974>.

b. In the left sidebar, select the version supported by your target Nexus Dashboard release.

The list of supported releases is available in the Release Notes.

c. In the main pane, click on the **Download** icon.

d. Click **Accept License Agreement**.

5. Launch the KVM console from CIMC GUI.



If you are unable to open the KVM console, you may need to update Java version.

Cisco Integrated Management Controller

admin@ C220-WMP250600S0

Refresh | Host | **Launch vKVM** | Ping | CIMC Reboot | Locator LED | ? | i

Server Properties

Product Name: SE-NODE-G2
Serial Number: WMP250600S0
PID: SE-NODE-G2
UUID: 09A2D89E-A6C0-4F6D-9C91-2665E18FF8DC
BIOS Version: C220M5.4.1.2a.0.0624200115
Description:
Asset Tag:

Cisco Integrated Management Controller (Cisco IMC) Information

Hostname: C220-WMP250600S0
IP Address: 172.28.185.116
MAC Address: 48:8B:0A:45:EC:D0
Firmware Version: 4.1(2a)
Current Time (UTC): Tue Mar 21 21:07:09 2023
Local Time: Tue Mar 21 21:07:09 2023 UTC +0000
Timezone: UTC [Select Timezone](#)

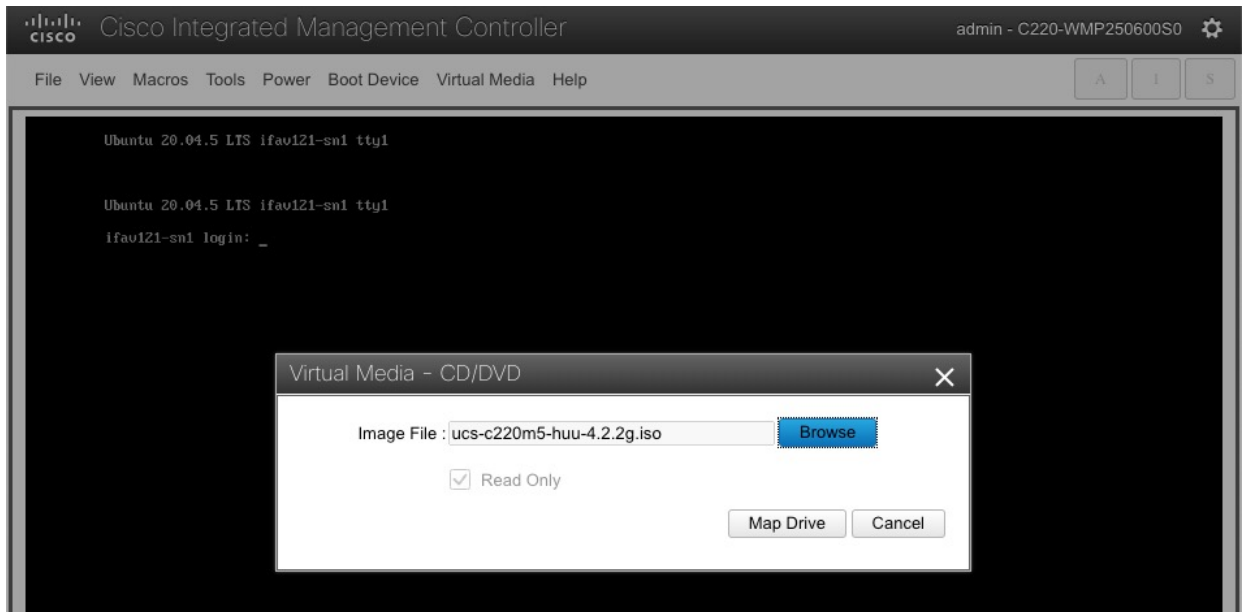
6. Mount the HUU ISO image you downloaded in Step 3.

a. From KVM console's **Virtual Media** menu, choose **Activate Virtual Devices**.

This adds virtual media options under the **Virtual Media** menu.

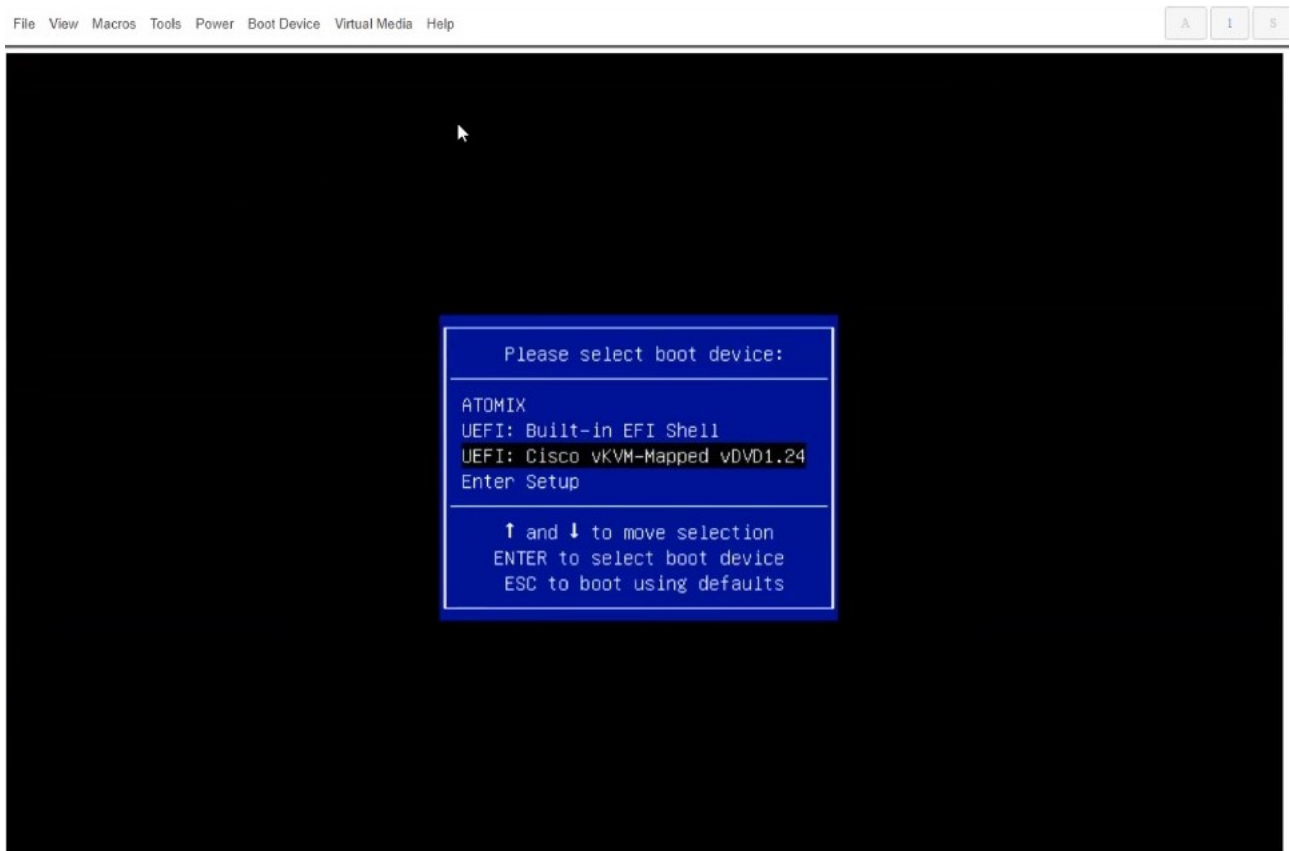
b. From KVM console's **Virtual Media** menu, choose **Map CD/DVD**.

