Cisco Nexus Dashboard Fabric Controller Installation and Upgrade Guide, Release 12.1.1p

First Published: 2022-08-15

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

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

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

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

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

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

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

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

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

CONTENTS

CHAPTER 1

Overview

Overview

Note Cisco Data Center Network Manager (DCNM) is renamed as Cisco Nexus Dashboard Fabric Controller (NDFC) from Release 12.0.1a.

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

Nexus Dashboard Fabric Controller primarily focuses on Control and Management for three primary market segments:

- LAN networking including VXLAN, Multi-Site, Classic Ethernet, and External Fabrics supporting Cisco Nexus switches running standalone NX-OS, with additional support for IOS-XR, IOS-XE, and adjacent Host, Compute, Virtual Machine, and Container Management systems.

- SAN networking for Cisco MDS and Cisco Nexus switches running standalone NX-OS, including support for integration with storage arrays and additionally Host, Compute, Virtual Machine, and Container Orchestration systems.

- Media Control for Multicast Video production networks running Cisco Nexus switches operated as standalone NX-OS, with additional integrations for 3rd party media control systems.

Previously, DCNM was an application server running on a VM deployed via OVA or ISO, a physical appliance deployed via ISO, or software installed on a qualified Windows or Linux machine. Cisco Nexus Dashboard Fabric Controller, Release 12 is available as an application running exclusively on top of the Cisco Nexus Dashboard Virtual or Physical Appliance.

Virtual Nexus Dashboard deployment with OVA is also referred to as virtual Nexus Dashboard (vND) deployment, while the deployment of Nexus Dashboard on physical appliance (Service Engine) is known as physical Nexus Dashboard (pND) deployment. To deploy Nexus Dashboard based on your requirement, refer to .

Beginning with Release 12, Cisco Nexus Dashboard Fabric Controller has a single installation mode. Post installation, it supports selection from multiple personas at run-time. After the Nexus Dashboard Fabric Controller Release 12.1.1p is installed, you can choose from one of the following personas:

- **Fabric Discovery**—Discover, Monitor, and Visualize LAN Deployments.
- **Fabric Controller**—LAN Controller for Classic Ethernet (vPC), Routed, VXLAN, and IP Fabric for Media Deployments.
- **SAN Controller**—SAN Controller for MDS and Nexus switches. Enhanced SAN Analytics with streaming telemetry.

Note For any given instance of Nexus Dashboard, only one version of NDFC service will be active. On the active NDFC service, you can configure only one persona at any given instance.

All features/services are modularized, broken into smaller microservices, and the required microservices are orchestrated based on the feature set or feature selections. Therefore, if any feature or microservice is down, only that microservice is restarted and recovered, resulting in minimal disruption.

In contrast to the previous DCNM Active-Standby HA model, Cisco NDFC introduces Active-Active HA deployment model utilizing all three nodes in a cluster for deploying microservices. This has significant improvement in both latency and effective resource utilization.

Note For NDFC to run on top of the virtual Nexus Dashboard (vND) instance, you must enable promiscuous mode on port groups that are associated with Nexus Dashboard interfaces where External Service IP addresses are specified. vND comprises of Nexus Dashboard management interface and data interface. By default, for LAN deployments, 2 external service IP addresses are required for the Nexus Dashboard management interface subnet. Therefore, you must enable promiscuous mode for the associated port-group. If inband management or Endpoint Locator (EPL) is enabled, you must specify External Service IP addresses in the Nexus Dashboard data interface subnet. You must also enable the promiscuous mode for the Nexus Dashboard data/fabric interface port-group. For NDFC SAN Controller, promiscuous mode must be enabled only on the Nexus Dashboard data interface associated port-group. For NDFC SAN Controller, promiscuous mode only needs to be enabled on the Nexus Dashboard data interface associated port-group. For more information, refer to .

