Cisco Nexus Dashboard Deployment Guide, Release 2.1.x

First Published: 2021-09-14

Last Modified: 2022-09-05

CONTENTS

New and Changed Information

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 in which 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	New Feature or Update	Where Documented
2.1.2	Added proxy configuration during initial cluster bootstrap.	Deploying as Physical Appliance, on page 25 Deploying in Linux KVM, on page 51 Deploying in VMware ESX, on page 33
2.1.2	Added descriptions of services and required ports	Deployment Overview, on page 3
2.1.1	First release of this document.	--

Deployment Overview and Requirements

Deployment 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.

Each Nexus Dashboard cluster typically consists of 3 `master` nodes. In addition, you can 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*.

Note   This document describes initial configuration of the 3-node cluster. After your cluster is up and running, you can configure and deploy additional nodes as described in the *Cisco Nexus Dashboard User Guide*, which is also available directly from the Nexus Dashboard GUI.

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 the initial deployment of the Nexus Dashboard software; hardware setup is described in the *Nexus Dashboard Hardware Setup Guide*, while other Nexus Dashboard operations procedures are described in the *Cisco Nexus Dashboard User 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.

Beginning with Release 2.1(2), 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, there are no changes to service installation and 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.

- 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 three VMware ESX virtual machines.

- Linux KVM (`.qcow2`)

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

- Amazon Web Services (`.ami`)

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

- Microsoft Azure (`.arm`)

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

Upgrading From Previous Versions of Nexus Dashboard

If you are already running a Nexus Dashboard, Release 2.0.1 or later, you can upgrade directly to the latest release while retaining the cluster configuration and applications, as described in Upgrading Nexus Dashboard, on page 79

Upgrading From Application Services Engine

If you are running Cisco Application Services Engine, you must upgrade to Nexus Dashboard release 2.0.2g or later as described in *Cisco Nexus Dashboard Deployment Guide, Release 2.0(x)* before upgrading to Nexus Dashboard release 2.1.x.

Cluster Sizing Guidelines

Nexus Dashboard supports co-hosting of applications. Depending on the type and number of applications you choose to run, you may be required to deploy additional worker nodes in your cluster. For cluster sizing information and recommended number of nodes based on specific use cases, see Cisco Nexus Dashboard Cluster Sizing.

After your initial 3-node cluster is up and running, you can configure and deploy additional nodes as described in the *Cisco Nexus Dashboard User Guide*, which is also available directly from the Nexus Dashboard GUI.

Supported Services

For the full list of supported applications and the associated compatibility and interoperability information, see the Nexus Dashboard and Services Compatibility Matrix.

Prerequisites 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.

Note 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.

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.

Table 2: External Network Purpose

Data Network	Management Network
• Nexus Dashboard node clustering • Service to service communication • Nexus Dashboard nodes to Cisco APIC, Cloud APIC, and NDFC/DCNM communication For example, the network traffic for services such as Nexus Dashboard Insights.	• Accessing Nexus Dashboard GUI • Accessing Nexus Dashboard CLI via SSH • DNS and NTP communication • Nexus Dashboard firmware upload • Accessing Cisco DC App Center (AppStore) If you want to use the Nexus Dashboard App Store to install services, https://dcappcenter.cisco.com must be reachable via the Management Network • Intersight device connector

The two networks have the following requirements:

- 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/DCNM 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 table bellow summarize service-specific requirements for the management and data networks.

Note 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.

Also note that if the two interfaces are in the same subnet, either of the IPs may be used as the source for traffic. For example, if you have remote authentication configured, you must add both management and data IPs to the list of permitted IP addresses on your external authentication provider because either of the two interfaces may be used as the source for authentication traffic.

Allocating persistent IP addresses 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.

Table 3: Service-Specific Network Requirements

Nexus Dashboard Service	Management Interface	Data Interface	Total Number of Persistent IPs	Support for Data and Management in the same subnet
Nexus Dashboard Orchestrator	Layer 3 adjacent	Layer 3 adjacent	N/A	Yes However, we recommend separate subnets for the two networks
Nexus Dashboard Insights without SFLOW/NetFlow (ACI fabrics)	Layer 3 adjacent	Layer 3 adjacent	N/A	Yes However, we recommend separate subnets for the two networks
Nexus Dashboard Insights without SFLOW/NetFlow (NDFC/DCNM 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	No
Nexus Dashboard Insights with SFLOW/NetFlow (ACI or NDFC/DCNM fabrics)	Layer 3 adjacent	Layer 2 adjacent	6 IPs in data interface network	No

Nexus Dashboard Service	Management Interface	Data Interface	Total Number of Persistent IPs	Support for Data and Management in the same subnet
Nexus Dashboard Fabric Controller	Layer 2 adjacent	Layer 2 adjacent	One of the following: • 2 IPs in management network if using the default **LAN Device Management Connectivity** setting • 2 IPs in data network if setting **LAN Device Management Connectivity** to `Data` Plus 1 IP per fabric for EPL in data network	No
Nexus Dashboard Data Broker	Layer 3 adjacent	N/A	N/A	Yes

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

✎

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

Table 4: RTT Requirements

Service	Connectivity	Maximum RTT
Nexus Dashboard Orchestrator	Between nodes	50 ms
	To sites	For APIC sites: 500 ms For NDFC/DCNM sites: 150 ms

Service	Connectivity	Maximum RTT
Nexus Dashboard Insights	Between nodes	50 ms
	To sites	50 ms
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.

Note 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 external to the cluster, which 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.

Communication Ports

The following ports are required by the Nexus Dashboard cluster and its services:

Table 5: Nexus Dashboard Communication Ports (Management Network)

Service	Port	Protocol	Direction In—towards the cluster Out—from the cluster towards the fabric or outside world	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
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 6: Nexus Dashboard Communication Ports (Data Network)

Service	Port	Protocol	Direction In—towards the cluster Out—from the cluster towards the fabric or outside world	Connection
SSH	22	TCP	Out	Inband of switches and APIC
HTTPS	443	TCP	Out	Inband of switches and APIC/NDFC/DCNM
VXLAN	4789	UDP	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 15223 30002-30006 30009-30010 30012 30014-30015 30018-30019 30025 30027	TCP	In/Out	Other cluster nodes
Kafka	30001	TCP	In/Out	Inband of switches and APIC/NDFC/DCNM
Infra-Service	30016 30017	TCP/UDP	In/Out	Other cluster nodes

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
Infra-Service	30500-30600	TCP/UDP	In/Out	Other cluster nodes

Table 7: Nexus Dashboard Insights Communication Ports (Data Network)

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
Show Techcollection	2022	TCP	In/Out	Inband of switches and APIC/NDFC/DCNM
Flow Telemetry	5640-5671	UDP	In	Inband of switches
TAC Assist	8884	TCP	In/Out	External
KMS	9989	TCP	In/Out	Other cluster nodes and ACI fabrics
SW Telemetry	5695 30000 57500 30570	TCP	In/Out	Other cluster nodes

Fabric Connectivity

The following sections describe how to connect your Nexus Dashboard cluster to your fabrics.

For on-premises APIC or NDFC/DCNM fabrics, you can connect the Nexus Dashboard cluster in one of 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 Cloud APIC fabrics, you will need to connect via a Layer 3 network.

Physical Node Cabling

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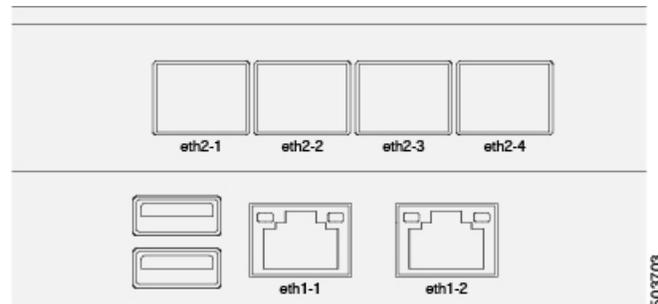
The interfaces are configured as Linux bonds (one for the data interfaces and one for the management interfaces) running in active-standby mode. All interfaces must be connected to individual host ports, PortChannel or vPC are not supported.

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 establish 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/DCNM fabrics, you must establish connectivity from the data interface to the in-band interface of each site's DCNM.

- 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 a 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/DCNM fabrics, if the data interface and DCNM's inband interface are in different subnets, you must add a route on NDFC/DCNM to reach the Nexus Dashboard's data network address.

You can add the route from the NDFC/DCNM 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 a 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 Layer 3 Network, Day-2 Operations Applications
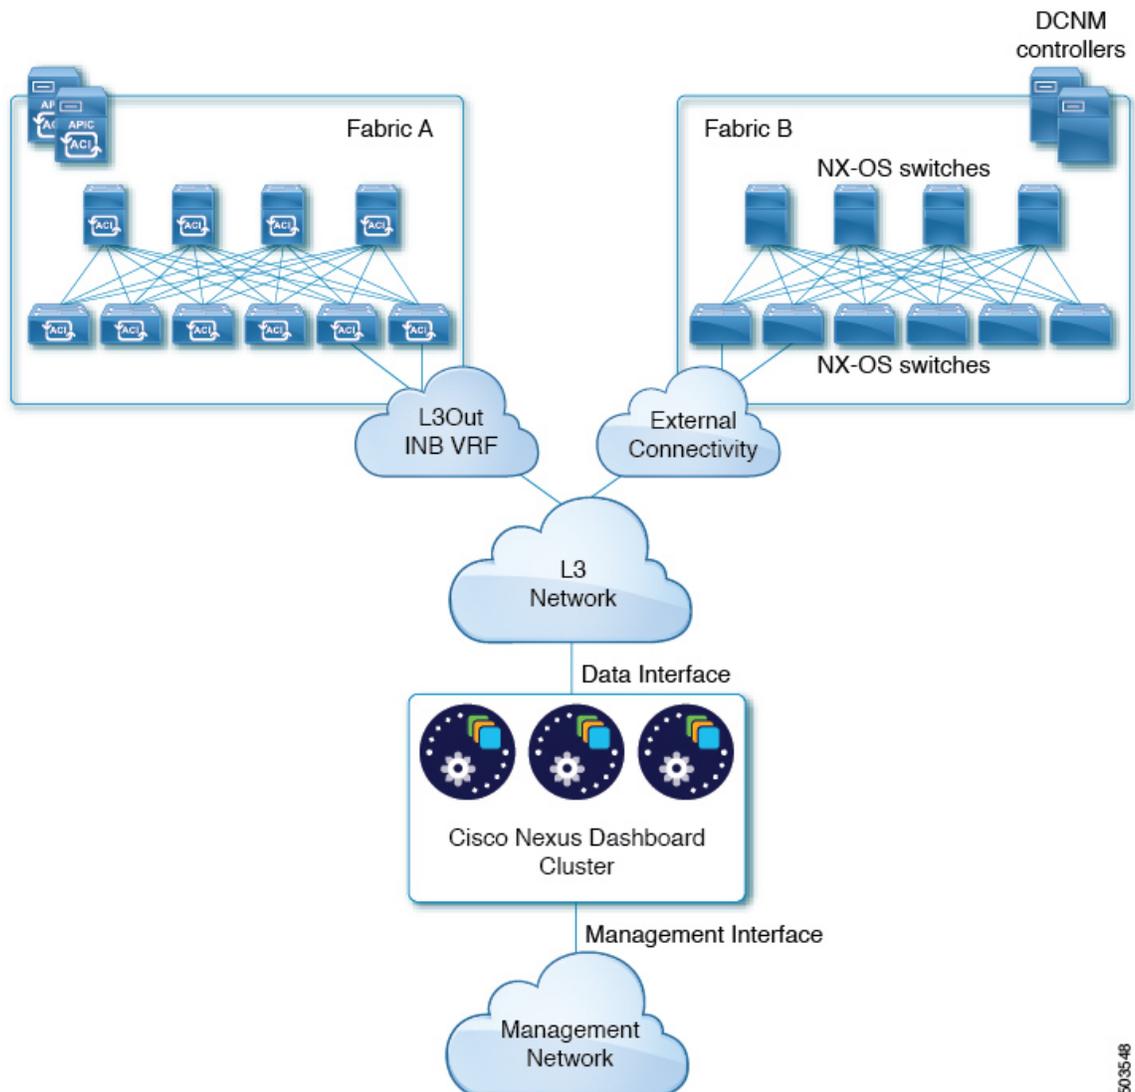
Figure 3: Connecting via Layer 3 Network, Nexus Dashboard Orchestrator
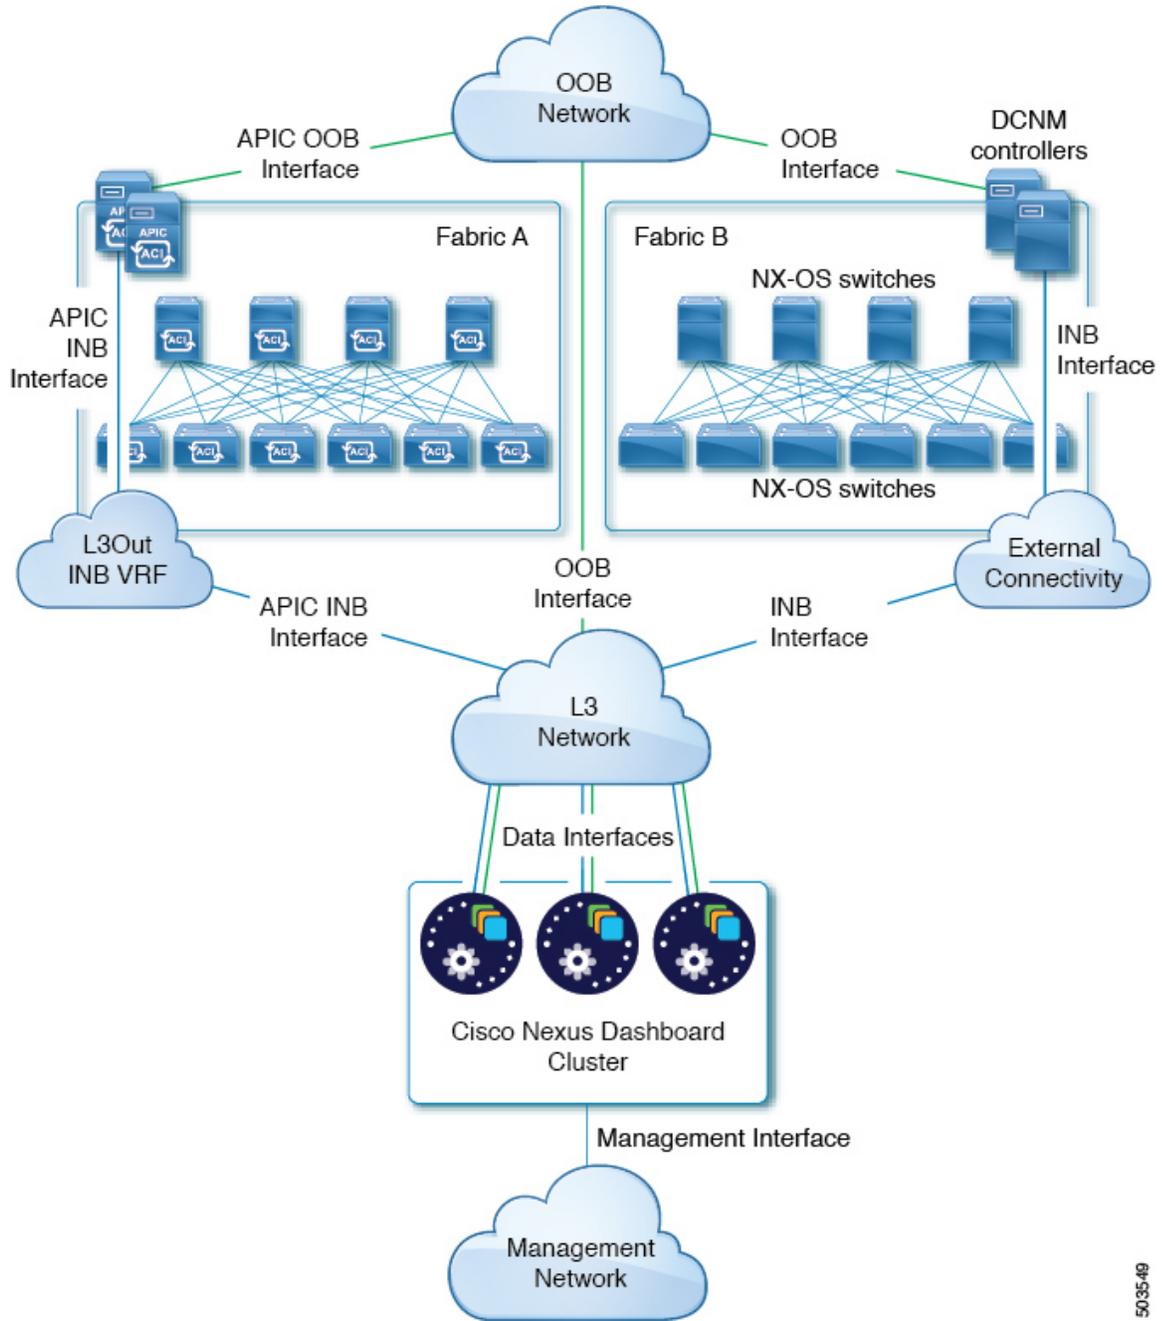
Connecting Nodes Directly to Leaf Switches

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, 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 Nexus Dashboard 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.