c. In the **Virtual Media - CD/DVD** dialog that opens, click **Browse** and choose the HUU image.

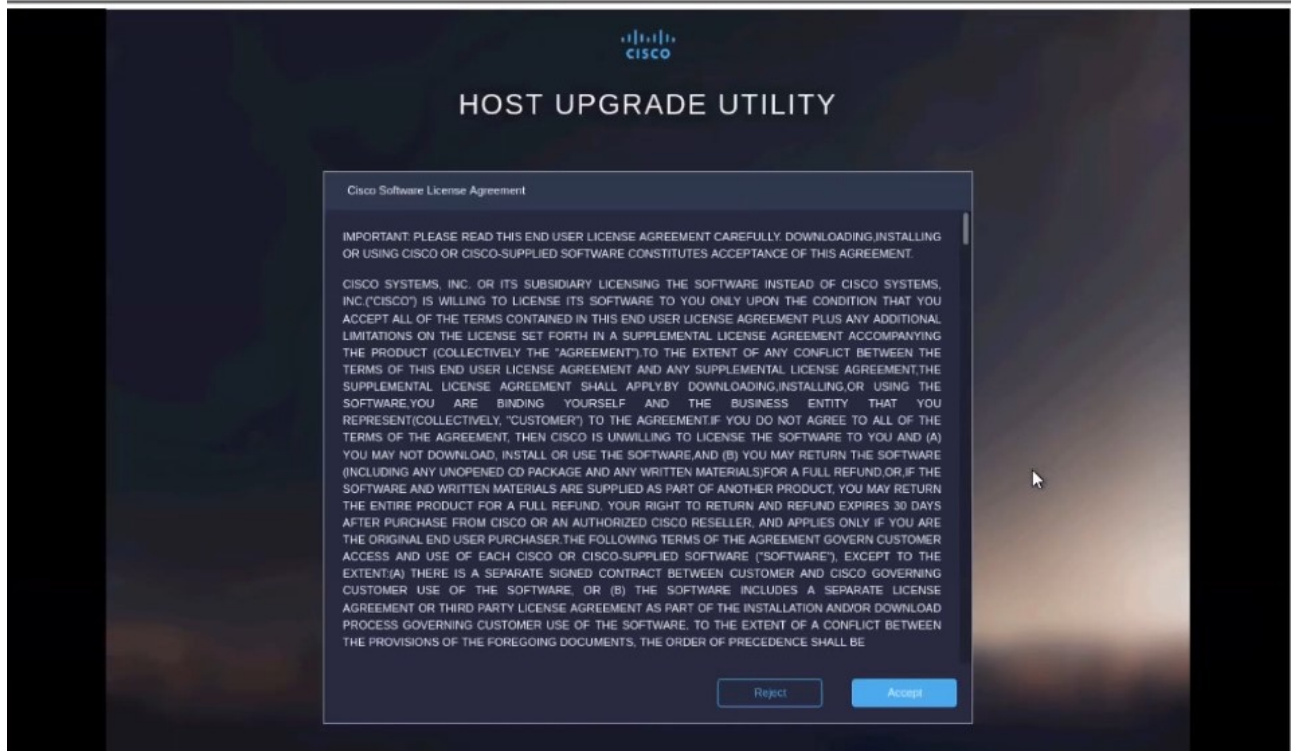


d. Finally, click **Map Drive**.