Cisco NDFC Release 12.1.1p supports hybrid cloud connectivity between on-prem and public cloud networks. Using Cisco Nexus Dashboard Orchestrator, connectivity is orchestrated between NDFC managed VXLAN fabric and Cloud Application Policy Infrastructure Controller (cAPIC) deployed on public cloud.

Note Only fresh installation of NDFC Release 12.1.1p is supported. You cannot upgrade to Release 12.1.1p from older NDFC or DCNM releases.

For more information, see Cisco Nexus Dashboard Fabric Controller (Formerly DCNM).

Change History

The following table shows the change history for this document.

Table 1: Change History

Date	Description
15 August 2022	Release 12.1.1p became available.

Deployment Options

The following deployment options are available for Cisco Nexus Dashboard Fabric Controller:

- NDFC on Single node (non-HA Cluster)

  On Single node Nexus Dashboard, you can deploy NDFC with the following personas:

  - Fabric Discovery for lab/non-production environments (<= 25 switches)
  - Fabric Controller for lab/non-production environments (<= 25 switches)
  - Fabric Controller in IP Fabric for Media controller mode for production environments
  - SAN Controller for production environments (<= 80 switches)

  Note Fabric Controller/Fabric Discovery deployment is for Lab purposes only. Do not deploy this in your production environment.

- NDFC on a 3-node Cluster (Active-Active HA mode)

  On 3-Node Nexus Dashboard, you can deploy NDFC with the following personas:

  - Fabric Discovery
  - Fabric Controller
  - SAN Controller with or without SAN Insights

- NDFC on a 5-node virtual Nexus Dashboard (vND) Cluster (Active-Active HA mode)

  On 5-Node Nexus Dashboard, you can deploy NDFC with the following personas:

  - Fabric Discovery
  - Fabric Controller

- NDFC on a 3-node/4-node/5-node physical Nexus Dashboard (pND) Cluster (Active-Active HA mode)

  On a 4-node or 5-node Nexus Dashboard, you can deploy Nexus Dashboard Insights (NDI) along with NDFC with the following personas:

  - Nexus Dashboard Insights and NDFC in Fabric Discovery persona (NDFC-Monitored mode) – 4 pND nodes
  - Nexus Dashboard Insights and NDFC in Fabric Controller persona (NDFC-Managed mode) – 5 pND nodes

- NDFC on a Nexus Dashboard running on top of Red Hat Enterprise Linux (RHEL)

  From Release 12.1.1e, on a 1-node or 3-node Nexus Dashboard on the RHEL server, you can deploy NDFC with the following personas:

  - SAN Controller with or without SAN Insights

- NDFC on a virtual Nexus Dashboard (vND) with KVM hypervisor

  From Release 12.1.1e, on a virtual Nexus Dashboard with KVM hypervisor, you can deploy NDFC with the following personas:

  Note You must create bridge interfaces on Linux before installing Nexus Dashboard on KVM with Centos7. Ensure that you use bridge interfaces and do not allow other interfaces during Nexus Dashboard installation.

  - Supports Fabric Controller, Fabric Discovery, and SAN Controller personas.

Refer to Nexus Dashboard Capacity Planning to determine the number of switches supported for each deployment.

In the 3-node and 5-node deployment, there are 3 Nexus Dashboard master nodes. In the 5-node deployment, the additional 2 nodes serve as worker nodes. The 3-node or 5-node cluster deployment is an active-active solution, that is, all nodes are utilized to run micro-services of Nexus Dashboard Fabric Controller. When a node fails, microservices running on the node, are moved to the other nodes. Nexus Dashboard Fabric Controller functions normally in a one-node failure scenario. However, it is expected that there will be a brief disruption to services that must be migrated on node failure. After the migration of services is complete, the supported scale will continue to be supported albeit at degraded performance. To restore optimal NDFC performance, a system running with one failed node is not the desired situation and must be rectified at the earliest. A 3-node or 5-node cluster cannot tolerate the failure of two Master nodes or all NDFC services will be disrupted.