Note that the "L3 Network" and the "Management Network" can be the same network infrastructure, for example in case the Nexus Dashboard nodes have the management and data network interfaces in the same subnet.

Figure 4: Connecting Directly to Leaf Switches, Day-2 Operations Applications
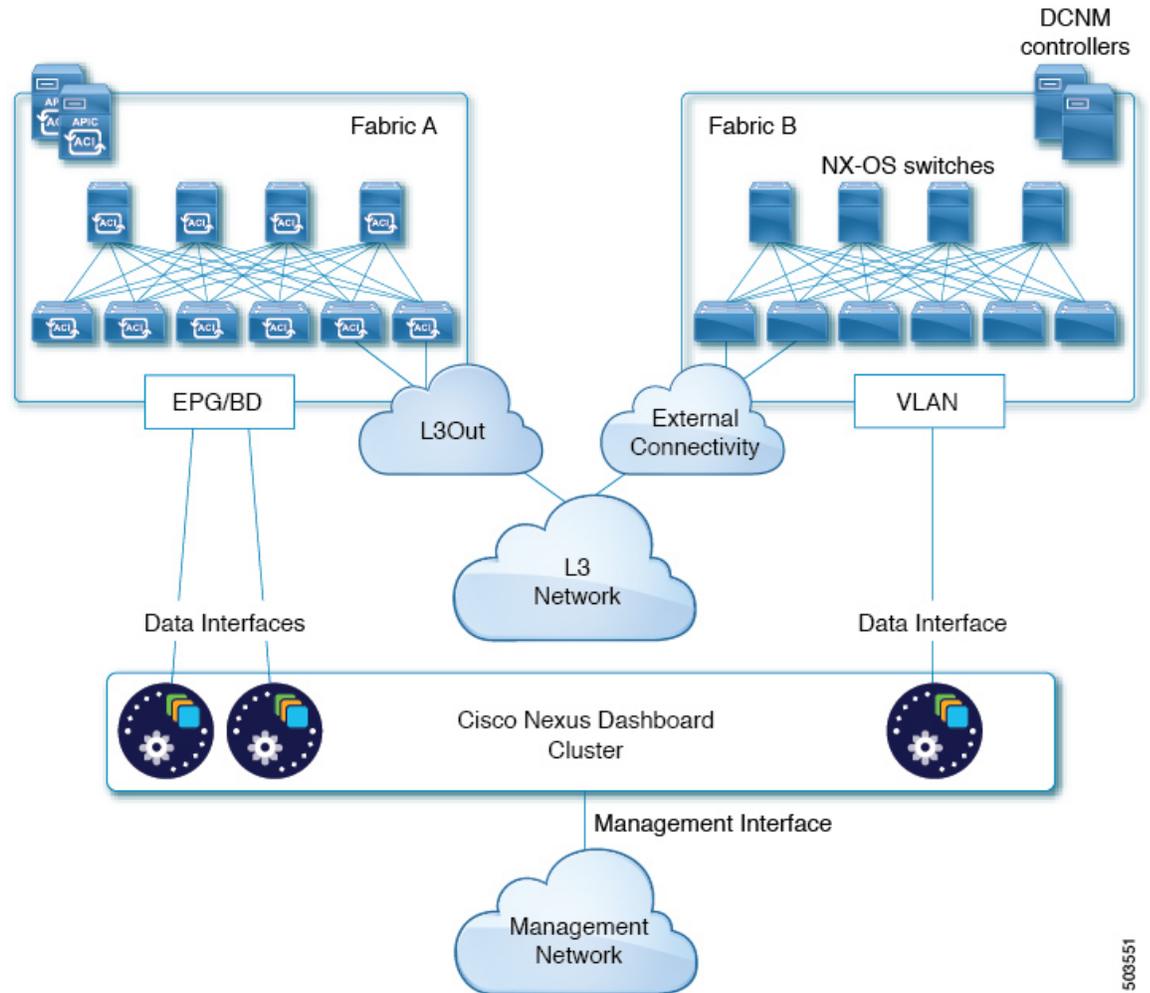
Figure 5: Connecting Directly to Leaf Switches, Nexus Dashboard Orchestrator
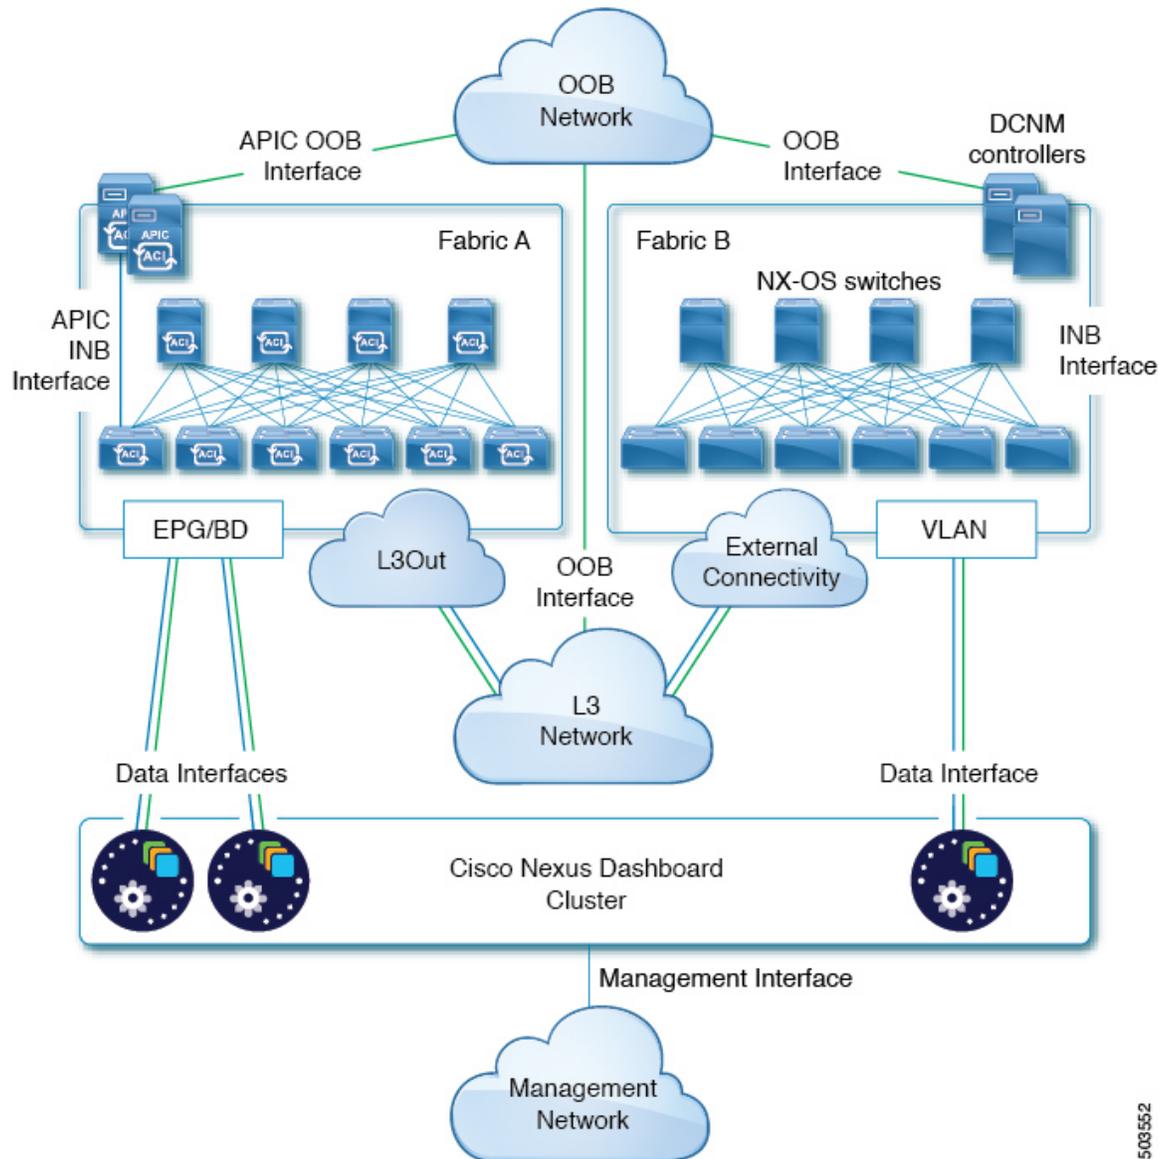
Node Distribution Across Sites

Nexus Dashboard supports distribution of cluster nodes across multiple sites. The following node distribution recommendations apply to both physical and virtual clusters.

Node Distribution for Nexus Dashboard Insights

For Nexus Dashboard Insights, we recommend centralized, single-site deployment. This service does not gain redundancy benefits from distributed cluster, which could instead expose the cluster to interconnection failures when nodes are in different sites.

Node Distribution for Nexus Dashboard Orchestrator

For Nexus Dashboard Orchestrator, we recommend a distributed cluster. Keep in mind that at least two Nexus Dashboard master nodes are required for the cluster to remain operational, so when deploying a Nexus Dashboard cluster across two sites, we recommend deploying a standby node in the site with the single master node as shown in the following figure:

Figure 6: Node Distribution Across Two Sites for Nexus Dashboard Orchestrator
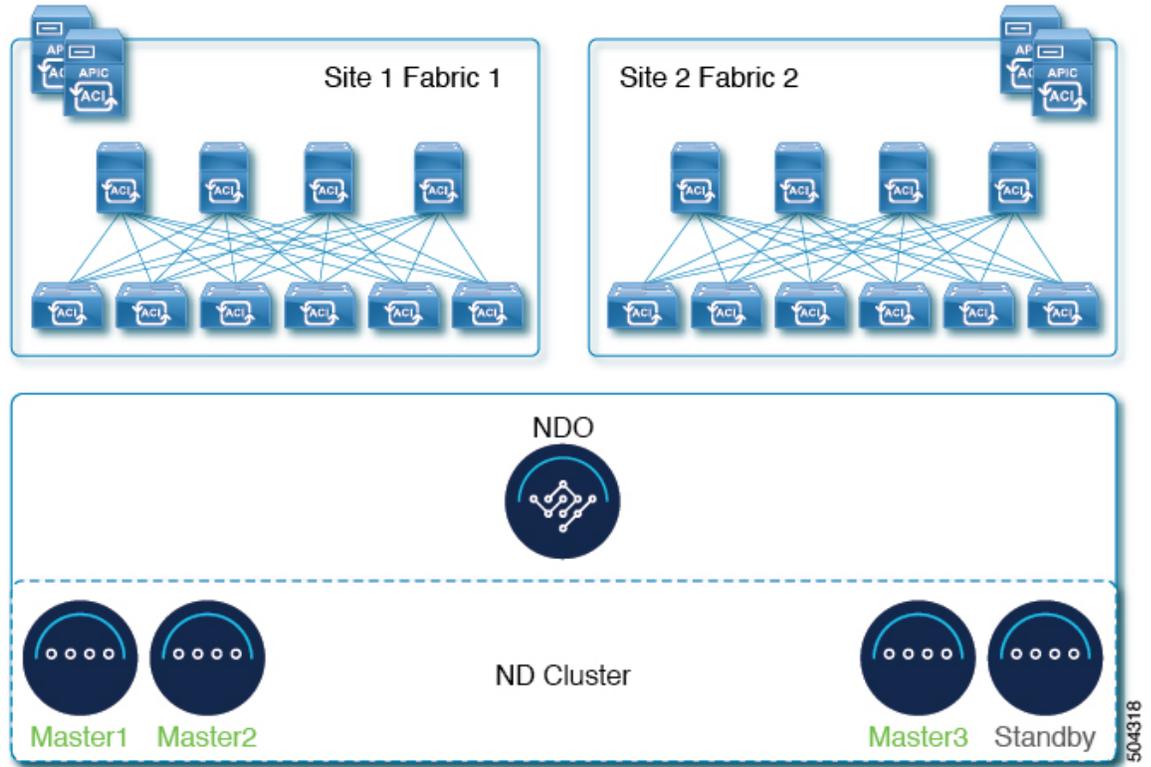
Node Distribution for Fabric Controller

For Nexus Dashboard Fabric Controller, we recommend a centralized, single-site deployment. This service does not support recovery in case if 2 `master` node are not available and thus gains no redundancy benefits from distributed cluster, which could instead expose the cluster to interconnection failures when nodes are in different sites.

Services Co-location Use Cases

This section describes a number of recommended deployment scenarios for specific single-service or multiple services co-hosting use cases.

Note This release does not support co-hosting services in Nexus Dashboard clusters that are deployed in Linux KVM, AWS, or Azure. All services co-hosting scenarios below apply for physical or VMware ESX cluster form factors only.

Single Site, Nexus Dashboard Insights and Orchestrator

In a single site scenario with Nexus Dashboard Insights and Orchestrator services, a single Nexus Dashboard cluster can be deployed with both services co-hosted on it.

Figure 7: Single Site, Nexus Dashboard Insights and Orchestrator
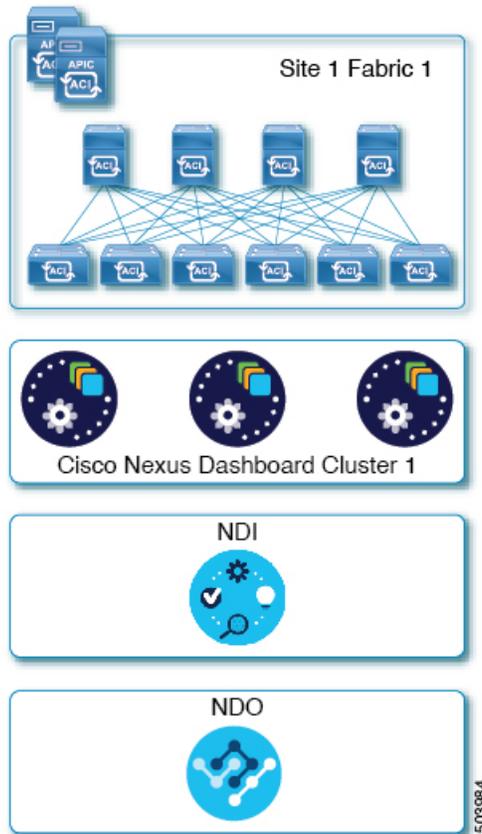
Multiple Sites, Single Cluster for Nexus Dashboard Insights and Orchestrator

In a multiple sites scenario with Nexus Dashboard Insights and Orchestrator services, a single Nexus Dashboard cluster can be deployed with both services co-hosted on it. In this case, the nodes can be distributed between the sites, however since the Insights service does not gain redundancy benefits from a distributed cluster and could instead be exposed to interconnection failures when nodes are in different sites, we recommend the deployment option on the left:

Figure 8: Multiple Sites, Single Cluster for Nexus Dashboard Insights and Orchestrator
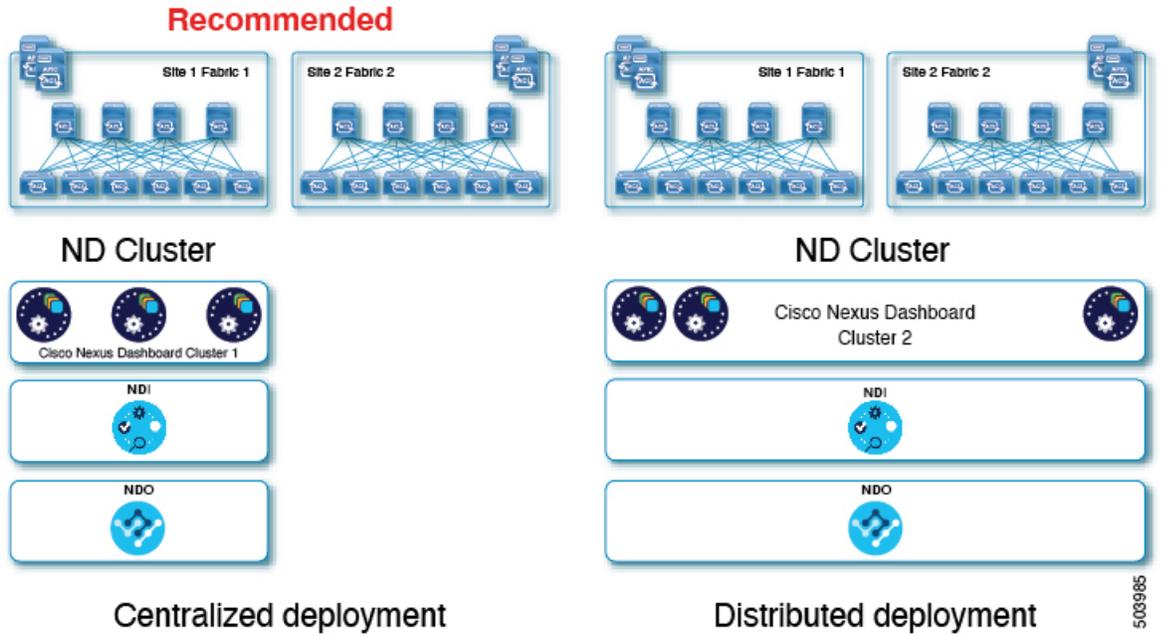
Multiple Sites, Multiple Clusters for Nexus Dashboard Insights and Orchestrator

In this case, we recommend deploying two Nexus Dashboard cluster, with one of them dedicated to the Nexus Dashboard Orchestrator service using the virtual or cloud form factor and the nodes distributed across the sites.