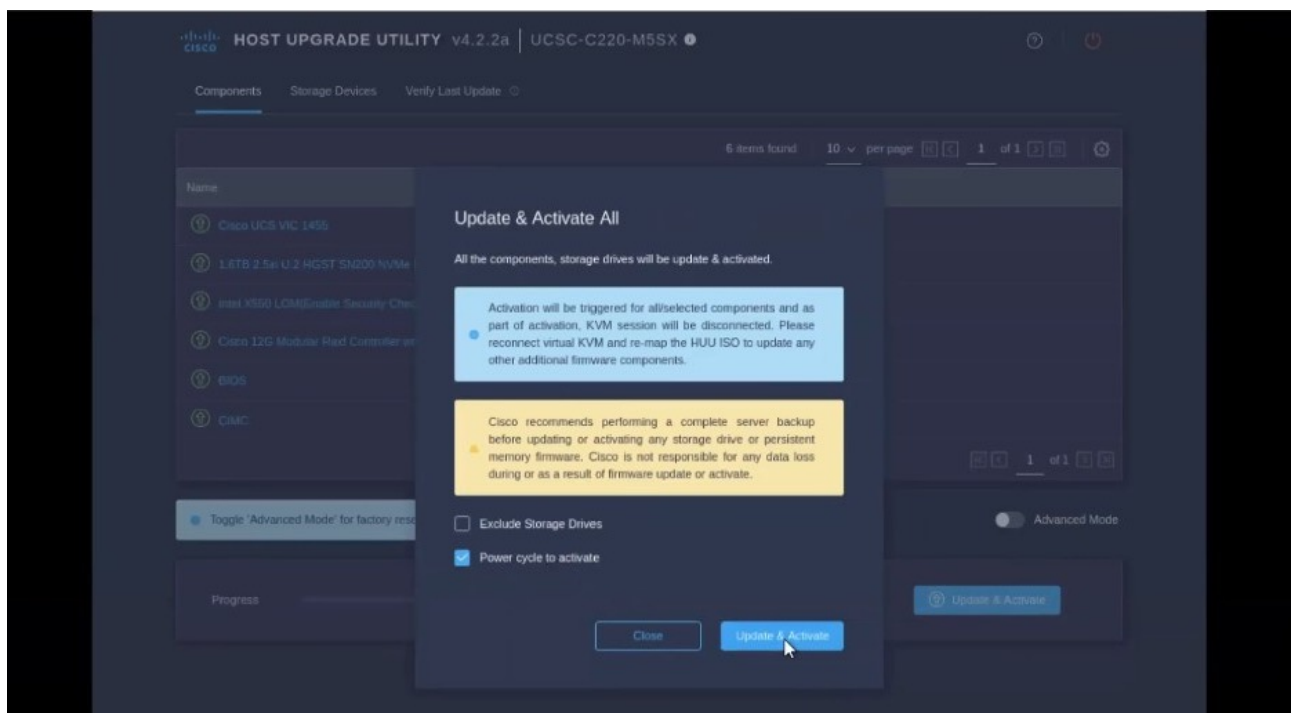
7. From KVM console's **Power** menu, choose **Power Cycle System** to reboot the server.
8. As the server is starting up, press **F6** to enter the boot menu and choose the **Cisco vKVM-Mapped vDVD**.



9. When prompted to accept Cisco Software License Agreement, choose **Accept**.



10. In the **Update & Activate All** dialog, choose **Update & Activate**.



You can verify that the upgrade was completed successfully through the GUI or by booting up the CIMC HUU and selecting the **Last Update Verify** option to ensure that all of the components were upgraded successfully.

11. After upgrade is completed, ensure that Trusted Platform Module State (TPM) is enabled.

You can check and enable it in the **BIOS > Configure BIOS > Security** menu.

Manual Cluster Upgrades

We recommend using the procedure described in [Firmware Management \(Cluster Upgrades\)](#) section to upgrade your cluster.

However, if you want to perform a manual upgrade of a single node (if you're adding a new node to the cluster but the node is running older firmware) or entire cluster (in case the GUI upgrade did not succeed), you can use the following steps instead.



If you're upgrading a single node that is running an older firmware to add it to an existing cluster, you will perform the following steps on that node only and not on the entire cluster.

1. Log in to the nodes you want to upgrade as **rescue-user**.
2. Copy the upgrade ISO image file into the **/tmp** directory on each node.
3. Start the upgrade on all nodes.

You can upgrade all nodes in parallel.

```
# acs installer update -f /tmp/nd-dk9.2.1.1a.iso
Warning: This command will initiate node update to new version.
Proceed? (y/n): y
Update in Progress ... Do not press Ctrl^C
```

4. Wait for the upgrade to complete.



You must wait for all nodes to finish upgrading before proceeding to the next step.

```
Update succeeded, reboot your host
```

5. Reboot one of the nodes.

Ensure that the upgrade is completed on all nodes as mentioned in the previous step before restarting any one node.

```
# acs reboot
This command will restart this device, Proceed? (y/n): y
```

6. Verify the node is healthy.

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

7. After the first node is successfully upgraded and healthy, reboot the other two nodes one at a time.



You must wait for the rebooted node to come up and ensure that the node is healthy using the `acs health` command before restarting the next node.

8. Once all nodes are up running the new version and are healthy, run post-upgrade tasks.

You can run the following command on all nodes in parallel.

```
# acs installer post-update
Warning: This command will run the post-update scripts. Proceed? (y/n): y
Update in Progress ... Do not press Ctrl^C
Post-update succeeded
```

9. Wait for the post-upgrade tasks to complete.

During this stage, the UI will show the progress, which looks similar to the initial cluster deployment. After the post-upgrade processes finish, you will be able to log in to the node as usual.

Re-Imaging Nodes

When you first receive the Nexus Dashboard physical hardware, it comes preloaded with the software image. If you simply want to configure the existing software, skip this section and proceed to [Managing Worker Nodes](#) or [Managing Standby Nodes](#).

If you are looking to manually upgrade the node to the latest software version, follow the instructions in [Manual Cluster Upgrades](#) instead.

This section describes how to redeploy the software stack on the Nexus Dashboard hardware. You may need to use the following steps in case of a catastrophic failure where you are no longer able to access the server's operating system and GUI, or in case you want to deploy a different release that does not support direct upgrade or downgrade from your existing cluster.



If you are planning to re-install an existing Nexus Dashboard cluster, you must clean up the site and app information first. In this case, ensure that the sites are disabled in all applications and removed from the ND cluster before bringing it down.

Before You Begin

- You must be able to connect to the server's CIMC using the Serial over LAN (SoL) port, so ensure that you have the server's CIMC IP address and an SSH client.

Detailed information about CIMC configuration is available at <https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-c-series-integrated-management-controller/products-installation-and-configuration-guides-list.html>

- Ensure that you are running a supported version of Cisco Integrated Management Controller (CIMC).

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

CIMC upgrade is described in detail in [Upgrading CIMC](#).

Installing Nexus Dashboard Using Remotely-Hosted Image

To re-install the Nexus Dashboard software:

1. Download the Cisco Nexus Dashboard image.
 - a. Browse to the Nexus Dashboard page and download the image.

<https://www.cisco.com/c/en/us/support/data-center-analytics/nexus-dashboard/series.html>

- b. Click the **Downloads** tab.
- c. Choose the Nexus Dashboard version you want to download.
- d. Download the Cisco Nexus Dashboard image (nd-dk9.<version>.iso).
- e. Host the image in a web server in your environment

You will need to provide an **http** URL when mounting the image.