For virtual Nexus Dashboard (vND) OVA deployments on ESXi environments, it is imperative that promiscuous mode is enabled on the port groups that are associated with Nexus Dashboard management and Nexus Dashboard data/fabric interfaces. Otherwise, some of the functionality such as SNMP trap, Image management, Endpoint Locator, SAN Insights, and so on, will not work.

Note that promiscuous mode settings are not required for the port group associated with the Data interface for Layer-3 adjacent network.

Note Nexus Dashboard cluster federation is not supported with Nexus Dashboard Fabric Controller.

Deployment Profile Simplification

Nexus Dashboard deployment profile simplification is intended to help streamline the onboarding of services against a given deployment scale and relieve the task of remembering the cross-connect of deployments.

Beginning with Cisco Nexus Dashboard Release 2.2.1h, resource profile selection has been reduced to several more intuitive parameters directly related to your deployment use case. These parameters, such as number of switches or flows describe the fabric size and use case intent, and allow the cluster to intelligently determine the resources needed for the service. The parameters are categorized as **Network Scale**.

NDFC selects an appropriate profile from among the predefined set of profiles to match the scale.

Note You must restart the services on the Nexus Dashboard after modifying the network scale parameters.

To view or modify the Network Scale parameters on Cisco Nexus Dashboard, perform the following steps:

1. Choose **Nexus Dashboard > Cluster Configuration > Network Scale**.
2. Click the edit icon to modify the network scale parameters.
3. In the **Number of Sites** field, provide the target number of sites for your deployment that this Nexus Dashboard cluster will manage.
4. In the **Number of Switches** field, provide the target number of switch nodes for your deployment.
5. In the **Flows per second** field, provide the target number of flows across sites for LAN/IPFM/SAN-Insights deployments or scale supported by NDFC/NDI cohosted setup.

From Release 12.1.1e, NDFC deployment profiles use a different naming convention for these deployment profiles which is more in line with the scale numbers that each profile supports.

On the fresh install of Nexus Dashboard, the **Network Scale** is empty. We recommend that you define the number of sites, switches, and flows per second in the Network Scale. In such a scenario, the service selects a default profile based on the number of cluster nodes.

If the available cluster compute capacity is less than the desired **Network Scale**, Cisco NDFC installation displays an error. You must resolve the network scale values on Nexus Dashboard and proceed to install NDFC. Note that the recommendations specified in the error message provide useful suggestions about remedial action.

Nexus Dashboard assigns profile names for supported scale values with NDFC.

Layer 3 Reachability Between Cluster Nodes

From Release 12.1.1e, NDFC can be deployed as a service on Nexus Dashboard with Layer 3 adjacent nodes. A sample NDFC Layer 3 adjacent Physical Connectivity topology is as shown in the following image.

When using Layer 3 adjacency between the Nexus Dashboard nodes on which the NDFC service is running, the persistent IP addresses are advertised using the Nexus Dashboard Data or Fabric interface. The Layer 3 Persistent IP subnet pool must be unique and will be advertised to the fabric using BGP on Nexus Dashboard. Cisco NDFC pods, such as EPL/SNMP Trap/SCP that requires Persistent IPs, are advertised as /32 BGP entries with the next hop of Nexus Dashboard Data Interface. Also, the BGP session between the Nexus Dashboard node and the uplink switches must be configured using directly connected links.

For information about persistent IP addresses, see Persistent IP Requirements for NDFC.

To deploy Layer 3 cluster connectivity, Nexus Dashboard nodes use BGP local and remote autonomous system configuration, along with Data Network gateway of the node to establish eBGP sessions with neighboring routers over the Data interface. As Nexus Dashboard nodes use gateway IPs to establish sessions, during Nexus Dashboard cluster configuration, the neighboring BGP peers must be Layer 2 adjacent. Peers without Layer 2 adjacent connectivity are not supported. You must configure the BGP network correctly to ensure that the Nexus Dashboard routes are transmitted correctly.