Figure 9: Multiple Sites, Multiple Clusters for Nexus Dashboard Insights and Orchestrator
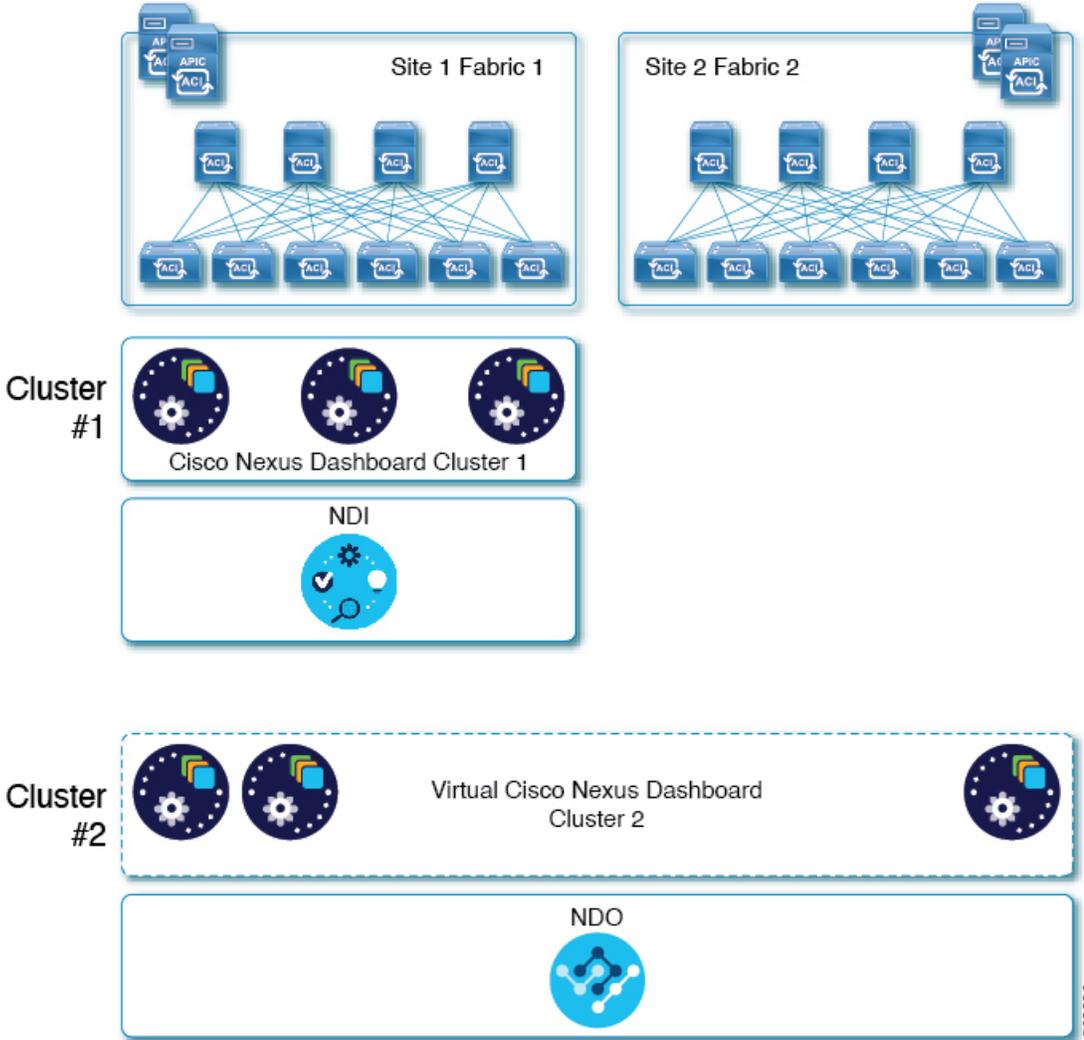
Pre-Installation Checklist

Before you proceed with deploying your Nexus Dashboard cluster, prepare the following information for easy reference during the process:

Table 8: Cluster Details

Parameters	Example	Your Entry
Cluster **Name**	`nd-cluster`	
NTP Server	`171.68.38.65`	
DNS Provider	`64.102.6.247 171.70.168.183`	

Parameters	Example	Your Entry
DNS Search Domain	cisco.com	
App Network	172.17.0.1/16	
Service Network	100.80.0.0/16	

Table 9: Node Details

Parameters	Example	Your Entry
For physical nodes, **CIMC** address and login information of the first node	10.195.219.84/24 Username: admin Password: Cisco1234	
For physical nodes, **CIMC** address and login information of the second node	10.195.219.85/24 Username: admin Password: Cisco1234	
For physical nodes, **CIMC** address and login information of the third node	10.195.219.86/24 Username: admin Password: Cisco1234	
Password used for each node's rescue-user and the initial GUI password. We recommend configuring the same password for all nodes in the cluster.	Welcome2Cisco!	
Management IP of the first node	192.168.9.172/24	
Management Gateway of the first node.	192.168.9.1	
Data Network IP of the first node	192.168.6.172/24	
Data Network Gateway of the first node	192.168.6.1	
(Optional) **Data Network VLAN** of the first node	101	
Management IP of the second node	192.168.9.173/24	
Management Gateway of the second node.	192.168.9.1	

Parameters	Example	Your Entry
Data Network IP of the second node	192.168.6.173/24	
Data Network Gateway of the second node	192.168.6.1	
(Optional) **Data Network VLAN** of the second node	101	
Management IP of the third node	192.168.9.174/24	
Management Gateway of the third node.	192.168.9.1	
Data Network IP of the third node	192.168.6.174/24	
Data Network Gateway of the third node	192.168.6.1	
(Optional) **Data Network VLAN** of the third node	101	

C H A P T E R 3

Deploying as Physical Appliance

Prerequisites and Guidelines

Before you proceed with deploying the Nexus Dashboard cluster, you must:

- Review and complete the general prerequisites described in Deployment Overview and Requirements, on page 3.

  Note that this document describes how to initially deploy the base Nexus Dashboard cluster. If you want to expand an existing cluster with additional nodes (such as `worker` or `standby`), see the "Infrastructure Management" chapter of the *Cisco Nexus Dashboard User Guide* instead, which is available from the Nexus Dashboard UI or online at *Cisco Nexus Dashboard User Guide*

  If you are looking to completely re-image the server, for example in case you cannot log in as the `rescue-user` for manual recovery, see the "Troubleshooting" chapter of the *Cisco Nexus Dashboard User Guide*.

  The guide is available from the Nexus Dashboard UI or online at *Cisco Nexus Dashboard User Guide*

- Review and complete any additional prerequisites described in the *Release Notes* for the services you plan to deploy.

- Ensure you are using the following hardware and the servers are racked and connected as described in *Cisco Nexus Dashboard Hardware Installation Guide*.

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

Table 10: Supported Hardware

PID	Hardware

SE-NODE-G2	• UCS C220 M5 Chassis
	
	• 2x 10-core 2.2G Intel Xeon Silver CPU
	
	• 256 GB of RAM
	
	• 4x 2.4TB HDDs
	
	  400GB SSD
	
	  1.2TB NVME drive
	
	• UCS Virtual Interface Card 1455 (4x25G ports)
	
	• 1050W power supply

SE-CL-L3	A cluster of 3x `SE-NODE-G2` appliances.

Note 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).

  Recommended version: CIMC, Release 4.2(2a).

  Minimum supported version: CIMC, Release 4.0(1a).

- Ensure that all nodes are running the same release version image.

- If your Nexus Dashboard hardware came with a different release image than the one you would like to deploy, we recommend deploying the cluster with the existing image first and then upgrading it to the desired release.

  For example, if the hardware you received came with Release 2.0.1 image pre-installed, but you want to deploy Release 2.1.1 instead, we recommend:

    - First, bring up the Release 2.0.1 cluster, as described in the following section.

    - Then upgrade to Release 2.1.1, as described in Upgrading Nexus Dashboard, on page 79.

You must have at least a 3-node cluster. Additional worker nodes can be added for horizontal scaling if required by the type and number of applications you will deploy. For maximum number of `worker` and `standby` nodes in a single cluster, see the *Release Notes* for your release.

Deploying Cisco Nexus Dashboard as Physical Appliance

When you first receive the Nexus Dashboard physical hardware, it comes preloaded with the software image. This section describes how to configure and bring up the initial 3-node Nexus Dashboard cluster.

Before you begin

- Ensure that you meet the requirements and guidelines described in .

Step 1   Configure the first node's basic information.

You only need to complete the following configuration on one of the nodes of the cluster. For the second and third master nodes, simply ensure that they are powered on, their CIMCs are configured with IP addresses and login credentials, and the CIMC IPs are reachable from the first node.

a) SSH into the node using CIMC management IP and use the `connect host` command to connect to the node's console.

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` CLI 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) 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
```

Step 2   Wait for the initial bootstrap process to complete.

After you provide and confirm management network information, the initial setup configures the networking and brings up the UI, which you will use to add two other nodes and complete the cluster deployment.

```
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.

Step 3   Open your browser and navigate to `https://<first-node-management-ip>` to open the GUI.

The rest of the configuration workflow takes place from the first node's GUI. You do not need to log in to or configure the other two nodes directly.

Enter the password you provided in a previous step and click **Begin Setup**

Step 4   Provide the **Cluster Details**.

In the **Cluster Details** screen of the initial setup wizard, provide the following information:

a) Provide the **Cluster Name** for this Nexus Dashboard cluster.

b) Click +**Add NTP Host** to add one or more NTP servers.

   You must provide an IP address, fully qualified domain name (FQDN) are not supported.

   After you enter the IP address, click the green checkmark icon to save it.

c) Click +**Add DNS Provider** to add one or more DNS servers.

   After you enter the IP address, click the green checkmark icon to save it.

d) Provide a **Proxy Server**.

   For clusters that do not have direct connectivity to Cisco cloud, we recommend configuring a proxy server to establish the connectivity, which will allow you to mitigate risk from exposure to non-conformant hardware and software in your fabrics.

   If you want to skip proxy configuration, click the information (**i**) icon next to the field, then click **Skip**.

e) (Optional) If your proxy server required authentication, change **Authentication required for Proxy** to `Yes` and provide the login credentials.

f) (Optional) Expand the **Advanced Settings** category and change the settings if required.

Under advanced settings, you can configure the following:

- Provide one or more search domains by clicking +**Add DNS Search Domain**.

  After you enter the IP address, click the green checkmark icon to save it.

- Provide custom **App Network** and **Service Network**.

  The application overlay network defines the address space used by the application's services running in the Nexus Dashboard. The field is pre-populated with the default `172.17.0.1/16` value.

  The services network is an internal network used by the Nexus Dashboard and its processes. The field is pre-populated with the default `100.80.0.0/16` value.

  Application and Services networks are described in the Prerequisites and Guidelines, on page 5 section earlier in this document.

g) Click **Next** to continue.

Step 5   In the **Node Details** screen, provide the node's information.

a) Click the **Edit** button next to the first node.

b) Provide the node's **Name**.

c) Provide the node's **Data Network** information.

The **Management Network** information is already pre-populated with the information you provided for the first node.

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.

d) (Optional) Provide IPv6 addresses 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.

Note   If you want to provide IPv6 information, you must do it during cluster bootstrap process. If you deploy the cluster using only IPv4 stack and want to add IPv6 information later, you would need to redeploy the cluster.

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

e) Click **Save** to save the changes.

Step 6   Click **Add Node** to add the second node to the cluster.

The **Node Details** window opens.

a) Provide the node's **Name**.

b) In the **CIMC Details** section, provide the node's CIMC IP address and login credentials, then click **Verify**.

The IP address and login credentials are used to configure that node.

c) Provide the node's **Management Network** information.

You must provide the management network IP address, netmask, and gateway.

d) Provide the node's **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.

e) (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.

Note If you want to provide IPv6 information, you must do it during cluster bootstrap process. If you deploy the cluster using only IPv4 stack and want to add IPv6 information later, you would need to redeploy the cluster.

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

f) Click **Save** to save the changes.

Step 7 Repeat the previous step to add the 3rd node.

Step 8 Click **Next** to continue.

Step 9 In the **Confirmation** screen, review the entered information and click **Configure** to create the cluster.

During the node bootstrap and cluster bring-up, the overall progress as well as each node's individual progress will be displayed in the UI.

It may take up to 30 minutes for the cluster to form and all the services to start. When cluster configuration is complete, the page will reload to the Nexus Dashboard GUI.

Step 10 Verify that the cluster is healthy.

It may take up to 30 minutes for the cluster to form and all the services to start.

After all three nodes are ready, you can log in to any one node via SSH and run the following command to verify cluster health:

a) Verify that the cluster is up and running.

You can check the current status of cluster deployment by logging in to any of the nodes and running the `acs health` command.

While the cluster is converging, you may see the following outputs:

```
$ acs health
k8s install is in-progress
```

```
$ acs health
k8s services not in desired state - [...]
```

```
$ acs health
k8s: Etcd cluster is not ready
```

When the cluster is up and running, the following output will be displayed:

```
$ acs health
All components are healthy
```

b) Log in to the Nexus Dashboard GUI.

After the cluster becomes available, you can access it by browsing to any one of your nodes' management IP addresses. The default password for the `admin` user is the same as the `rescue-user` password you chose for the first node of the Nexus Dashboard cluster.

Step 11 If you plan to host multiple applications in the same cluster, configure deployment profiles for the App Infra Services.

If you plan to host only a single application in your Nexus Dashboard cluster, skip this step.

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

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

Deploying in VMware ESX

Prerequisites and Guidelines

Before you proceed with deploying the Nexus Dashboard cluster in VMware ESX, you must:

- Ensure that the ESX form factor supports your scale and services requirements.

  Scale and services support and co-hosting vary based on the cluster form factor and the specific services you plan to deploy. You can use the Nexus Dashboard Capacity Planning tool to verify that the virtual form factor satisfies your deployment requirements.

  Note Some services (such as Nexus Dashboard Fabric Controller) may require only a single ESX virtual node for one or more specific use cases. In that case, the capacity planning tool will indicate the requirement and you can simply skip the additional node deployment step in the following section.

  However, note that if you have to deploy a mix of App and Data nodes, for example if you plan to deploy Nexus Dashboard Insights or co-host multiple services in the same cluster, you must ensure that the Data nodes are deployed first as the initial cluster's 3 master nodes. Then you can add the App nodes as the `worker` nodes, as described in the *Cisco Nexus Dashboard User Guide*.

- Review and complete the general prerequisites described in Deployment Overview and Requirements, on page 3.

  Note that this document describes how to initially deploy the base Nexus Dashboard cluster. If you want to expand an existing cluster with additional nodes (such as `worker` or `standby`), see the "Infrastructure Management" chapter of the *Cisco Nexus Dashboard User Guide* instead, which is available from the Nexus Dashboard UI or online at *Cisco Nexus Dashboard User Guide*

- Review and complete any additional prerequisites described in the *Release Notes* for the services you plan to deploy.

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

- Data Node—node profile designed for data-intensive applications, such Nexus Dashboard Insights

- App Node—node profile designed for non-data-intensive applications, such Nexus Dashboard Orchestrator

Ensure you have enough system resources:

Table 11: Deployment Requirements

Nexus Dashboard Version	Data Node Requirements	App Node Requirements

Release 2.1.x	• VMware ESXi 6.5, 6.7, or 7.0
	
	• VMware vCenter 6.x, if deploying using vCenter
	
	• Each VM requires:
	
	  • 32 vCPUs
	
	  • 128GB of RAM
	
	  • 3TB SSD storage for the data volume and an additional 50GB for the system volume
	
	  All `Data` nodes must be deployed on SSD or faster storage.
	