2. Deploy the ISO to the server.

This step requires you to connect to the server's CIMC. Detailed information about CIMC configuration is available at <https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-c-series-integrated-management-controller/products-installation-and-configuration-guides-list.html>.

- a. SSH into the server's CIMC.
- b. Connect to the virtual media.

```
C220-WZP21510DHS# scope vmedia
C220-WZP21510DHS /vmedia #
```

- c. Map the Nexus Dashboard image you downloaded to the **CIMC-Mapped vDVD**.

```
C220-WZP21510DHS /vmedia # map-www image http://<ip-address>/<path>
<image>
```

For example:

```
C220-WZP21510DHS /vmedia # map-www image http://172.31.131.47/images nd-
dk9.2.0.1.iso
```

- d. Verify that the image is mounted.

```
C220-WZP21510DHS /vmedia # show mappings
```

Volume	Map-Status	Drive-Type	Remote-Share	Remote-File	Mount-Type
image	OK	CD	[<ip>/<path>]	nd-dk9.2.0.1.iso	www

- e. Reboot the server and connect to its console.

```
C220-WZP23150D4C /vmedia # exit
C220-WZP23150D4C# scope chassis
C220-WZP23150D4C /chassis # power cycle
C220-WZP23150D4C /chassis # exit
C220-WZP23150D4C# connect host
CISCO Serial Over LAN:
Press Ctrl+x to Exit the session
```

- f. Select the boot device.

Watch the boot process until you see the following message:

```
Press <F2> Setup, <F6> Boot Menu, <F7> Diagnostics, <F8> Cisco IMC
Configuration, <F12> Network Boot
```

Then press F6 and select the virtual media device where you mounted the image (**Cisco CIMC-Mapped vDVD1**):

```
/-----\
| Please select boot device: |
|-----|
| (Bus 05 Dev 00)PCI RAID Adapter |
| UNIGEN PHF16H0CM1-DTE PMAP |
| Cisco vKVM-Mapped vHDD1.22 |
| Cisco CIMC-Mapped vHDD1.22 |
| Cisco vKVM-Mapped vDVD1.22 |
| Cisco CIMC-Mapped vDVD1.22 |
| Cisco vKVM-Mapped vFDD1.22 |
| UEFI: Built-in EFI Shell |
| IBA GE Slot 0100 v1585 |
| IBA GE Slot 0101 v1585 |
| Enter Setup |
|-----|
| ^ and v to move selection |
| ENTER to select boot device |
| ESC to boot using defaults |
\-----/
```

g. Configure the networking.

When the server first boots, you will see the following output:

```
+ '[' -z http://172.31.131.47/nd-dk9.2.0.1.iso ']'
++ awk -F '/' '{print $4}'
+ urlip=172.31.131.47
+ '[' -z 172.31.131.47 ']'
+ break
+ '[' -n http://172.31.131.47/nd-dk9.2.0.1.iso ']'
+ set +e
+ configured=0
+ '[' 0 -eq 0 ']'
+ echo 'Configuring network interface'
Configuring network interface
+ echo 'type static, dhcp, bash for a shell to configure networking, or url to re-enter
the url: '
type static, dhcp, bash for a shell to configure networking, or url to re-enter the url:
+ read -p '? ' ntype
? static ①
+ case $ntype in
+ configure_static
+ echo 'Available interfaces'
Available interfaces
+ ls -l /sys/class/net
total 0
lrwxrwxrwx. 1 root root 0 Apr 26 01:21 enp1s0f0 ->
../devices/pci0000:3a/0000:3a:00.0/0000:3b:00.0/net/enp1s0f0
lrwxrwxrwx. 1 root root 0 Apr 26 01:21 enp1s0f1 ->
../devices/pci0000:3a/0000:3a:00.0/0000:3b:00.1/net/enp1s0f1
lrwxrwxrwx. 1 root root 0 Apr 26 01:21 enp1s0f4 ->
../devices/pci0000:5d/0000:5d:00.0/0000:5e:00.0/0000:5f:01.0/0000:61:00.0/000
0:62:00.0/0000:63:00.0/net/enp1s0f4
lrwxrwxrwx. 1 root root 0 Apr 26 01:21 enp1s0f5 ->
../devices/pci0000:5d/0000:5d:00.0/0000:5e:00.0/0000:5f:01.0/0000:61:00.0/000
0:62:00.0/0000:63:00.1/net/enp1s0f5
lrwxrwxrwx. 1 root root 0 Apr 26 01:21 lo -> ../devices/virtual/net/lo
+ read -p 'Interface to configure: ' interface
Interface to configure: enp1s0f0 ②
+ read -p 'address: ' addr
address: 172.23.53.59/21 ③
+ read -p 'gateway: ' gw
gateway: 172.23.48.1 ④
+ ip addr add 172.23.53.59/23 dev enp1s0f0
```

```
+ ip link set enp1s0f0 up
+ ip route add default via 172.23.48.1
RTNETLINK answers: Network is unreachable
++ seq 1 2
+ for count in '$(seq 1 2)'
+ ping -c 1 172.31.131.47
```

- ① For IP address, enter **dchp** if there is a DHCP server in your environment or **static**.
- ② For the interface, enter the first management port (**enp1s0f0**).
- ③ If you chose **static**, provide the IP address for the connection.
- ④ If you chose **static**, provide the gateway for the connection.

3. After the server boots from the provided image, select the only available installation option.

It may take up to 20 minutes for the installation process to complete.

After the image is deployed, you can add the node to your cluster as described in [Managing Worker Nodes](#) or [Managing Standby Nodes](#).

Rebuilding Existing Cluster

In some cases, you may need to re-build an existing cluster, for example if you want to change the data network's subnet or the nodes' data IP addresses, which requires redeploying the cluster.

1. Back up the Nexus Dashboard cluster configuration as described in [Backup and Restore](#).
2. Back up the configuration for all services deployed in your cluster.

For NDO, see **Operations > Backup and Restore** in the [Nexus Dashboard Orchestrator Configuration Guide](#).

For NDI, see **Operations > Backup and Restore** in the [Nexus Dashboard Insights User Guide](#).

For NDFC, see **Operations > Backup and Restore** in the [NDFC Fabric Controller Configuration Guide](#).