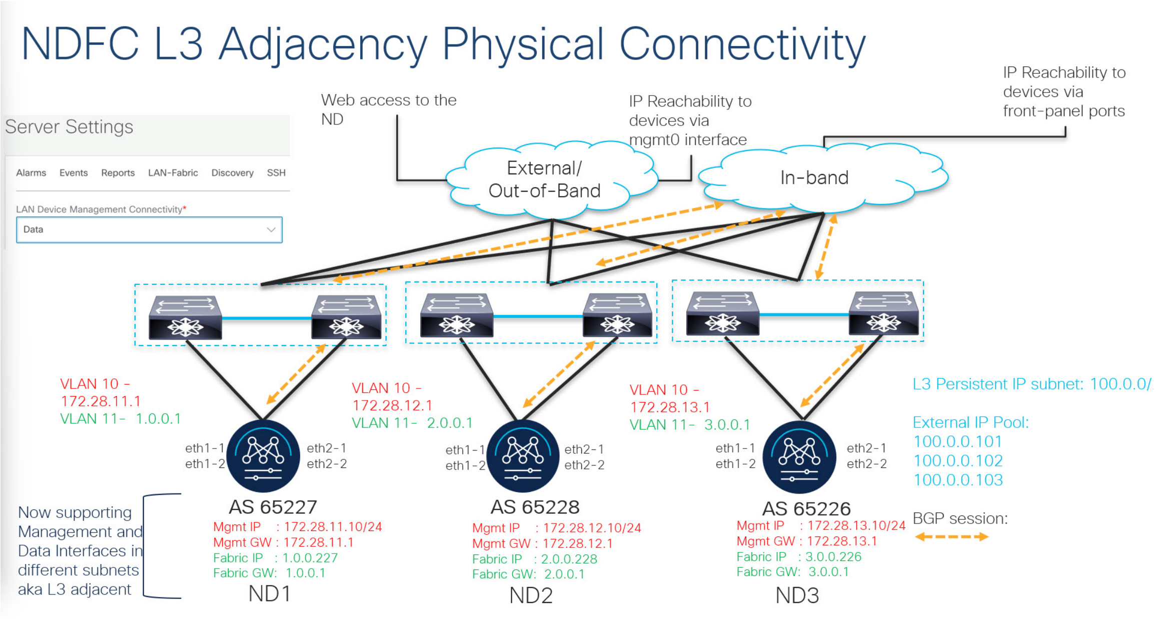

Upgrade or modification from an existing Layer-2 adjacent Nexus Dashboard cluster to a Layer-3 adjacent cluster is not supported. When using Layer 3 adjacency, NDFC service is supported only when the switch connectivity is through the Nexus Dashboard Data interface. Choose NDFC **UI** > **Settings** > **Admin** tab. From the **LAN Device Management Connectivity** drop-down list, select **Data**.