	• We recommend that each Nexus Dashboard node is deployed in a different ESXi server.
		• VMware ESXi 6.5, 6.7, or 7.0
		
		• VMware vCenter 6.x, if deploying using vCenter
		
		• Each VM requires:
		
		  • 16 vCPUs
		
		  • 64GB of RAM
		
		  • 500GB HDD or SSD storage for the data volume and an additional 50GB for the system volume
		
		  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.
		
		• We recommend that each Nexus Dashboard node is deployed in a different ESXi server.

- After each node's VM is deployed, ensure that the VMware Tools periodic time synchronization is disabled as described in the deployment procedure in the next section.

- VMware vMotion is not supported for Nexus Dashboard cluster nodes.

- VMware Distributed Resource Scheduler (DRS) is not supported for Nexus Dashboard cluster nodes.

- You can choose to deploy the nodes directly in ESXi or using vCenter.

If you want to deploy using vCenter, following the steps described in Deploying Cisco Nexus Dashboard Using VMware vCenter, on page 36.

If you want to deploy directly in ESXi, following the steps described in Deploying Cisco Nexus Dashboard Directly in VMware ESXi, on page 44.

ESX Host Network Connectivity

If you plan to install Nexus Dashboard Insights or Fabric Controller service and use the Persistent IPs feature, you must ensure that the ESX host where the cluster nodes are deployed has a single logical uplink. In other words, it is connected via a single link, PC, or vPC and not a dual Active/Active (`A/A`) or Active/Standby (`A/S`) link without PC/vPC.

The following diagrams summarize the supported and unsupported network connectivity configurations for the ESX host where the nodes are deployed:

- In case the ESX host is connected directly, the following configurations are supported:

    - `A/A` uplinks of Port-Group or virtual switch with PC or vPC

    - Single uplink of Port-Group or virtual switch

    - Port-Channel used for the uplink of Port-Group or virtual switch.

    `A/A` or `A/S` uplinks of Port-Group or virtual switch without PC or vPC are not supported

    Figure 10: ESX Host Connectivity (Direct)
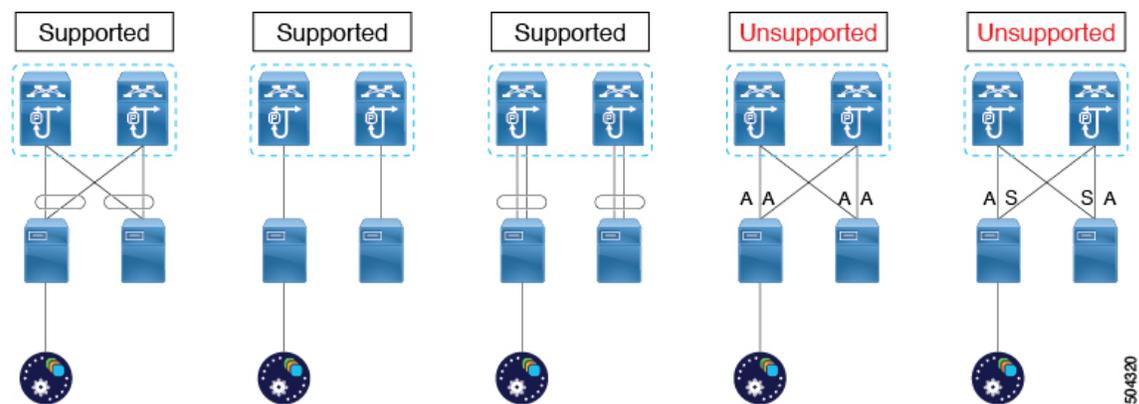
- In case the ESX host is connected via a UCS Fabric Interconnect (or equivalent), the following configurations are supported:

    - `A/S` uplinks of Port-Group or virtual switch at UCS Fabric Interconnect level without PC or vPC

      In this case, the `Active/Standby` links are based on the server technology, such as Fabric Failover for Cisco UCS and not at the ESXi hypervisor level.

    - Single uplink of Port-Group or virtual switch

    `A/A` or `A/S` uplinks of Port-Group or virtual switch at the hypervisor level without PC or vPC are not supported

Figure 11: ESX Host Connectivity (with Fabric Interconnect)
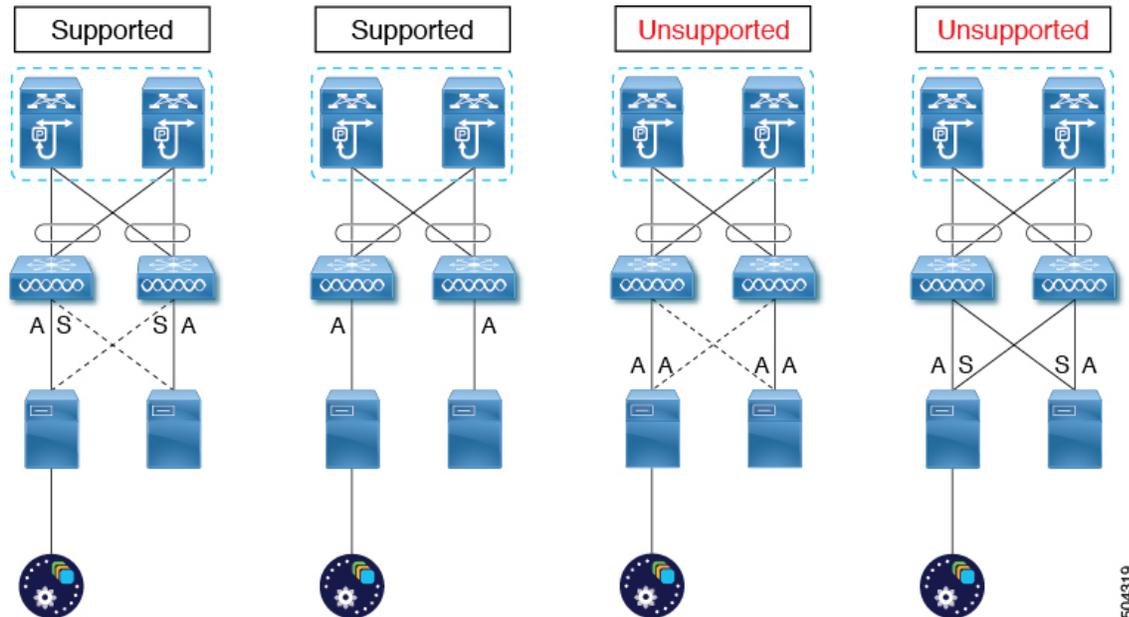
Deploying Cisco Nexus Dashboard Using VMware vCenter

This section describes how to deploy Cisco Nexus Dashboard cluster using VMware vCenter. If you prefer to deploy directly in ESXi, follow the steps described in Deploying Cisco Nexus Dashboard Directly in VMware ESXi, on page 44 instead.

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 33.

Step 1   Obtain the Cisco Nexus Dashboard OVA image.

a) Browse to the Software Download page.

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 appropriate Cisco Nexus Dashboard image (`nd-dk9.<version>.ova`).

For Data nodes, download the `nd-dk9.<version>-data.ova`.

For App nodes, download the `nd-dk9.<version>-app.ova`.

Step 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.

Step 3   Start the new VM deployment.
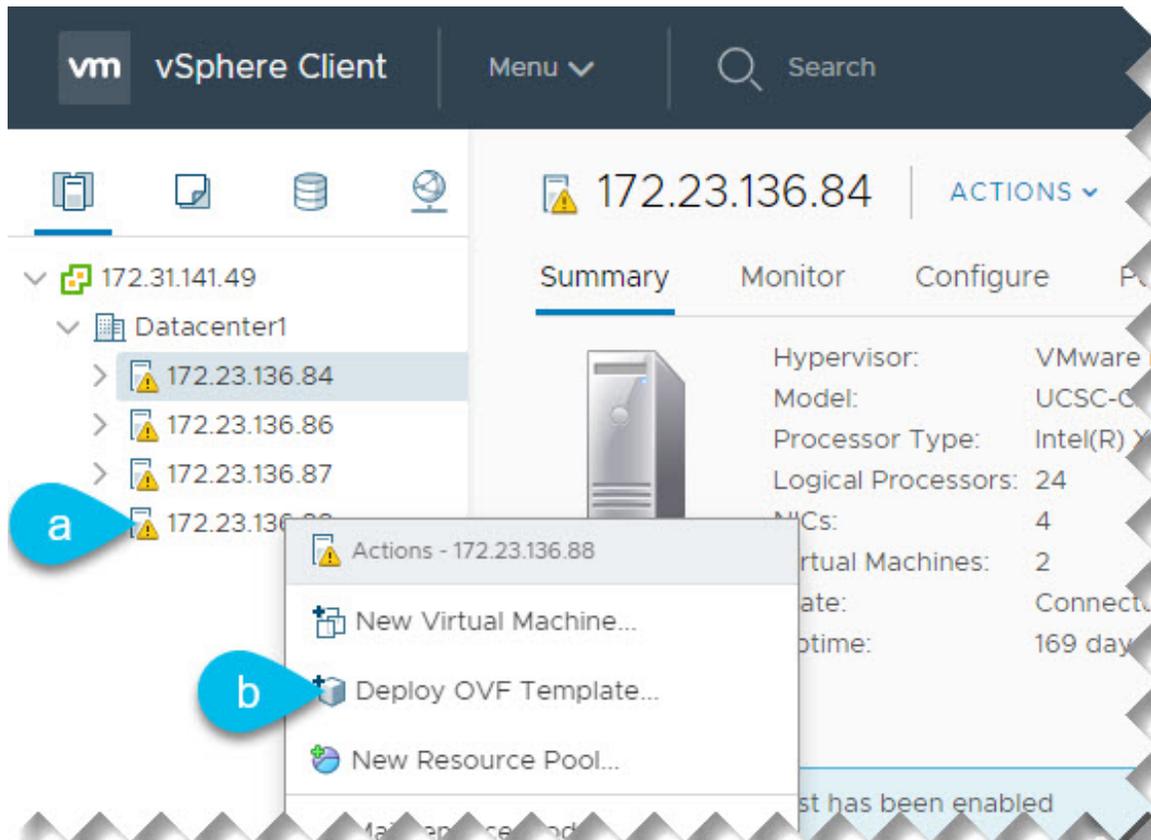
a) Right-click the ESX host where you want to deploy.

b) Then select **Deploy OVF Template...**

The **Deploy OVF Template** wizard appears.

Step 4   In the **Select an OVF template** screen, provide the OVA image.

a) Provide the image.

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

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

b) Click **Next** to continue.

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

a) Provide the name for the virtual machine.
b) Select the location for the virtual machine.
c) Click **Next** to continue

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

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

b) Click **Next** to continue

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

Step 8 In the **Select storage** screen, provide the storage information.

a) From the **Select virtual disk format** dropdown, select **Thick Provision Lazy Zeroed**.

b) Select the datastore for the virtual machine.

We recommend a unique datastore for each node.

c) Click **Next** to continue

Step 9   In the **Select networks** screen, accept default values and click **Next** to continue.

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

Step 10   In the **Customize template** screen, provide the required information.

a) Provide the size for the node's data volume.

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**.

We recommend configuring the same password for all nodes, however you can choose to provide different passwords for the second and third node. If you provide different passwords, the first node's password will be used as the initial password of the `admin` user in the GUI.

c) Provide the **Management Network** IP address, netmask, and gateway.

d) If this is the first node you are deploying, check the **Cluster Leader** checkbox.

Only a single node in the cluster must be leader. You will use the cluster leader's management IP address to complete cluster creation using a GUI wizard in your browser.

e) Click **Next** to continue.

Step 11   In the **Ready to complete** screen, verify that all information is accurate and click **Finish** to begin deploying the first node.

Step 12   Repeat previous steps to deploy the second and third nodes.

You do not need to wait for the first node deployment to complete, you can begin deploying the other two nodes simultaneously.

Note	The steps to deploy the second and third nodes are identical with the only exception being that you must leave the **Cluster Leader** checkbox unchecked.

Step 13   Wait for all three VMs to finish deploying.

Step 14   Ensure that the VMware Tools periodic time synchronization is disabled, then start the VMs.

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.

Step 15   Open your browser and navigate to `https://<first-node-management-ip>` to open the GUI.

The rest of the configuration workflow takes place from the first node's (`Cluster Leader`) GUI. You do not need to log in to or configure the other two nodes directly.

Enter the password you provided in a previous step and click **Begin Setup**
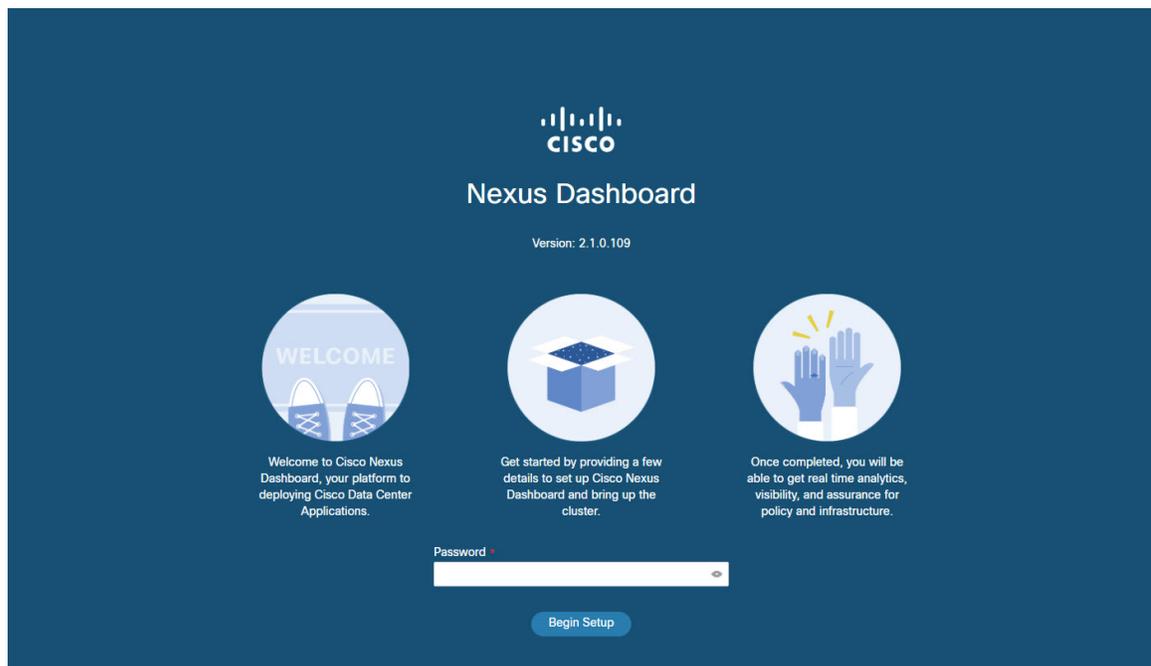
Step 16   Enter the password you provided for the first node and click **Begin Setup**.

Step 17   Provide the **Cluster Details**.

In the **Cluster Details** screen of the initial setup wizard, provide the following information:

a) Provide the **Cluster Name** for this Nexus Dashboard cluster.

b) Click +**Add NTP Host** to add one or more NTP servers.

You must provide an IP address, fully qualified domain name (FQDN) are not supported.

After you enter the IP address, click the green checkmark icon to save it.

c) Click **+Add DNS Provider** to add one or more DNS servers.

After you enter the IP address, click the green checkmark icon to save it.

d) Provide a **Proxy Server**.

For clusters that do not have direct connectivity to Cisco cloud, we recommend configuring a proxy server to establish the connectivity, which will allow you to mitigate risk from exposure to non-conformant hardware and software in your fabrics.

If you want to skip proxy configuration, click the information (**i**) icon next to the field, then click **Skip**.

e) (Optional) If your proxy server required authentication, change **Authentication required for Proxy** to `Yes` and provide the login credentials.

f) (Optional) Expand the **Advanced Settings** category and change the settings if required.

Under advanced settings, you can configure the following:

- Provide one or more search domains by clicking **+Add DNS Search Domain**.

  After you enter the IP address, click the green checkmark icon to save it.

- Provide custom **App Network** and **Service Network**.

  The application overlay network defines the address space used by the application's services running in the Nexus Dashboard. The field is pre-populated with the default `172.17.0.1/16` value.

  The services network is an internal network used by the Nexus Dashboard and its processes. The field is pre-populated with the default `100.80.0.0/16` value.

  Application and Services networks are described in the Prerequisites and Guidelines, on page 5 section earlier in this document.

g) Click **Next** to continue.

Step 18   In the **Node Details** screen, provide the node's information.

a) Click the **Edit** button next to the first node.
b) Provide the node's **Name**.
c) Provide the node's **Data Network** information.

The **Management Network** information is already pre-populated with the information you provided for the first node.

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.

d) (Optional) Provide IPv6 addresses 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.