3. If your cluster is deployed as a physical appliance...
 - a. Log in to each node as **rescue-user**.
 - b. On each node, run the **acs reboot factory-reset**.

This resets the node to factory settings and reboots it.

- c. Redeploy the cluster using the same hardware.

You can follow the same procedure as you did when you first deployed the cluster, which is described in the "Deploying as Physical Appliance" chapter of the [Nexus Dashboard Deployment Guide](#)

4. If your cluster is deployed in virtual machines (VMs)...
 - a. Power down existing VMs.

You can keep the existing cluster's VMs until you deploy a new cluster and restore services and their configuration in it. Then you can simply delete the old cluster's VMs.

b. Redeploy a brand new cluster.

You can follow the same procedure as you did when you first deployed the cluster, which is described in the "Deploying in VMware ESX" or "Deploying in Linux KVM" chapter of the [Nexus Dashboard Deployment Guide](#)

5. Restore Nexus Dashboard configuration as described in [\[Backups and Restore\]](#).
6. Install the service(s) you had deployed previously as described [\[Service Management\]](#).
7. Restore each service's configuration from the backups you created in Step 1.

For NDO, see **Operations > Backup and Restore** in the [Nexus Dashboard Orchestrator Configuration Guide](#).

For NDI, see **Operations > Backup and Restore** in the [Nexus Dashboard Insights User Guide](#).

For NDFC, see **Operations > Backup and Restore** in the [NDFC Fabric Controller Configuration Guide](#).

AppStore Errors

When attempting to access the **Services > AppStore** tab in the Nexus Dashboard GUI, you may encounter the following error:

```
{
  "error": "There was a problem proxying the request"
}
```

Cause

When a master node where the AppStore service is running fails, it may take up to 5 minutes for the AppStore services to relocate to another master node

Resolution

Simply wait for the services to recover and refresh the page.

Event Export

Syslog events are not reaching the intended external events monitoring service.

Cause

Most common cause of this issue is not configured or improperly configured Syslog destination server.

Resolution

Ensure that the external server configuration in **Cluster Configuration > Syslog** is correct. For more information, see [Cluster Configuration](#).

Cause 2

Remote server is allowing traffic from only a specific set of IP addresses and the traffic from the Nexus Dashboard nodes' IP addresses is not allowed.

Resolution 2

Update your external server's configuration to allow traffic from the Nexus Dashboard cluster nodes.

Factory Reset

You can reset the entire physical cluster by running the following command on each node:

```
# acs reboot factory-reset
```



Doing this will lose all cluster configuration and applications and you will need to rebuild the cluster.

If you have a virtual or cloud Nexus Dashboard cluster, we recommend simply deleting the existing VMs and re-deploying the entire cluster instead of resetting all the nodes, as described in the [Cisco Nexus Dashboard Deployment Guide](#).

Changing Node IP Addresses

Changing the data network IP address is not supported. If you want to change the data IP address for the cluster nodes, you must re-create the cluster.

If you are running a single-node cluster, changing the management IP address is also not supported without re-creating the cluster.

If you are running a multi-node cluster, you can change the management IP addresses of one or more nodes as follows:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the main navigation menu, select **System Resources > Nodes**.
3. From the (...) menu next to the node, choose **Edit Node**.

Note that you can only change the IP address of a node that you are not currently logged in to. To change the IP of the current node, navigate to a different node's management IP address, log in, and repeat this procedure for the last node.

4. Update the **Management Network Address** and **Management Network Gateway** for the node.

For example, **172.31.140.58/24** and **172.31.140.1** respectively.

5. Click **Save**.

The changes will take effect immediately and you can access the nodes using the new IP addresses.

Cluster Configuration Errors

When you configure or change the proxy server in Nexus Dashboard, you may see a number of `cisco-mso service: Replicaset() not in desired state` errors in the **Cluster Configuration** page.

Cause

The errors are displayed while the service is restarting and will resolve on their own within 30-60 seconds.

Resolution

Simply wait for the services to recover and refresh the page.

Two-Factor Authentication (2FA) Not Prompting for Login Info

After the initial login using two-factor authentication, subsequent login attempts do not ask for username and password information and present a blank page instead.

Cause

The cookie timeout configured for the OIDC application is longer than the authentication token timeout set in the Nexus Dashboard.

Resolution

Clear your browser cache and the authentication process will work as expected.

Red Hat Enterprise Linux (RHEL) Deployments

You can view the installation logs by logging into your RHEL system and checking the `/logs/ndlinux/` directory.

In order to run the common Nexus Dashboard troubleshooting commands described in the [Troubleshooting](#) sections, you must first access the Nexus Dashboard environment.

To access the Nexus Dashboard environment from your RHEL system:

1. Log in to your RHEL system using the Nexus Dashboard user you provided in the YAML configuration file during installation.
2. Run the `attach-nd` command to access the Nexus Dashboard environment.

```
/usr/bin/attach-nd
```

After you access the Nexus Dashboard environment, you can use all the common Nexus Dashboard commands described in the [Troubleshooting](#) section of this guide.

Unable to Connect to Site After APIC Configuration Import

When you onboard a Cisco APIC site to Nexus Dashboard, APIC configuration is updated to reflect the onboarding. If you subsequently import an earlier configuration in APIC, the site may show as unavailable in Nexus Dashboard or services.

Cause

Earlier site configuration does not contain information specific to the Nexus Dashboard cluster where it is onboarded.

Resolution

We recommend exporting APIC configuration after the site is onboarded in Nexus Dashboard for any future config restores.

To resolve the issue after it occurs, you can re-register the site in the Nexus Dashboard GUI:

1. Log in to your Nexus Dashboard cluster.
2. Navigate to **Admin Console > Sites**
3. From the **Actions (...)** menu next to the site, select **Edit Site**.
4. In the **Site Edit** screen, check the **Re-register Site** checkbox and provide the site details again.
5. Click **Save**.

Re-Adding Same Master Node to Physical Cluster

This section describes how to re-add a master node to a physical cluster. This scenario can happen if the node was accidentally or deliberately removed via configuration reset (such as **acs reboot factory-reset**) or vMedia re-install.

If you have a standby node in your cluster, simply convert the standby into a master node as described in [Replacing Single Master Node with Standby Node](#) and then add the old master node as a new standby node as described in [Adding Standby Nodes](#).