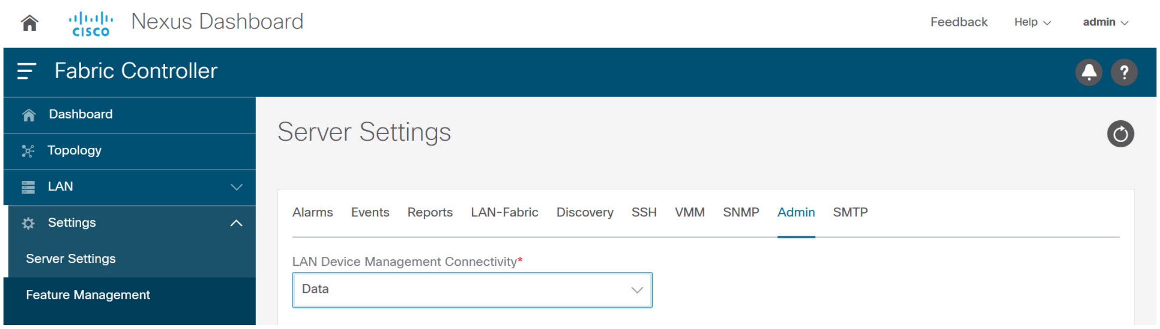

LAN Device Management Connectivity*

Data

Nexus Dashboard uses eBGP to publish up-to-date reachability of /32 routes for reaching NDFC features using external service IPs obtained from the Persistent IP subnet. If a node or network fails, the external IPs are not reachable until recovery is complete (if the network can recover itself). After the microservices on the failed node are brought up on one of the existing nodes on the cluster, the eBGP peering from that node will automatically advertise the corresponding /32 persistent IP reachability to the rest of the network, by that means, autorepairing the service disruption.

The following table provides information about different scenarios about Layer 3 adjacent cluster nodes connectivity.

Network details	Support provided
Modify or upgrade from Layer 2 adjacency to Layer 3 adjacency	Not supported; the cluster must be redeployed if necessary.
Modify or upgrade from Layer 3 adjacency to Layer 2 adjacency	Not supported; the cluster must be redeployed if necessary.
NDFC to Switch connectivity over the management interface	Supported (The traffic initiated by the switch to NDFC is routed via the Data Interface)
NDFC to Switch connectivity over Data interface	Supported
Nexus Dashboard BGP traffic over the management interface	Not supported
Cisco Nexus Dashboard BGP traffic over Data interface	Supported
Nexus Dashboard BGP peer L2-Adjacent	Supported
Nexus Dashboard BGP peer L3-Adjacent	Not supported

See Cisco Nexus Dashboard User Guide, Release 2.2.x for more information.

Appendix

The following images show different NDFC connectivity

NDFC Connectivity - I

LAN

Device reachability from NDFC, for both OOB and Inband device access, is via ND Data interface

- ND Management interface used for external web access interface only

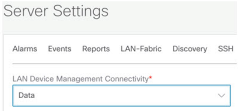

Server Settings

LAN Device Management Connectivity*: Data

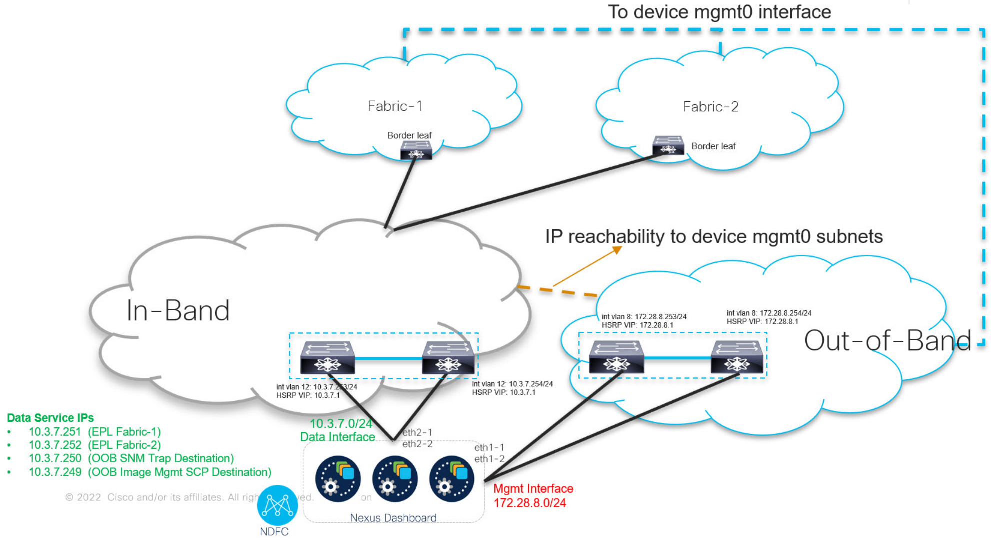

Data Service IPs

- 10.3.7.251 (EPL Fabric-1)
- 10.3.7.252 (EPL Fabric-2)
- 10.3.7.250 (OOB SNM Trap Destination)
- 10.3.7.249 (OOB Image Mgmt SCP Destination)