Note   If you want to provide IPv6 information, you must do it during cluster bootstrap process. If you deploy the cluster using only IPv4 stack and want to add IPv6 information later, you would need to redeploy the cluster.

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

e) Click **Save** to save the changes.

Step 19   Click **Add Node** to add the second node to the cluster.

The **Node Details** window opens.

a) Provide the node's **Name**.

b) In the **Credentials** section, provide the node's **Management Network** IP address and login credentials, then click **Verify**.

The IP address and login credentials are used to pull that node's information.

c) Provide the node's **Data Network** IP address and gateway.

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.

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.

d) (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.

Note	If you want to provide IPv6 information, you must do it during cluster bootstrap process. If you deploy the cluster using only IPv4 stack and want to add IPv6 information later, you would need to redeploy the cluster.
	All nodes in the cluster must be configured with either only IPv4 or dual IPv4/IPv6 stack.

e) Click **Save** to save the changes.

Step 20   Repeat the previous step to add the 3rd node.

Step 21   Click **Next** to continue.

Step 22   In the **Confirmation** screen, review the entered information and click **Configure** to create the cluster.

During the node bootstrap and cluster bring-up, the overall progress as well as each node's individual progress will be displayed in the UI.

It may take up to 30 minutes for the cluster to form and all the services to start. When cluster configuration is complete, the page will reload to the Nexus Dashboard GUI.

Step 23   Verify that the cluster is healthy.

It may take up to 30 minutes for the cluster to form and all the services to start.

After all three nodes are ready, you can log in to any one node via SSH and run the following command to verify cluster health:

a) Verify that the cluster is up and running.

You can check the current status of cluster deployment by logging in to any of the nodes and running the `acs health` command.

While the cluster is converging, you may see the following outputs:

```
$ acs health
k8s install is in-progress

$ acs health
k8s services not in desired state - [...]

$ acs health
k8s: Etcd cluster is not ready
```

When the cluster is up and running, the following output will be displayed:

```
$ acs health
All components are healthy
```

b) Log in to the Nexus Dashboard GUI.

After the cluster becomes available, you can access it by browsing to any one of your nodes' management IP addresses. The default password for the `admin` user is the same as the `rescue-user` password you chose for the first node of the Nexus Dashboard cluster.

Step 24 If you plan to host multiple applications in the same cluster, configure deployment profiles for the App Infra Services.

If you plan to host only a single application in your Nexus Dashboard cluster, skip this step.

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

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

Deploying Cisco Nexus Dashboard Directly in VMware ESXi

This section describes how to deploy Cisco Nexus Dashboard cluster directly in VMware ESXi. If you prefer to deploy using vCenter, follow the steps described in Deploying Cisco Nexus Dashboard Directly in VMware ESXi, on page 44 instead.

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 33.

Step 1 Obtain the Cisco Nexus Dashboard OVA image.

a) Browse to the Software Download page.

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 appropriate Cisco Nexus Dashboard image (`nd-dk9.<version>.ova`).

For Data nodes, download the `nd-dk9.<version>-data.ova`.

For App nodes, download the `nd-dk9.<version>-app.ova`.

Step 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.

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

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

Step 5 In the **Select OVF and VMDK files** screen, provide the virtual machine name (for example, `nd-node1`) and the OVA image you downloaded in the first step, then click **Next**.

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

Step 7 In the **Select OVF and VMDK files** screen, provide the virtual machine name (for example, `nd-node1`) and the OVA image you downloaded in the first step, then click **Next**.

Step 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.

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

Step 10 Repeat previous steps to deploy the second and third nodes.

You do not need to wait for the first node deployment to complete, you can begin deploying the other two nodes simultaneously.

Step 11 Wait for all three VMs to finish deploying.

Step 12 Ensure that the VMware Tools periodic time synchronization is disabled, then start the VMs.

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.

Step 13 Open one of 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 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) For the first node only, designate it as the "Cluster Leader".

You will log into the cluster leader node to finish configuration and complete cluster creation.

```
Is this the cluster leader?: y
```

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, 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
Cluster leader: no

Re-enter config? (y/N): n
```

Step 14   Open your browser and navigate to `https://<first-node-management-ip>` to open the GUI.

The rest of the configuration workflow takes place from the first node's (`Cluster Leader`) GUI. You do not need to log in to or configure the other two nodes directly.

Enter the password you provided in a previous step and click **Begin Setup**
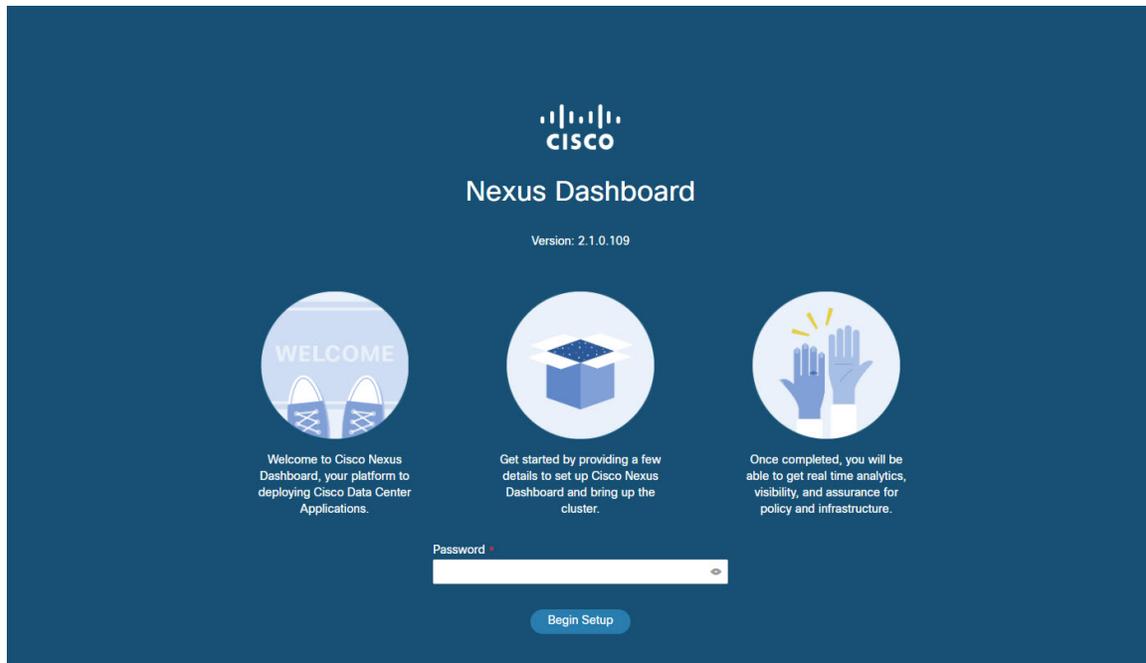
Step 15   Enter the password you provided for the first node and click **Begin Setup**.

Step 16   Provide the **Cluster Details**.

In the **Cluster Details** screen of the initial setup wizard, provide the following information:

a) Provide the **Cluster Name** for this Nexus Dashboard cluster.

b) Click +**Add NTP Host** to add one or more NTP servers.

You must provide an IP address, fully qualified domain name (FQDN) are not supported.

After you enter the IP address, click the green checkmark icon to save it.

c) Click +**Add DNS Provider** to add one or more DNS servers.

After you enter the IP address, click the green checkmark icon to save it.

d) Provide a **Proxy Server**.

For clusters that do not have direct connectivity to Cisco cloud, we recommend configuring a proxy server to establish the connectivity, which will allow you to mitigate risk from exposure to non-conformant hardware and software in your fabrics.

If you want to skip proxy configuration, click the information (**i**) icon next to the field, then click **Skip**.

e) (Optional) If your proxy server required authentication, change **Authentication required for Proxy** to `Yes` and provide the login credentials.

f) (Optional) Expand the **Advanced Settings** category and change the settings if required.

Under advanced settings, you can configure the following:

- Provide one or more search domains by clicking +**Add DNS Search Domain**.

  After you enter the IP address, click the green checkmark icon to save it.

- Provide custom **App Network** and **Service Network**.

  The application overlay network defines the address space used by the application's services running in the Nexus Dashboard. The field is pre-populated with the default `172.17.0.1/16` value.

  The services network is an internal network used by the Nexus Dashboard and its processes. The field is pre-populated with the default `100.80.0.0/16` value.

  Application and Services networks are described in the Prerequisites and Guidelines, on page 5 section earlier in this document.

g) Click **Next** to continue.

Step 17 In the **Node Details** screen, provide the node's information.

a) Click the **Edit** button next to the first node.

b) Provide the node's **Name**.

c) Provide the node's **Data Network** information.

The **Management Network** information is already pre-populated with the information you provided for the first node.

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.

d) (Optional) Provide IPv6 addresses 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.

Note If you want to provide IPv6 information, you must do it during cluster bootstrap process. If you deploy the cluster using only IPv4 stack and want to add IPv6 information later, you would need to redeploy the cluster.

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

e) Click **Save** to save the changes.

Step 18 Click **Add Node** to add the second node to the cluster.

The **Node Details** window opens.

a) Provide the node's **Name**.

b) In the **Credentials** section, provide the node's **Management Network** IP address and login credentials, then click **Verify**.

The IP address and login credentials are used to pull that node's information.

c) Provide the node's **Data Network** IP address and gateway.

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.

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.

d) (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.

Note   If you want to provide IPv6 information, you must do it during cluster bootstrap process. If you deploy the cluster using only IPv4 stack and want to add IPv6 information later, you would need to redeploy the cluster.

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

e) Click **Save** to save the changes.

Step 19   Repeat the previous step to add the 3rd node.

Step 20   Click **Next** to continue.

Step 21   In the **Confirmation** screen, review the entered information and click **Configure** to create the cluster.

During the node bootstrap and cluster bring-up, the overall progress as well as each node's individual progress will be displayed in the UI.

It may take up to 30 minutes for the cluster to form and all the services to start. When cluster configuration is complete, the page will reload to the Nexus Dashboard GUI.

Step 22   Verify that the cluster is healthy.

It may take up to 30 minutes for the cluster to form and all the services to start.

After all three nodes are ready, you can log in to any one node via SSH and run the following command to verify cluster health:

a) Verify that the cluster is up and running.

You can check the current status of cluster deployment by logging in to any of the nodes and running the `acs health` command.

While the cluster is converging, you may see the following outputs:

```
$ acs health
k8s install is in-progress
```

```
$ acs health
k8s services not in desired state - [...]
```

```
$ acs health
k8s: Etcd cluster is not ready
```

When the cluster is up and running, the following output will be displayed:

```
$ acs health
All components are healthy
```

b) Log in to the Nexus Dashboard GUI.

After the cluster becomes available, you can access it by browsing to any one of your nodes' management IP addresses. The default password for the `admin` user is the same as the `rescue-user` password you chose for the first node of the Nexus Dashboard cluster.

Step 23   If you plan to host multiple applications in the same cluster, configure deployment profiles for the App Infra Services.

If you plan to host only a single application in your Nexus Dashboard cluster, skip this step.

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

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

Deploying in Linux KVM

Prerequisites and Guidelines

Before you proceed with deploying the Nexus Dashboard cluster in Linux KVM, you must:

- Ensure that the KVM form factor supports your scale and services requirements.

  Scale and services support and co-hosting vary based on the cluster form factor. You can use the Nexus Dashboard Capacity Planning tool to verify that the virtual form factor satisfies your deployment requirements.

- Review and complete the general prerequisites described in Deployment Overview and Requirements, on page 3.

- Review and complete any additional prerequisites described in the *Release Notes* for the services you plan to deploy.

- Ensure you have enough system resources:

Table 12: Deployment Requirements

Orchestrator Version	Requirements
Release 2.1.x	

Orchestrator Version	Requirements

	• KVM deployments are supported for Nexus Dashboard Fabric Controller and Nexus Dashboard Orchestrator services only.
	
	Specific versions of the required OS and libraries for each service are listed below.
	
	• For Nexus Dashboard Fabric Controller:
	
	    • You must deploy in CentOS 7.9
	
	    • You must have the supported versions of Kernel and KVM:
	
	        • Kernel version
	        `3.10.0-957.el7.x86_64`
	
	        • KVM version
	        `libvirt-4.5.0-23.el7_7.1.x86_64`
	
	• For Nexus Dashboard Fabric Orchestrator:
	
	    • You must deploy in CentOS 7.7
	
	    • You must have the supported versions of Kernel and KVM:
	
	        • Kernel `3.10.0-1062.el7.x86_64`
	
	        • KVM `libvirt 4.5.0`
	
	• 16 vCPUs
	
	• 64 GB of RAM
	
	• 550 GB disk
	
	Each node requires a dedicated disk partition
	
	• The disk must have I/O latency of 20ms or less.
	
	To verify the I/O latency:
	
	1. Create a test directory.
	
	    For example, `test-data`.
	
	2. Run the following command:
	
	    `# fio --rw=write --ioengine=sync --fdatasync=1 --directory=test-data --size=22m --bs=2300 --name=mytest`
	
	3. After the command is executed, confirm that the `99.00th=[<value>]` in the `fsync/fdatasync/sync_file_range` section is under 20ms.

Orchestrator Version	Requirements
	• We recommend that each Nexus Dashboard node is deployed in a different KVM hypervisor.

Deploying Cisco Nexus Dashboard in Linux KVM

This section describes how to deploy Cisco Nexus Dashboard cluster in Linux KVM.

Before you begin

- Ensure that you meet the requirements and guidelines described in .

Step 1 Download the Cisco Nexus Dashboard image.

a) Browse to the Software Download page.

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

b) Click **Nexus Dashboard Software**.
c) From the left sidebar, choose the Nexus Dashboard version you want to download.
d) Download the Cisco Nexus Dashboard image for Linux KVM (nd-dk9.<version>.qcow2).

Step 2 Copy the image to the Linux KVM servers where you will host the nodes.

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

```
# scp nd-dk9.2.1.1a.qcow2 root@<kvm-host-ip>:/home/nd-base
```

The following steps assume you copied the image into the /home/nd-base directory.

Step 3 Create the required disk images for the first node.

You will create a snapshot of the base qcow2 image you downloaded and use the snapshots as the disk images for the nodes' VMs. You will also need to create a second disk image for each 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-dk9.2.1.1a.qcow2 with the location of the base image you created in the previous step.

```
# qemu-img create -f qcow2 -b /home/nd-base/nd-dk9.2.1.1a.qcow2 /home/nd-node1/nd-node1-disk1.qcow2
```

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-node1-disk2.qcow2 500G
```

Step 4   Repeat the previous step to create the disk images for the second and third nodes.

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

- For node1, `/home/nd-node1/` directory with two disk images:

  - `/home/nd-node1/nd-node1-disk1.qcow2`, which is a snapshot of the base `qcow2` image you downloaded in Step 1.

  - `/home/nd-node1/nd-node1-disk2.qcow2`, which is a new 500GB disk you created.

- For node2, `/home/nd-node2/` directory with two disk images:

  - `/home/nd-node2/nd-node2-disk1.qcow2`, which is a snapshot of the base `qcow2` image you downloaded in Step 1.

  - `/home/nd-node2/nd-node2-disk2.qcow2`, which is a new 500GB disk you created.

- For node3, `/home/nd-node3/` directory with two disk images:

  - `/home/nd-node1/nd-node3-disk1.qcow2`, which is a snapshot of the base `qcow2` image you downloaded in Step 1.

  - `/home/nd-node1/nd-node3-disk2.qcow2`, which is a new 500GB disk you created.

Step 5   Create the first node's VM.

a) Open the KVM console and click **New Virtual Machine**.

   You can open the KVM console from the command line using the `virt-manager` command.

b) In the **New VM** screen, choose **Import existing disk image option** and click **Forward**.

c) In the **Provide existing storage path** field, click **Browse** and select the `nd-node1-disk1.qcow2` file.

   We recommend that each node's disk image is stored on its own disk partition.

d) Choose `Generic` for the **OS type** and **Version**, then click **Forward**.

e) Specify 64GB memory and 16 CPUs, then click **Forward**.

f) Enter the **Name** of the virtual machine, for example `nd-node1` and check the **Customize configuration before install** option. Then click **Finish**.

   Note   You must select the **Customize configuration before install** checkbox to be able to make the disk and network card customizations required for the node.

   The VM details window will open.

In the VM details window, change the NIC's device model:

a) Select **NIC <mac>**.

b) For **Device model**, choose `e1000`.

c) For **Network Source**, choose the bridge device and provide the name of the "mgmt" bridge.

In the VM details window, add a second NIC:

a) Click **Add Hardware**.

b) In the **Add New Virtual Hardware** screen, select **Network**.

c) For **Network Source**, choose the bridge device and provide the name of the created "data" bridge.

d) Leave the default **Mac address** value.

e) For **Device model**, choose `e1000`.