If you need to completely replace (RMA) a master node due to hardware failure and do not have a standby node available, follow the procedure described in [\[Replacing Single Physical Master Node without Standby Node\]](#) instead.

To re-add the master node to the same cluster:

1. Ensure that the node is reset to factory settings.

If the node is in a bad state, log in to the node as **rescue-user** and reset the node using the following command:

```
# acs reboot factory-reset
```

2. Log in to the Nexus Dashboard GUI using the management IP address of one of the healthy nodes.

3. Navigate to **System Resources > Nodes**.

The node you want to replace will be listed as **Inactive** in the UI.

4. From the actions (...) menu for the node, select **Register**.

Register Node page will open.

5. In the **Register Node** page, provide the required information and click **Validate**.

For physical nodes, you need to provide the CIMC IP address and login information.

For virtual nodes, the management IP address will be retained and you need to provide only the password for the **rescue-user**.

6. Ensure the rest of the node information is accurate.

7. Click **Register** to re-register the node and re-add it as a **master** node to the cluster.

It will take up to 20 minutes to bootstrap, configure, and re-add the node. After it's done, the node will show as an **Active** master node in the UI.

Replacing a Single Virtual Master Node Without a Standby Node

This section describes how to recover from a master node failure in a VMware ESX or Linux KVM virtual Nexus Dashboard cluster. The procedure involves deploying a brand new Nexus Dashboard node using the same form factor as the node which you are replacing and joining it as a master node to the remaining cluster.

1. Ensure that the failed node's VM is powered down.
2. Bring up a new Nexus Dashboard node.

Bringing up an additional node in VMware ESX is described in [Deploying Additional Virtual Nodes in VMware ESX](#). Note that you must bring up a node of the same type (**OVA-App** or **OVA-Data**) as the node you are replacing.

Bringing up an additional node in Linux KVM is described in [Deploying Additional Virtual Nodes in Linux KVM](#).



Ensure that you use the same exact network configuration settings as you used for the failed node.

3. Power on the new node's VM and wait for it to boot up.
4. Log in to the Nexus Dashboard GUI.

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

5. Replace the node.
 - a. From the left navigation pane, select **System Resources > Nodes**.

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

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

The **Replace** window will open.

- c. Provide the **Management IP Address** and **Password** for the node, then click **Verify**.

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

- d. Click **Replace**.

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

Replacing a Single Physical Master Node Without a Standby Node

The following section describes how to recover from a single master node failure in a physical Nexus Dashboard cluster without a standby node. This procedure is for hardware issues that require it to be physically replaced. If the node is simply in a bad software state, you can use the **acs reboot clean** commands instead and re-add the same node to the cluster as described in [Re-Adding Same Master Node to Physical Cluster](#).

If your cluster has a standby node configured, we recommend using the steps described in [Replacing Single Master Node with Standby Node](#) instead.

Before you begin

- Ensure that at least 2 master nodes are healthy.

If two of the master nodes are unavailable, you will need to manually restore the cluster as described in [Replacing Two Master Nodes with Standby Nodes](#)

- Ensure that the master node you want to replace is powered off.
- Prepare and deploy the new node as described in [Deploying Additional Physical Nodes](#).
- Ensure that you have the same CIMC IP address and login information on the new node as you configured for the failed node.

The remaining master nodes will use the CIMC information to restore configuration to the new node.

- Ensure that the new node is powered on and note down its serial number.

To replace a single failed master node:

1. Log in to your Nexus Dashboard GUI using the management IP of one of the other **master** nodes.
2. From the main navigation menu, select **System Resources > Nodes**.
3. In the nodes list, find the **Serial** number of the node you want to replace and ensure that the node's **Status** shows **Inactive**.
4. In the Nexus Dashboard's **Nodes** screen, select the inactive node by clicking the checkbox next to it.
5. From the **Actions** menu, select **Replace**.

6. In the **New Serial Number** field, provide the serial number of the new node and click **Replace**.

After the process is completed, you will see the serial number of the old node updated to the new node's serial number and the status will change to **Active** once the new master has successfully joined the cluster.

Replacing Worker or Standby Nodes

When replacing a failed worker or standby node, you can simply delete the **Inactive** node from the GUI and then deploy a brand new worker or standby node as you typically would.

Before You begin

- Ensure that the worker node you want to replace is powered off.

To replace a failed worker or standby node:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the main navigation menu, select **System Resources > Nodes**.
3. In the nodes list, find the **Serial** number of the node you want to replace and ensure that the node's **Status** shows **Inactive**.
4. Select the inactive node by clicking the checkbox next to it.
5. From the **Actions** menu, select **Delete**.

This will remove the failed node from the list.

6. Power on the new node and add it as a new **worker** or **standby** node to the cluster as described in [Managing Worker Nodes](#) or [Managing Standby Nodes](#).

You can use the same configuration parameters as you used to set up the old node.

Initial Cluster Bootstrap Issues

This section describes the different stages of the initial cluster bootstrap process and summarizes some common issues you may run into when first deploying your Nexus Dashboard cluster.

After you bring up the nodes and provide each node's information during the GUI setup, the initial bootstrap process goes through a number of stages to bring up the nodes, configure the required information, and create the cluster. The bootstrap screen allows you to track the progress and indicates any issues that may come up:




Figure 17. Bootstrap Progress

- **Bootstrap Master Node** and **Bootstrap Peer Nodes**—bring up the first master node with the management and data networks IP addresses you provided. Then brings up the 2nd and 3rd master nodes with their respective IPs.

If the process fails at one of these stages, connect to each node’s console and verify that all the information you provided is correct. You can view the configuration you provided using the **acs system-config** command.

You can also check the bootstrap logs (**/logs/k8/install.log**) for additional details.

Typically, you can resolve any issues caused by misconfiguration by resetting the node using **acs reboot factory-reset** and restarting the setup process.

- **Discover Cluster Members**—establishes connectivity between all master nodes in the cluster over the data network.

Failures at this stage typically indicate misconfiguration of the data network IP address and the node being unable to reach its other 2 peers.

You can use `acs cluster masters` command on any of the nodes to confirm the data IP you have provided.