NDFC Connectivity - II

LAN

Device reachability from NDFC, for both OOB and Inband device access, is via ND Data interface

- ND Management interface used for external web access interface only

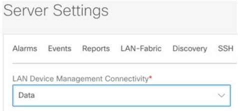

Server Settings

LAN Device Management Connectivity*: Data

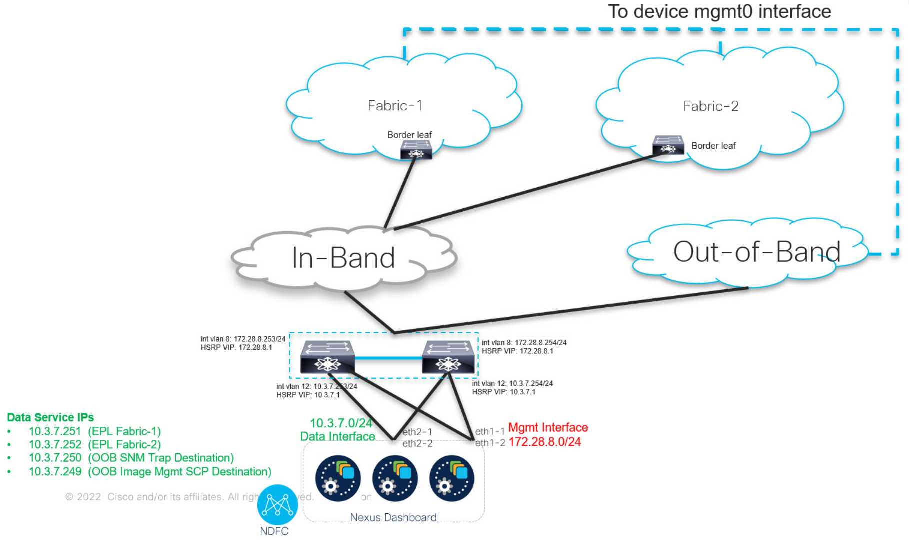

Data Service IPs

- 10.3.7.251 (EPL Fabric-1)
- 10.3.7.252 (EPL Fabric-2)
- 10.3.7.250 (OOB SNM Trap Destination)
- 10.3.7.249 (OOB Image Mgmt SCP Destination)