In the VM details window, add the second disk image:

a) Click **Add Hardware**.

b) In the **Add New Virtual Hardware** screen, select **Storage**.

c) For the disk's bus driver, choose `IDE`.

d) Select **Select or create custom storage**, click **Manage**, and select the `nd-node1-disk2.qcow2` file you created.

e) Click **Finish** to add the second disk.

Finally, click **Begin Installation** to finish creating the node's VM.

Step 6   Repeat the previous step to create the VMs for the second and third nodes, then start all VMs.

Step 7   Open one of 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 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) For the first node only, designate it as the "Cluster Leader".

You will log into the cluster leader node to finish configuration and complete cluster creation.

```
Is cluster leader?: y
```

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, 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
Cluster leader: no

Re-enter config? (y/N): n
```

Step 8   Repeat previous step to configure the initial information for the second and third nodes.

You do not need to wait for the first node configuration to complete, you can begin configuring the other two nodes simultaneously.

The steps to deploy the second and third nodes are identical with the only exception being that you must indicate that they are not the **Cluster Leader**.

Step 9   Wait for the initial bootstrap process to complete on all nodes.

After you provide and confirm management network information, the initial setup on the first node (`Cluster Leader`) configures the networking and brings up the UI, which you will use to add two other nodes and complete the cluster deployment.

```
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.
```

Step 10   Open your browser and navigate to `https://<first-node-management-ip>` to open the GUI.

The rest of the configuration workflow takes place from the first node's (`Cluster Leader`) GUI. You do not need to log in to or configure the other two nodes directly.

Enter the password you provided in a previous step and click **Begin Setup**
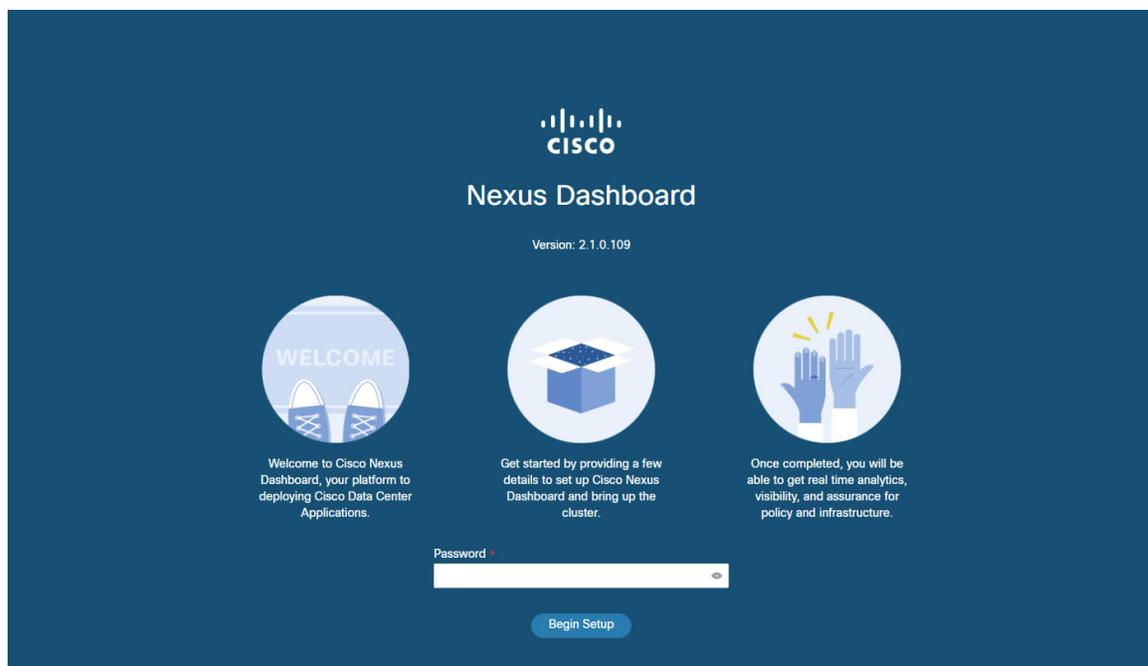
Step 11   Enter the password you provided for the first node and click **Begin Setup**.

Step 12   Provide the **Cluster Details**.

In the **Cluster Details** screen of the initial setup wizard, provide the following information:

a) Provide the **Cluster Name** for this Nexus Dashboard cluster.

b) Click +**Add NTP Host** to add one or more NTP servers.

   You must provide an IP address, fully qualified domain name (FQDN) are not supported.

After you enter the IP address, click the green checkmark icon to save it.

c) Click +**Add DNS Provider** to add one or more DNS servers.

After you enter the IP address, click the green checkmark icon to save it.

d) Provide a **Proxy Server**.

For clusters that do not have direct connectivity to Cisco cloud, we recommend configuring a proxy server to establish the connectivity, which will allow you to mitigate risk from exposure to non-conformant hardware and software in your fabrics.

If you want to skip proxy configuration, click the information (**i**) icon next to the field, then click **Skip**.

e) (Optional) If your proxy server required authentication, change **Authentication required for Proxy** to `Yes` and provide the login credentials.

f) (Optional) Expand the **Advanced Settings** category and change the settings if required.

Under advanced settings, you can configure the following:

- Provide one or more search domains by clicking +**Add DNS Search Domain**.

  After you enter the IP address, click the green checkmark icon to save it.

- Provide custom **App Network** and **Service Network**.

  The application overlay network defines the address space used by the application's services running in the Nexus Dashboard. The field is pre-populated with the default `172.17.0.1/16` value.

  The services network is an internal network used by the Nexus Dashboard and its processes. The field is pre-populated with the default `100.80.0.0/16` value.

  Application and Services networks are described in the Prerequisites and Guidelines, on page 5 section earlier in this document.

g) Click **Next** to continue.

Step 13 In the **Node Details** screen, provide the node's information.

a) Click the **Edit** button next to the first node.
b) Provide the node's **Name**.
c) Provide the node's **Data Network** information.

The **Management Network** information is already pre-populated with the information you provided for the first node.

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.

d) (Optional) Provide IPv6 addresses 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.

Note   If you want to provide IPv6 information, you must do it during cluster bootstrap process. If you deploy the cluster using only IPv4 stack and want to add IPv6 information later, you would need to redeploy the cluster.

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

e) Click **Save** to save the changes.

Step 14   Click **Add Node** to add the second node to the cluster.

The **Node Details** window opens.

a) Provide the node's **Name**.

b) In the **Credentials** section, provide the node's **Management Network** IP address and login credentials, then click **Verify**.

The IP address and login credentials are used to pull that node's information.

c) Provide the node's **Data Network** IP address and gateway.

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.

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.

d) (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.

> **Note**   If you want to provide IPv6 information, you must do it during cluster bootstrap process. If you deploy the cluster using only IPv4 stack and want to add IPv6 information later, you would need to redeploy the cluster.
>
> All nodes in the cluster must be configured with either only IPv4 or dual IPv4/IPv6 stack.

e) Click **Save** to save the changes.

Step 15   Repeat the previous step to add the 3rd node.

Step 16   Click **Next** to continue.

Step 17   In the **Confirmation** screen, review the entered information and click **Configure** to create the cluster.

During the node bootstrap and cluster bring-up, the overall progress as well as each node's individual progress will be displayed in the UI.

It may take up to 30 minutes for the cluster to form and all the services to start. When cluster configuration is complete, the page will reload to the Nexus Dashboard GUI.

Step 18   Verify that the cluster is healthy.

It may take up to 30 minutes for the cluster to form and all the services to start.

After all three nodes are ready, you can log in to any one node via SSH and run the following command to verify cluster health:

a) Verify that the cluster is up and running.

You can check the current status of cluster deployment by logging in to any of the nodes and running the `acs health` command.

While the cluster is converging, you may see the following outputs:

```
$ acs health
k8s install is in-progress

$ acs health
k8s services not in desired state - [...]
```

```
$ acs health
k8s: Etcd cluster is not ready
```

When the cluster is up and running, the following output will be displayed:

```
$ acs health
All components are healthy
```

b) Log in to the Nexus Dashboard GUI.

After the cluster becomes available, you can access it by browsing to any one of your nodes' management IP addresses. The default password for the `admin` user is the same as the `rescue-user` password you chose for the first node of the Nexus Dashboard cluster.

CHAPTER 6

Deploying in Amazon Web Services

- Prerequisites and Guidelines, on page 61
- Deploying the Cisco Nexus Dashboard in AWS, on page 63

Prerequisites and Guidelines

Before you proceed with deploying the Nexus Dashboard cluster in Amazon Web Services (AWS), you must:

- Ensure that the AWS form factor supports your scale and services requirements.

  Scale and services support and co-hosting vary based on the cluster form factor. You can use the Nexus Dashboard Capacity Planning tool to verify that the cloud form factor satisfies your deployment requirements.

- Review and complete the general prerequisites described in the Deployment Overview, on page 3.

- Review and complete any additional prerequisites described in the *Release Notes* for the services you plan to deploy.

- Have appropriate access privileges for your AWS account.

  You must be able to launch multiple instances of Elastic Compute Cloud (`m5.2xlarge`) to host the Nexus Dashboard cluster.

- Have at least 6 AWS Elastic IP addresses.

  A typical Nexus Dashboard deployment consists of 3 nodes with each node requiring 2 AWS Elastic IP addresses for the management and data networks.

  By default, your AWS account has lower elastic IP limit, so you may need to request an increase. To request IP limit increase:

  1. In your AWS console, navigate to **Computer** > **EC2**.

  2. In the EC2 Dashboard, click **Network & Security** > **Elastic IPs** and note how many Elastic IPs are already being used.

  3. In the EC2 Dashboard, click **Limits** and note the maximum number of **EC2-VPC Elastic IPs** allowed.

     Subtract the number of IPs already being used from the limit to get. Then if necessary, click **Request limit increase** to request additional Elastic IPs.

- Create a Virtual Private Cloud (VPC).

A VPC is an isolated portion of the AWS cloud for AWS objects, such as Amazon EC2 instances. To create a VPC:

1. In your AWS console, navigate to **Networking & Content Delivery Tools** > **VPC**.

2. In the VPC Dashboard, click **Your VPCs** and choose **Create VPC**. Then provide the **Name Tag** and **IPv4 CIDR block**.

   The CIDR block is a range of IPv4 addresses for your VPC and must be in the `/16` to `/24` range. For example, `10.9.0.0/16`.

• Create an Internet Gateway and attach it to the VPC.

Internet Gateway is a virtual router that allows a VPC to connect to the Internet. To create an Internet Gateway:

   • In the VPC Dashboard, click **Internet Gateways** and choose **Create internet gateway**. Then provide the **Name Tag**.

   • In the **Internet Gateways** screen, select the Internet Gateway you created, then choose **Actions** > **Attach to VPC**. Finally, from the **Available VPCs** dropdown, select the VPC you created and click **Attach internet gateway**.

• Create a routes table.

Routes table is used for connecting the subnets within your VPC and Internet Gateway to your Nexus Dashboard cluster. To create a routes table:

   • In the VPC Dashboard, click **Route Tables**, choose the **Routes** tab, and click **Edit routes**.

   • In the **Edit routes** screen, click **Add route** and create a `0.0.0.0/0` destination. From the **Target** dropdown, select `Internet Gateway` and choose the gateway you created. Finally, click **Save routes**.

• Create a key pair.

A key pair consists of a private key and a public key, which are used as security credentials to verify your identity when connecting to an EC2 instance. To create a key pair:

   • Navigate to **All services** > **Compute** > **EC2**.

   • In the EC2 Dashboard, click **Network & Security** > **Key Pairs**. Then click **Create Key Pairs**.

   • Provide a name for your key pair, select the **pem** file format, and click **Create key pair**.

     This will download the `.pem` private key file to your system. Move the file to a safe location, you will need to use it the first time you log in to an EC2 instance's console.

✎

Note   By default only PEM-based login is enabled for each node. To be able to SSH into the nodes using a password, as required by the GUI setup wizard, you will need to explicitly enable password-based logins by logging in to each node using the generated key and running the required command as described in the setup section below.

Deploying the Cisco Nexus Dashboard in AWS

This section describes how to deploy Cisco Nexus Dashboard cluster in Amazon Web Services (AWS).

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 61.

Step 1 Subscribe to Cisco Nexus Dashboard product in AWS Marketplace.

a) Log into your AWS account and navigate to the AWS Management Console

The Management Console is available at https://console.aws.amazon.com/.

b) Navigate to **Services** > **AWS Marketplace Subscriptions**.

c) Click **Manage Subscriptions**.

d) Click **Discover products**.

e) Search for **Cisco Nexus Dashboard** and click the result.

f) In the product page, click **Continue to Subscribe**.

g) Click **Accept Terms**.

It may take a couple of minutes for the subscription to be processed.

h) Finally click **Continue to Configuration**.

Step 2 Select software options and region.

a) From the **Delivery Method** dropdown, select `Cisco Nexus Dashboard for Cloud`.

b) From the **Software Version** dropdown, select the version you want to deploy.

c) From the **Region** dropdown, select the regions where the template will be deployed.

This must be the same region where you created your VPC.

d) Click **Continue to Launch**.

The product page appears, which shows a summary of your configuration and enables you to launch the cloud formation template.

Step 3 From the **Choose Action**, select `Launch CloudFormation` and click **Launch**.

The **Create stack** page appears.

Step 4 Create stack.

a) In the **Prerequisite - Prepare template** area, select `Template is ready`.

b) In the **Specify Template** area, select `Amazon S3 URL` for the template source.

The template will be populated automatically.

c) Click **Next** to continue.

The **Specify stack details** page appears.

Step 5 Specify stack details.

    a) Provide the **Stack name**.

    b) From the **VPC identifier** dropdown, select the VPC you created.

        For example, `vpc-038f83026b6a48e98(10.176.176.0/24)`.

    c) In the **ND cluster Subnet block**, provide the VPC subnet CIDR block.

        Choose a subnet from the VPC CIDR that you defined. You can provide a smaller subnet or use the whole CIDR. The CIDR can be a `/24` or `/25` subnet and will be segmented to be used across the availability zones.

        For example, `10.176.176.0/24`.

    d) From the **Availability Zones** dropdown, select one or more available zones.

        We recommend you choose 3 availability zones. For regions that support only 2 availability zones, 2nd and 3rd nodes of the cluster will launch in the second availability zone.

    e) From the **Number of Availability Zones** dropdown, select the number of zones you added in the previous substep.

        Ensure that the number matches the number of availability zones you selected in the previous substep.

    f) Enable **Data Interface EIP support**.

        This field enables external connectivity for the node. External connectivity is required for communication with Cisco ACI fabrics outside AWS.

    g) In the **Password** and **Confirm Password** fields, provide the password.

        This password will be used for the Nexus Dashboard's `rescue-user` login, as well as the initial password for the GUI's `admin` user.

    h) From the **SSH key pair** dropdown, select the key pair you created.

    i) In the **Access control** field, provide the external network allowed to access the cluster.

        For example, `0.0.0.0/0` to be able to access the cluster from anywhere.

    j) Click **Next** to continue.

Step 6   In the **Advanced options** screen, simply click **Next**.

Step 7   In the **Review** screen, verify template configuration and click **Create stack**.

Step 8   Wait for the deployment to complete, then start the VMs.

        You can view the status of the instance deployment in the **CloudFormation** page, for example `CREATE_IN_PROGRESS`. You can click the refresh button in the top right corner of the page to update the status.

        When the status changes to `CREATE_COMPLETE`, you can proceed to the next step.

Step 9   Note down all nodes' public IP addresses.

a) After all instances are deployed, navigate to the AWS console's **EC2** > **Instances** page.

b) Note down which node is labeled as `FirstMaster`.

You will use this node's public IP address to complete cluster configuration.

c) Note down all nodes' public IP addresses.

You will provide this information to the GUI bootstrap wizard in the following steps.

Step 10   Enable password-based login on all nodes.

By default only PEM-based login is enabled for each node. To be able to SSH into the nodes using a password, as required by the GUI setup wizard, you will need to explicitly enable password-based logins.

| **Note** | You must enable password-based login on all nodes before proceeding to cluster bootstrap described in the following steps or you will not be able to complete the cluster configuration. |

a) SSH into one of the instances using its public IP address and the PEM file.

Use the PEM file you created for this as part of Prerequisites and Guidelines, on page 61.