If the command does not return any information, use `ip addr` to check the data interface's (`bond0br`) IP address and ensure that all nodes' IPs are reachable from the other nodes.

```
$ ip addr
[..]
6: bond0br: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state
UP group default qlen 1000
    link/ether 52:54:00:e1:93:06 brd ff:ff:ff:ff:ff:ff
    inet 10.195.255.165/24 brd 10.195.255.255 scope global bond0br
        valid_lft forever preferred_lft forever
    inet6 fe80::5054:ff:fee1:9306/64 scope link
        valid_lft forever preferred_lft forever
[..]
```

- **Setup Security**—sets up Key Management Service (KMS) to enable data encryption between the nodes.

If the `acs cluster masters` command returns `ca cert not found` error, it indicates a KMS issue. For additional details, check the `/logs/kms` logs.

- **Setup Initial Cluster** and **Bootstrap Kubernetes Cluster**—any failures during these stages typically indicate Kubernetes issues.

You can get additional details from the logs in `/logs/k8` on each node.

- After the **Bootstrap** stages are complete, the process advances to the **Cluster Bringup** stages.

From **Initialize System** to the **Wait for infra services to be ready** stages finalize the cluster creation by bringing up the remaining services.

At this stage, you can use the `acs health` command on any of the nodes to see which service is not coming up correctly. Then check the specific service's logs in `/logs/k8_infra/<service>`

Multi-Cluster Connectivity Issues

The following sections list common issues with multi-cluster connectivity.

For additional information about connecting multiple clusters together, see [Multi-Cluster Connectivity](#).

Non-Primary Cluster Unable to Reconnect

If you clean reboot and redeploy a cluster that was part of a multi-cluster connectivity group, the group's primary cluster will not be able to recognize it and will indicate that the cluster remains unreachable.

To resolve this issue, disconnect and reconnect the cluster:

1. Log in to the primary cluster.
2. Remove the cluster you re-deployed from the group.

This is described in [Disconnecting Clusters](#).

3. Re-add the cluster to the group.

This is described in [Connecting Multiple Clusters](#).

Non-Primary Cluster Redeployed with Older Version

If for any reason you redeploy one of the non-primary clusters in the group with a version of Nexus Dashboard that does not support this feature, the primary cluster will still be able to connect to that cluster, but will not be able to retrieve any information and the UI will remain blank.

To resolve this issue, remove that cluster from the group:

1. Log in to the primary cluster as a local **admin** user.

If you log in with the remote user shared across all clusters, the UI page will remain blank.

2. Remove the cluster you re-deployed from the group.

This is described in [Disconnecting Clusters](#).

3. Log out and log back in using the remote user you use to manage the multi-cluster connectivity and verify that UI loads correctly.

Generating Private Key, Creating CSR, and Obtaining CA-Signed Certificate

This section provides an example of how to generate a private key, create a certificate signing request (CSR), and obtain a certificate signed by a Certificate Authority (CA) for use in your Nexus Dashboard cluster.

If you want to generate both a key and a self-signed certificate, skip this section and follow the steps described in [Generating Private Key and Self-Signed Certificate](#) instead.

The configuration steps required to add the keys and certificates in the Nexus Dashboard GUI are described in the [Security](#) chapter.

1. Generate private key.

You can generate the private key on any platform that has OpenSSL installed or you can SSH into one of your Nexus Dashboard nodes as the **rescue-user** and perform these steps there.

```
[rescue-user@localhost ~]$ openssl genrsa -out nd.key 2048
Generating RSA private key, 2048 bit long modulus
.....
.....
e is 65537 (0x10001)
```



```
[rescue-user@localhost ~]$ ls
nd.key
```

2. Generate your CSR signed with the private key you generated in the first step.
 - a. Create the CSR configuration file (`csr.cfg`) with the required information.

An example configuration file is shown below:

```
[req]
default_bits = 2048
distinguished_name = req_distinguished_name
req_extensions = req_ext
prompt = no
[req_distinguished_name]
countryName = US
stateOrProvinceName = Texas
localityName = Plano
organizationName = CSS
organizationalUnitName = DC
commonName = nd.dc.css
emailAddress = no-reply@mydomain.com
[req_ext]
subjectAltName = @alt_names
[alt_names]
DNS.1 = *.dc.css
IP.1 = 10.0.0.96
IP.2 = 10.0.0.97
```

- b. Generate your CSR.

```
[rescue-user@localhost ~]$ openssl req -new -key nd.key -out nd.csr -config
csr.cfg
[rescue-user@localhost ~]$ ls
csr.cfg nd.csr nd.key
```

You can view the generated CSR using the following command:

```
[rescue-user@localhost ~]$ openssl req -in nd.csr -text -noout
```

3. Obtain a CA-signed certificate.

In production deployments, you will provide the CSR (`ca.csr`) from the previous step to a public CA, such as IdenTrust or DigiCert, to obtain the CA-signed certificate (`ca.crt`).

4. Verify the signed certificate.

The following command assumes you copied the CA-signed certificate (`ca.crt`) into the same folder as the private key you generated.

```
[rescue-user@localhost ~]$ openssl verify -verbose -CAfile ca.crt nd.crt
nd.crt: OK
```

5. Add the contents of the generated files in your Nexus Dashboard's GUI.

Following the steps described in [Security Configuration](#), where you will need to provide the contents of the following 3 files generated in the previous steps:

- o Private key (`nd.key`)
- o Certificate Authority's (CA) public certificate (`ca.crt`)
- o CA-signed certificate (`nd.crt`)

Generating Private Key and Self-Signed Certificate

This section provides an example of how to generate a private key and custom certificates should you want to use them in your Nexus Dashboard cluster.

If you want to use a CA-signed certificate, skip this section and follow the steps described in [Creating CSR, and Obtaining CA-Signed Certificate](#).

The configuration steps required to add the keys and certificates in the Nexus Dashboard GUI are described in the [Security](#) chapter.