```
# ssh -i <pem-file-name>.pem rescue-user@<node-public-ip>
```

b) Enable password-based login.

On each node, run the following command:

```
# acs login-prompt enable
```

c) Repeat this step for the other two instances.

Step 11   Open your browser and navigate to `https://<first-node-public-ip>` to open the GUI.

| **Note** | You must use the public IP address of the first node (`FirstMaster`) or cluster configuration cannot be completed. |

The rest of the configuration workflow takes place from the first node's GUI. You do not need to log in to or configure the other two nodes directly.

Enter the password you provided during template deployment and click **Begin Setup**
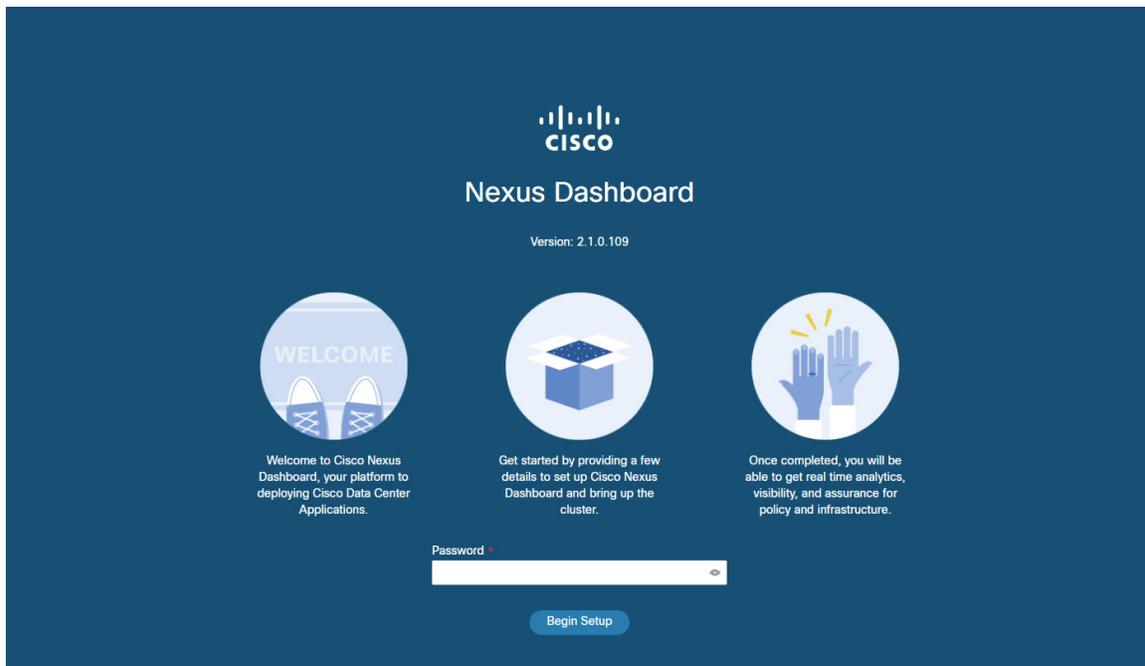
Step 12   Enter the password you provided for the first node and click **Begin Setup**.

Step 13   Provide the **Cluster Details**.

In the **Cluster Details** screen of the initial setup wizard, provide the following information:

a) Provide the **Cluster Name** for this Nexus Dashboard cluster.

b) Click +**Add NTP Host** to add one or more NTP servers.

You must provide an IP address, fully qualified domain name (FQDN) are not supported.

After you enter the IP address, click the green checkmark icon to save it.

c) Click +**Add DNS Provider** to add one or more DNS servers.

After you enter the IP address, click the green checkmark icon to save it.

d) Provide a **Proxy Server**.

For clusters that do not have direct connectivity to Cisco cloud, we recommend configuring a proxy server to establish the connectivity, which will allow you to mitigate risk from exposure to non-conformant hardware and software in your fabrics.

If you want to skip proxy configuration, click the information (**i**) icon next to the field, then click **Skip**.

e) (Optional) If your proxy server required authentication, change **Authentication required for Proxy** to `Yes` and provide the login credentials.

f) (Optional) Expand the **Advanced Settings** category and change the settings if required.

Under advanced settings, you can configure the following:

• Provide one or more search domains by clicking +**Add DNS Search Domain**.

After you enter the IP address, click the green checkmark icon to save it.

• Provide custom **App Network** and **Service Network**.

The application overlay network defines the address space used by the application's services running in the Nexus Dashboard. The field is pre-populated with the default `172.17.0.1/16` value.

The services network is an internal network used by the Nexus Dashboard and its processes. The field is pre-populated with the default `100.80.0.0/16` value.

Application and Services networks are described in the Prerequisites and Guidelines, on page 5 section earlier in this document.

    g) Click **Next** to continue.

Step 14   In the **Node Details** screen, provide the node's information.

    a) Click the **Edit** button next to the first node.

    b) Provide the node's **Name**.

The **Management Network** and **Data Network** information will be already populated from the VNET subnet you have configured before deploying the cluster.

The cluster creates six subnets from the given VPC CIDR, from which the data and management networks will be allocated for the cluster's three nodes.

    c) Leave IPv6 addresses and VLAN fields blank.

Cloud Nexus Dashboard clusters do not support these options.

    d) Click **Save** to save the changes.

Step 15   Click **Add Node** to add the second node to the cluster.

The **Node Details** window opens.

    a) Provide the node's **Name**.

    b) In the **Credentials** section, provide the node's **Public IP Address** and the password you provided during template deployment, then click **Verify**.

The IP address and password are used to pull that node's **Management Network** and **Data Network** information, which will be populated in the fields below.

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

Step 16   Repeat the previous step to add the 3rd node.

Step 17   Click **Next** to continue.

Step 18   In the **Confirmation** screen, review the entered information and click **Configure** to create the cluster.

During the node bootstrap and cluster bring-up, the overall progress as well as each node's individual progress will be displayed in the UI.

It may take up to 30 minutes for the cluster to form and all the services to start. When cluster configuration is complete, the page will reload to the Nexus Dashboard GUI.

Step 19   Verify that the cluster is healthy.

It may take up to 30 minutes for the cluster to form and all the services to start.

After all three nodes are ready, you can log in to any one node via SSH and run the following command to verify cluster health:

    a) Verify that the cluster is up and running.

You can check the current status of cluster deployment by logging in to any of the nodes and running the `acs health` command.

While the cluster is converging, you may see the following outputs:

```
$ acs health
k8s install is in-progress
```

```
$ acs health
k8s services not in desired state - [...]
```

```
$ acs health
k8s: Etcd cluster is not ready
```

When the cluster is up and running, the following output will be displayed:

```
$ acs health
All components are healthy
```

b) Log in to the Nexus Dashboard GUI.

After the cluster becomes available, you can access it by browsing to any one of your nodes' management IP addresses. The default password for the `admin` user is the same as the `rescue-user` password you chose for the first node of the Nexus Dashboard cluster.

Step 20 Verify that the cluster is healthy.

It may take up to 30 minutes for the cluster to form and all the services to start.

After all three nodes are ready, you can log in to any one node via SSH and run the following command to verify cluster health:

a) Verify that the cluster is up and running.

You can check the current status of cluster deployment by logging in to any of the nodes and running the `acs health` command.

While the cluster is converging, you may see the following outputs:

```
$ acs health
k8s install is in-progress
```

```
$ acs health
k8s services not in desired state - [...]
```

```
$ acs health
k8s: Etcd cluster is not ready
```

When the cluster is up and running, the following output will be displayed:

```
$ acs health
All components are healthy
```

b) Log in to the Nexus Dashboard GUI.

After the cluster becomes available, you can access it by browsing to any one of your nodes' management IP addresses. The default password for the `admin` user is the same as the `rescue-user` password you chose for the first node of the Nexus Dashboard cluster.

Deploying in Microsoft Azure

Prerequisites and Guidelines

Before you proceed with deploying the Nexus Dashboard cluster in Microsoft Azure, you must:

- Ensure that the Azure form factor supports your scale and services requirements.

  Scale and services support and co-hosting vary based on the cluster form factor. You can use the Nexus Dashboard Capacity Planning tool to verify that the cloud form factor satisfies your deployment requirements.

- Review and complete the general prerequisites described in the Deployment Overview, on page 3.

- Review and complete any additional prerequisites described in the *Release Notes* for the services you plan to deploy.

- Have appropriate access privileges for your Azure account and subscription.

- Have created a resource group for your Nexus Dashboard cluster resources.

> **Note**　The resource group must be empty and not contain any existing objects. Resource groups with existing objects cannot be used for Nexus Dashboard deployment.

  To create a resource group:

  - In the Azure portal, navigate to **All Resources** > **Resource Groups**.

  - Click +**Add** to create a new resource group.

  - In the **Create a resource group** screen, provide the name of the subscription you will use for your Nexus Dashboard cluster, the name for the resource group (for example, `nd-cluster`), and the region.

- Create an SSH key pair.

A key pair consists of a private key and a public key, you will be asked to provide the public key when creating the Nexus Dashboard nodes.

✎

Note   You will need to use the same machine where you create the public key for a one-time login into each node to enable general SSH login during cluster deployment procedure.

Creating SSH keys is described in Generating an SSH Key Pair in Linux or MacOS, on page 70 and Generating an SSH Key Pair in Windows, on page 71 sections below.

Generating an SSH Key Pair in Linux or MacOS

These procedures describe how to generate an SSH public and private key pair in Linux or MacOS. For instructions on generate an SSH public and private key pair in Windows, see Generating an SSH Key Pair in Windows, on page 71.

Step 1   On your Linux virtual machine or Mac, create a public and private key pair using ssh-keygen, directing the output to a file.

```
# ssh-keygen -f filename
```

For example:

```
# ssh-keygen -f azure_key
```

Output similar to the following appears. Press the Enter key without entering any text when you are asked to enter a passphrase (leave the field empty so that there is no passphrase).

```
Generating public/private rsa key pair.
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in azure_key.
Your public key has been saved in azure_key.pub.
The key fingerprint is:
SHA256:gTsQIIAadjgNsgcguifIloh4XGpVWMdcXVV6U0dyBNs
...
```

Step 2   Locate the public and private key files that you saved.

```
# ls
```

Two files should be displayed, where:

- The file with the .pub suffix contains the public key information

- The file with the same name, but with no suffix, contains the private key information

For example, if you directed the output to a file named azure_key, you should see the following output:

```
# ls
azure_key
azure_key.pub
```

In this case:

- The azure_key.pub file contains the public key information

- The `azure_key` file contains the private key information

Step 3   Open the public key file and copy the public key information from that file, without the `username@hostname` information at the end.

| Note | The private key file is not used in the installation process. However, you might need it for other reasons, such as logging into your Nexus Dashboard nodes through SSH. |

Generating an SSH Key Pair in Windows

These procedures describe how to generate an SSH public and private key pair in Windows. For instructions on generate an SSH public and private key pair in Linux, see .

Step 1   Download and install the PuTTY Key Generator (puttygen):

https://www.puttygen.com/download-putty

Step 2   Run the PuTTY Key Generator by navigating to **Windows** > **Start Menu** > **All Programs** > **PuTTY** > **PuTTYgen**.

You will see a window for the PuTTY Key Generator on your screen.
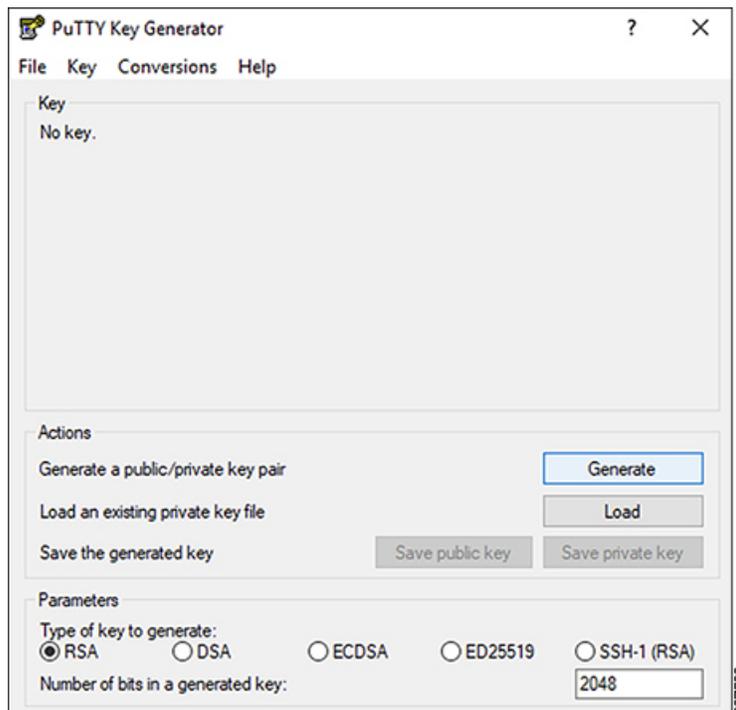
Step 3   Click **Generate**.

A screen appears, asking you to move the mouse over the blank area to generate a public key.

Step 4   Move your cursor around the blank area to generate random characters for a public key.
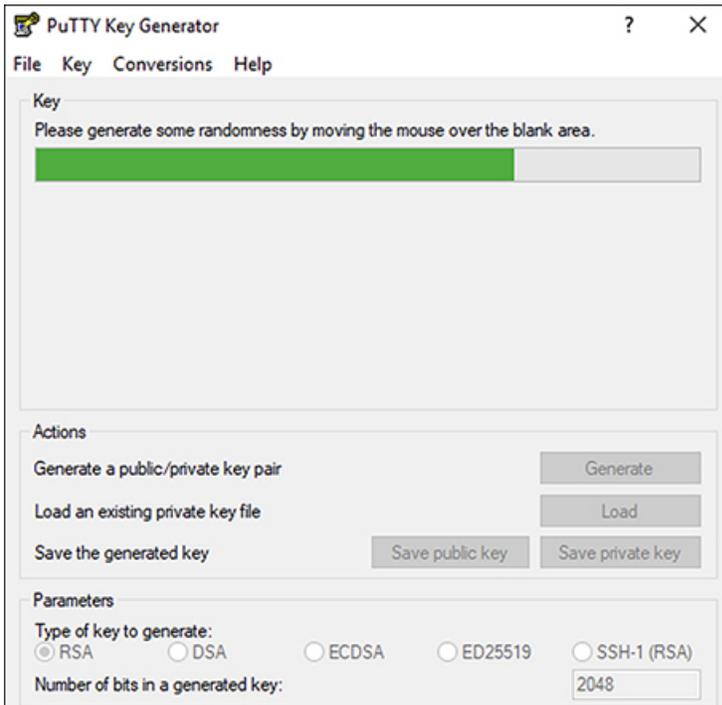
Step 5   Save the public key.

a) Navigate to a folder on your laptop where you want to save the public key file and create a text file for this public key.

b) Copy the information in the PuTTY Key Generator.

Copy the public key information in the window, with these inclusions and exclusions:

- Including the **ssh-rsa** text at the beginning of the public key.

- Excluding the following text string at the end:

  `== rsa-key-<date-stamp>`

Truncate the key so that it does not include the **== rsa-key-<date-stamp>** text string at the end.

Note   In the next set of procedures, you will paste the public key information into the Azure ARM template. If the form does not accept the key in this format, add == back to the end of the key, as this format is required in some regions.

If the key is not in the correct format, the Nexus Dashboard will not complete its installation.
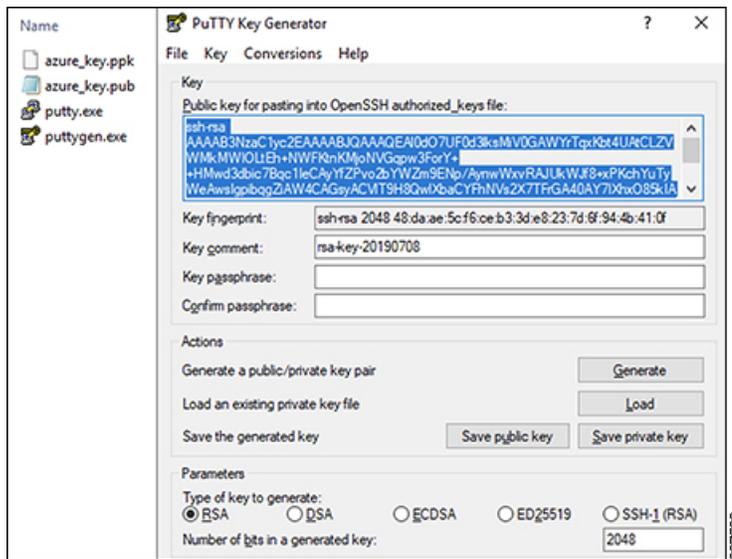
c) Paste the information in the public key text file that you created in 5.a, on page 72 and save the file, giving it a unique file name.

This public key text file will now contain a key that is on a single line of text. You will need the information in this public key text file in the next set of procedures.

Note   Do not save the public key using the **Save public key** option in the PuTTY Key Generator. Doing so saves the key in a format that has multiple lines of text, which is not compatible with the Nexus Dashboard deployment process.

Step 6   Save the private key.

a) Click **Save private key**.

A screen appears, asking if you want to save the file without a passphrase. Click **Yes** on this screen.

b) Navigate to a folder on your laptop and save the private key file, giving it a unique file name.

Note   The private key file is not used in the installation process. However, you might need it for other reasons, such as logging into your Nexus Dashboard nodes through SSH.

Deploying the Cisco Nexus Dashboard in Azure

This section describes how to deploy Cisco Nexus Dashboard cluster in Microsoft Azure.

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 69.

Step 1   Subscribe to Cisco Nexus Dashboard product in Azure Marketplace.

    a) Log into your Azure account and browse to https://azuremarketplace.microsoft.com

    b) In the search field, type `Cisco Nexus Dashboard` and select the option that is presented.

        You will be re-directed to the Nexus Dashboard Azure Marketplace page.

    c) Click **Get it now**.

    d) In the **Select a plan** dropdown, select the version and click **Create**.

Step 2   Provide **Basic** information.

    a) From the **Subscription** dropdown, select the subscription you want to use for this.

    b) From the **Resource group** dropdown, select the resource group you created for this as part of Prerequisites and Guidelines, on page 69.

    c) From the **Region** dropdown, select the region where the template will be deployed.

    d) In the **Password** and **Confirm Password** fields, provide the admin password for the nodes.

        This is the same password that will be used for the `rescue-user` on each node.

    e) In the **SSH public key** field, paste the public key from the key pair you generated as part of the Prerequisites and Guidelines, on page 69 section.

    f) Click **Next** to proceed to the next screen.

Step 3   Provide **ND Settings** information.

    a) Provide the **Cluster Name**.

    b) In the **Image Version** dropdown, verify that the correct version is selected.

    c) In the **Virtual Network Name** field, provide the name for a VNET that will be created for your cluster.

        The VNET must not already exist and will be created for you during deployment. If you provide an already existing VNET, the deployment cannot proceed.

    d) In the **Subnet Address Prefix** field, provide a subnet within the VNET.

        The subnet must be a `/24` subnet and it must be different from the default VNET subnet you defined when creating the VNET.

    e) In the **External Subnets** field, provide the external network allowed to access the cluster.

        For example, `0.0.0.0/0` to be able to access the cluster from anywhere.

    f) Click **Next** to proceed to the next screen.

Step 4   In the **Review + create** page, review information and click **Create** to deploy the cluster.

Step 5   Wait for the deployment to complete, then start the VMs.

Step 6   Note down all nodes' public IP addresses.

    After all instances are deployed, navigate to the Azure console, select each VM, and note down all nodes' public IP addresses. You will provide this information to the GUI bootstrap wizard in the following steps.

    Also note which is the "first" node, which will be indicated by the node's VM name `vm-`**`node1`**`-<cluster-name>`. You will use this node's public IP address to complete cluster configuration.

Step 7   Enable password-based login on all nodes.

    By default only key-based login is enabled for each node. To be able to SSH into the nodes using a password, as required by the GUI setup wizard, you will need to explicitly enable password-based logins.

Note	You must enable password-based login on all nodes before proceeding to cluster bootstrap described in the following steps or you will not be able to complete the cluster configuration.

a) SSH in to one of the nodes as `rescue-user`.

Note	You must use the same machine as you used to create the public key for the deployment during the section.

You can log in as `rescue-user` using the password you provided in template's **Basic** settings:

ssh rescue-user@*<node-public-ip>*

b) Enable password-based login.

acs login-prompt enable

c) Repeat this step for the other two nodes.

Step 8   Open your browser and navigate to `https://<first-node-public-ip>` to open the GUI.

Note	You must use the public IP address of the first node (`vm-`**node1**`-`*<cluster-name>*) or cluster configuration cannot be completed.

The rest of the configuration workflow takes place from the first node's GUI. You do not need to log in to or configure the other two nodes directly.

Enter the password you provided during template deployment and click **Begin Setup**
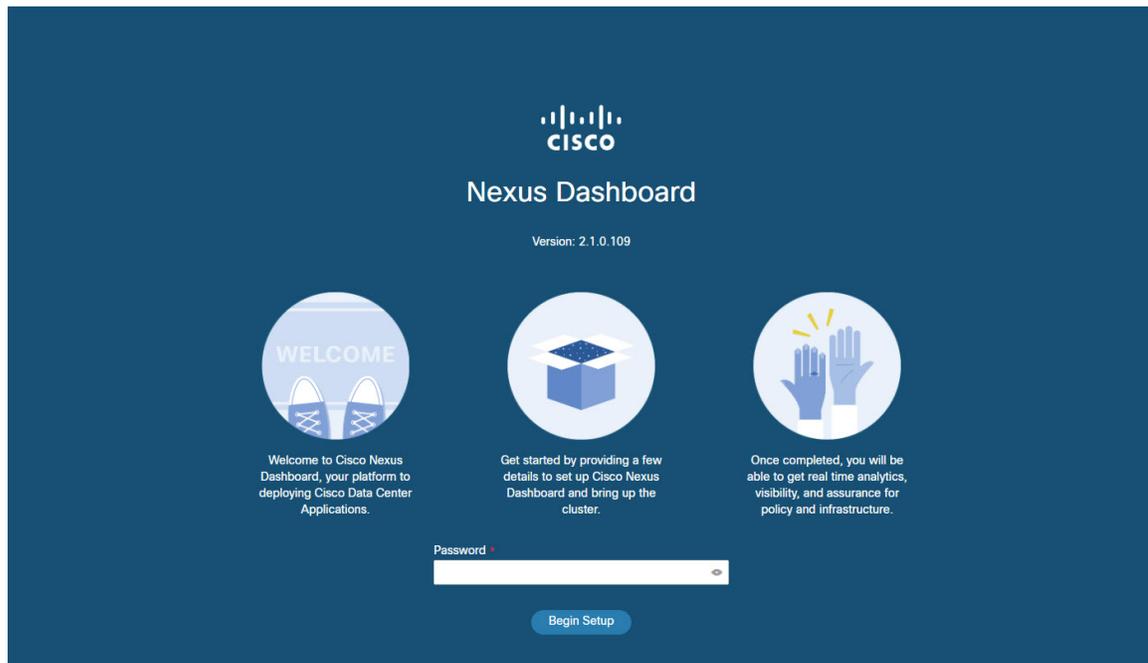
Step 9   Enter the password you provided for the first node and click **Begin Setup**.

Step 10   Provide the **Cluster Details**.

In the **Cluster Details** screen of the initial setup wizard, provide the following information:

a) Provide the **Cluster Name** for this Nexus Dashboard cluster.

b) Click **+Add NTP Host** to add one or more NTP servers.

You must provide an IP address, fully qualified domain name (FQDN) are not supported.

After you enter the IP address, click the green checkmark icon to save it.

c) Click **+Add DNS Provider** to add one or more DNS servers.

After you enter the IP address, click the green checkmark icon to save it.

d) Provide a **Proxy Server**.

For clusters that do not have direct connectivity to Cisco cloud, we recommend configuring a proxy server to establish the connectivity, which will allow you to mitigate risk from exposure to non-conformant hardware and software in your fabrics.

If you want to skip proxy configuration, click the information (**i**) icon next to the field, then click **Skip**.

e) (Optional) If your proxy server required authentication, change **Authentication required for Proxy** to `Yes` and provide the login credentials.

f) (Optional) Expand the **Advanced Settings** category and change the settings if required.

Under advanced settings, you can configure the following:

- Provide one or more search domains by clicking **+Add DNS Search Domain**.

    After you enter the IP address, click the green checkmark icon to save it.

- Provide custom **App Network** and **Service Network**.

    The application overlay network defines the address space used by the application's services running in the Nexus Dashboard. The field is pre-populated with the default `172.17.0.1/16` value.

    The services network is an internal network used by the Nexus Dashboard and its processes. The field is pre-populated with the default `100.80.0.0/16` value.

    Application and Services networks are described in the Prerequisites and Guidelines, on page 5 section earlier in this document.

g) Click **Next** to continue.

Step 11   In the **Node Details** screen, provide the node's information.

a) Click the **Edit** button next to the first node.

b) Provide the node's **Name**.

The **Management Network** and **Data Network** information will be already populated from the VNET subnet you have configured before deploying the cluster.

The cluster creates six subnets from the given VNET, from which the data and management networks will be allocated for the cluster's three nodes.

c) Leave IPv6 addresses and VLAN fields blank.

Cloud Nexus Dashboard clusters do not support these options.

d) Click **Save** to save the changes.

Step 12   Click **Add Node** to add the second node to the cluster.

The **Node Details** window opens.

a) Provide the node's **Name**.

b) In the **Credentials** section, provide the node's **Public IP Address** and the password you provided during template deployment, then click **Verify**.

The IP address and password are used to pull that node's **Management Network** and **Data Network** information, which will be populated in the fields below.

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

Step 13 Repeat the previous step to add the 3rd node.

Step 14 Click **Next** to continue.

Step 15 In the **Confirmation** screen, review the entered information and click **Configure** to create the cluster.

During the node bootstrap and cluster bring-up, the overall progress as well as each node's individual progress will be displayed in the UI.

It may take up to 30 minutes for the cluster to form and all the services to start. When cluster configuration is complete, the page will reload to the Nexus Dashboard GUI.

Step 16 Verify that the cluster is healthy.

It may take up to 30 minutes for the cluster to form and all the services to start.

After all three nodes are ready, you can log in to any one node via SSH and run the following command to verify cluster health:

a) Verify that the cluster is up and running.

You can check the current status of cluster deployment by logging in to any of the nodes and running the `acs health` command.

While the cluster is converging, you may see the following outputs:

```
$ acs health
k8s install is in-progress

$ acs health
k8s services not in desired state - [...]

$ acs health
k8s: Etcd cluster is not ready
```

When the cluster is up and running, the following output will be displayed:

```
$ acs health
All components are healthy
```

b) Log in to the Nexus Dashboard GUI.

After the cluster becomes available, you can access it by browsing to any one of your nodes' management IP addresses. The default password for the `admin` user is the same as the `rescue-user` password you chose for the first node of the Nexus Dashboard cluster.

Upgrading Nexus Dashboard

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. Regardless of whether you deployed your cluster using physical servers, VMware ESX OVA, or in Azure or AWS cloud, you will use the target release's ISO image to upgrade.

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

- You must be running Cisco Nexus Dashboard, Release 2.0.1d or later.

  If you are running Cisco Application Services Engine, you must upgrade to Nexus Dashboard as described in *Cisco Nexus Dashboard Deployment Guide, Release 2.0.x* before upgrading to release 2.1.1. We recommend upgrading your Application Services Engine cluster to Nexus Dashboard release 2.0.2h and then upgrading it to release 2.1.x.

- Ensure that your current Nexus Dashboard cluster is healthy.

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

- Create a backup of the existing configuration.

- If you are running Nexus Insights service, you must disable it before the upgrade and re-enable it after the upgrade completes successfully.

  After you disable the service, ensure the cluster stabilizes and is healthy before you proceed with the upgrade.

Before you re-enable the service, ensure that the App Infra Services deployment profiles are configured appropriate for your deployment, as described in the "App Infra Services" section of the *Cisco Nexus Dashboard User Guide*.

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

- If you are upgrading from Release 2.1.1 or earlier, you may need to clear your browser cache for the new event monitoring page to properly show in the UI.

- After upgrading to this release, we recommend upgrading all the applications to their latest versions.

  For a complete list of Nexus Dashboard and services interoperability support, see the Nexus Dashboard and Services Compatibility Matrix.

- Downgrading from Release 2.1.x is not supported.

Upgrading Nexus Dashboard

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

Before you begin

- Ensure that you have completed the prerequisites described in

Step 1   Download the Nexus Dashboard image.

a)   Browse to the Software Download page.

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

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

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

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

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

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

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

Step 3   Upload the new image to the cluster.
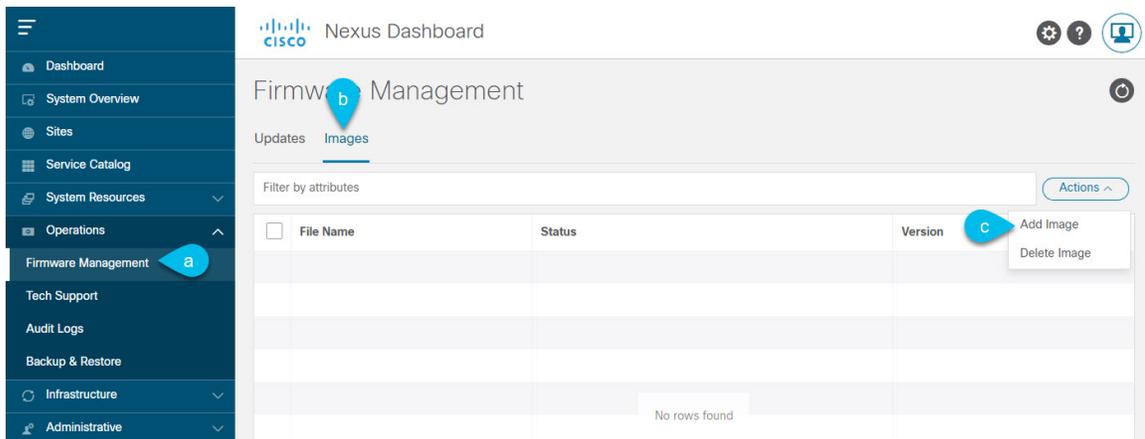
a) Navigate to **Operations** > **Firmware Management**.

b) Select the **Images** tab.

c) From the **Actions** menu, select **Add Image**.

Step 4   Select the new image.

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

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

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

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

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

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

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

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

Step 6   Set up the update.
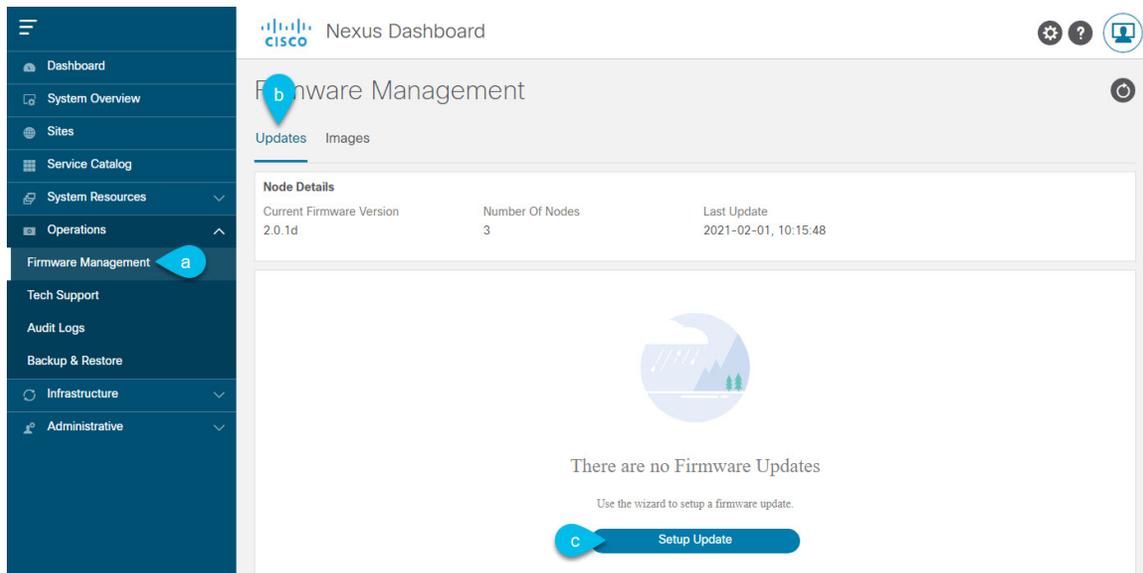
a) Navigate to **Operations** > **Firmware Management**.

b) Select the **Updates** tab.

c) Click **Setup Update**.

The **Firmware Update** screen opens.

Step 7 Choose the upgrade image.

a) In the **Firmware Update** > **Version selection** screen, select the firmware version you uploaded and click **Next**.

b) In the **Firmware Update** > **Confirmation** screen, verify the details and click **Begin Install**.

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

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

Step 8 Activate the new image.

a) Navigate back to the **Operations** > **Firmware Management** screen

b) In the **Last Update Status** tile, click **View Details**.

c) Click **Activate**.

d) In the **Activation Confirmation** window, click **Continue**.

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

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

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

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

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

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

The profile you choose depends on your use case scenario:

- If you plan to run only the Nexus Dashboard Orchestrator service, convert all nodes to the `App` node profile.

- If you plan to run Nexus Dashboard Insights or co-host applications, you must convert the nodes to the `Data` profile.

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

a) Bring down one of the nodes.

You must replace one node at a time.

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

When deploying the new node, you must use the same exact network configuration parameters as the node you are replacing. You must also ensure that the **Cluster Leader** checkbox in the OVF parameters is left unchecked.

c) Log in to the existing Nexus Dashboard GUI.

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

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

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

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

The **Replace** window will open.

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

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

g) Click **Replace**.

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

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

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

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

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

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