1. Generate private key.

You can generate the private key on any platform that has OpenSSL installed or you can SSH into one of your Nexus Dashboard nodes as the `rescue-user` and perform these steps there.

```
[rescue-user@localhost ~]$ openssl genrsa -out nd.key 2048
Generating RSA private key, 2048 bit long modulus
.....
.....
e is 65537 (0x10001)
[rescue-user@localhost ~]$ ls
nd.key
```

2. Generate Certificate Authority (CA) key.

To generate a self-signed CA, for example for lab and testing purposes, run the following command:

```
[rescue-user@localhost ~]$ openssl genrsa -out ca.key 2048
```

```
Generating RSA private key, 2048 bit long modulus
```

```
.....
```

```
.....
```

```
e is 65537 (0x10001)
```

```
[rescue-user@localhost ~]$ ls
```

```
ca.key nd.key
```

3. Generate CSR for the CA.

```
[rescue-user@localhost ~]$ openssl req -new -key ca.key -subj
```

```
"/CN=Self/C=US/O=Private/ST=Texas" -out ca.csr
```

```
[rescue-user@localhost ~]$ ls
```

```
ca.csr ca.key nd.key
```

You can view the generated CSR using the following command:

```
[rescue-user@localhost ~]$ openssl req -in ca.csr -text -noout
```

4. Create self-signed root certificate.

```
[rescue-user@localhost ~]$ openssl x509 -req -in ca.csr -signkey ca.key
```

```
-CAcreateserial -out ca.crt -days 3650
```

```
Signature ok
```

```
subject=/CN=Self/C=US/O=Private/ST=Texas
```

```
Getting Private key
```

```
[rescue-user@localhost ~]$ ls
```

```
ca.crt ca.csr ca.key nd.key
```

You can view the generated root certificate using the following command:

```
[rescue-user@localhost ~]$ openssl x509 -in ca.crt -text -noout
```

5. Generate your CSR signed with the private key you generated in the first step.

a. Create the CSR configuration file (`csr.cfg`) with the required information.

An example configuration file is shown below:

```
[req]
default_bits = 2048
distinguished_name = req_distinguished_name
req_extensions = req_ext
prompt = no
```

```
[req_distinguished_name]
countryName = US
stateOrProvinceName = Texas
localityName = Plano
organizationName = CSS
organizationalUnitName = DC
commonName = nd.dc.css
emailAddress = no-reply@mydomain.com
[req_ext]
subjectAltName = @alt_names
[alt_names]
DNS.1 = *.dc.css
IP.1 = 10.0.0.96
IP.2 = 10.0.0.97
```

b. Generate your CSR.

```
[rescue-user@localhost ~]$ openssl req -new -key nd.key -out nd.csr -config
csr.cfg
[rescue-user@localhost ~]$ ls
ca.crt ca.csr ca.key csr.cfg nd.csr nd.key
```

You can view the generated CSR using the following command:

```
[rescue-user@localhost ~]$ openssl req -in nd.csr -text -noout
```

6. Self-sign the certificate you generated.

```
[rescue-user@localhost ~]$ openssl x509 -req -in nd.csr -CA ca.crt -CAkey ca.key
-CAcreateserial -out nd.crt -days 3600
Signature ok
subject=/C=US/ST=Texas/L=Plano/O=CSS/OU=DC/CN=nd.dc.css/emailAddress=no-
reply@mydomain.com
Getting CA Private Key
[rescue-user@localhost ~]$ ls
ca.crt ca.csr ca.key ca.srl csr.cfg nd.crt nd.csr nd.key
```

7. Verify the signed certificate.

```
[rescue-user@localhost ~]$ openssl verify -verbose -CAfile ca.crt nd.crt
nd.crt: OK
```

8. Add the contents of the generated files in your Nexus Dashboard's GUI.

Following the steps described in [Security Configuration](#), where you will need to provide the contents of the following 3 files generated in the previous steps:

- o Private key ([nd.key](#))
- o Certificate Authority's (CA) public certificate ([ca.crt](#))
- o CA-signed certificate ([nd.crt](#))

Updating NDO Configuration After Replacing Switch Devices Managed by NDFC

If your Nexus Dashboard Fabric Controller (NDFC) fabric is managed by Nexus Dashboard Orchestrator (NDO) and you replace one or more devices that are managed by the NDFC, you must ensure that NDO is aware of the new switch serial numbers.

The following sections provide a summary of the steps required to synchronize the new fabric device's information with NDO.

Replacing a Core or Route Server (RS) Device

1. Log in to NDFC.
2. To replace a physical switch in a Fabric when using NDFC Easy Fabric mode, follow the Return Material Authorization (RMA) steps mentioned in the [Cisco NDFC Fabric Controller Configuration Guide](#).
3. Log in to NDO.
4. Navigate to **Infrastructure > Site Connectivity**.
5. Click **Refresh** on the **Control Plane Configuration** in the **General Settings** page where the RS/Core is present.
6. Click **Deploy**.

Replacing a Leaf Switch

1. Log in to NDFC.
2. To replace a physical switch in a Fabric when using NDFC Easy Fabric mode, follow the Return Material Authorization (RMA) steps mentioned in the [Cisco NDFC Fabric Controller Configuration Guide](#).
3. Log in to NDO.
4. Navigate to **Application Management > Schema** and click the Schema/Template for that Site/Device.
5. Re-import VRF/Network that was present on the device:
 - a. In the **View Overview** drop-down list, select the template.
 - b. In the **Template Properties** section, click the VRF/Network from the **VRFs** box.
 - c. Select the site from the **Import** drop-down list.

- d. Select the VRF after clicking **VRF**.
- e. Click **Import**.

Replacing Border Gateway (BGW) Devices

1. Log in to NDFC.
2. To replace a physical switch in a Fabric when using NDFC Easy Fabric mode, follow the Return Material Authorization (RMA) steps mentioned in the [Cisco NDFC Fabric Controller Configuration Guide](#).
3. Log in to NDO.
4. Navigate to **Infrastructure > Site Connectivity**.
5. Click **Refresh** on the site where BGW is present and click **Deploy**.
6. Navigate to **Application Management > Schema** and click the Schema/Template for that Site/Device.
7. Re-import VRF/Network that was present on the device:
 - a. In the **View Overview** drop-down list, select the template.
 - b. In the **Template Properties** section, click the VRF/Network from the **VRFs** box.
 - c. Select the site from the **Import** drop-down list.
 - d. Select the VRF after clicking **VRF**.
 - e. Click **Import**.