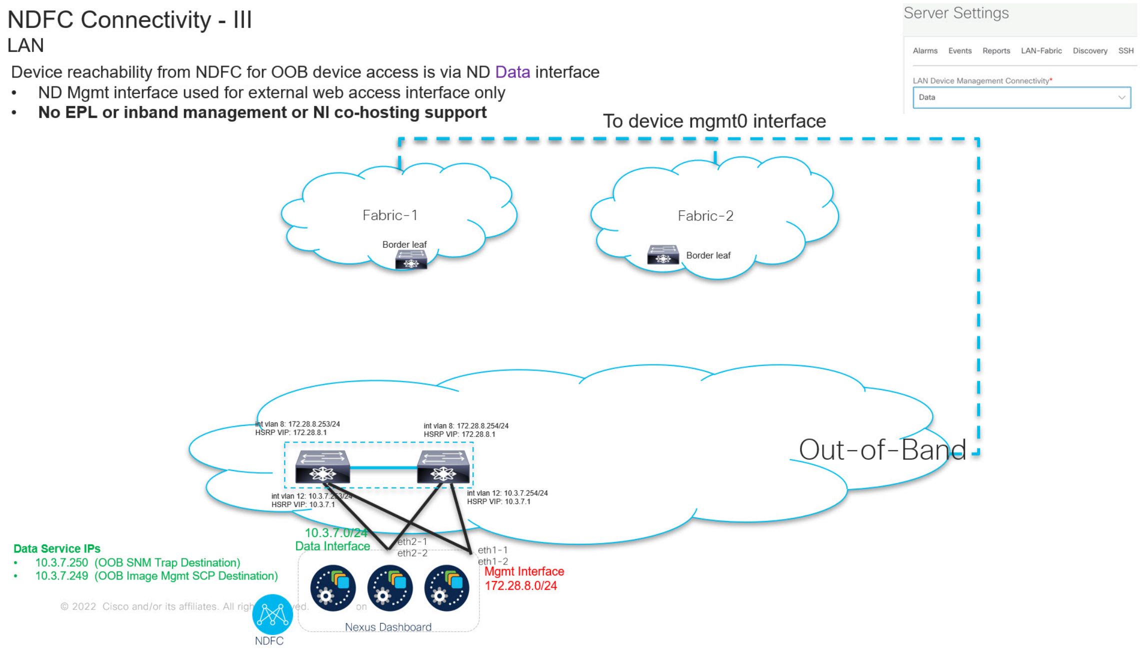

NDFC Connectivity - III

LAN

Device reachability from NDFC for OOB device access is via ND Data interface

- ND Mgmt interface used for external web access interface only
- **No EPL or inband management or NI co-hosting support**

Server Settings

Alarms Events Reports LAN-Fabric Discovery SSH

LAN Device Management Connectivity*: Data

Data Service IPs

- 10.3.7.250 (OOB SNM Trap Destination)
- 10.3.7.249 (OOB Image Mgmt SCP Destination)

CHAPTER 2

System Requirements

System Requirements

This chapter lists the tested and supported hardware and software specifications for Cisco Nexus Dashboard Fabric Controller architecture. The application is in English locales only.

The following sections describes the various system requirements for the proper functioning of your Cisco Nexus Dashboard Fabric Controller, Release 12.1.1p.

Note We recommend that you do not upgrade any underlying third-party software separately. All the necessary software components will be updated during the inline upgrade procedure. Upgrading the components outside of Nexus Dashboard Fabric Controller upgrade causes functionality issues.

- Cisco Nexus Dashboard Version Compatibility
- Nexus Dashboard Server Resource (CPU/Memory) Requirements
- Nexus Dashboard Networks
- Nexus Dashboard Fabric Controller Ports
- Supported Latency
- Supported Web Browsers
- Other Supported Software

Cisco Nexus Dashboard Version Compatibility

Cisco Nexus Dashboard Fabric Controller (NDFC) requires Nexus Dashboard version 2.2.1h or higher. If you try to upload NDFC 12.1.1p on a Nexus Dashboard version earlier than 2.2.1h, you will not be allowed to upload the application. To download the correct version of Nexus Dashboard, visit Software Download – Nexus Dashboard.

Nexus Dashboard Server Resource (CPU/Memory) Requirements

The following table provides information about Server Resource (CPU/Memory) Requirements to run NDFC on top of Nexus Dashboard. Refer to Nexus Dashboard Capacity Planning to determine the number of switches supported for each deployment.

Table 2: Server Resource (CPU/Memory) Requirements to run NDFC on top of Nexus Dashboard

Deployment Type	Node Type	CPUs	Memory	Storage (Throughput: 40-50MB/s)
Fabric Discovery	Virtual Node (vND) – app OVA	16vCPUs	64GB	550GB SSD
	Physical Node (pND) (PID: SE-NODE-G2)	2x 10-core 2.2G Intel Xeon Silver CPU	256 GB of RAM	4x 2.4TB HDDs 400GB SSD 1.2TB NVME drive
Fabric Controller	Virtual Node (vND) – app OVA	16vCPUs	64GB	550GB SSD
	Physical Node (pND) (PID: SE-NODE-G2)	2x 10-core 2.2G Intel Xeon Silver CPU	256 GB of RAM	4x 2.4TB HDDs 400GB SSD 1.2TB NVME drive

Deployment Type	Node Type	CPUs	Memory	Storage (Throughput: 40-50MB/s)
SAN Controller	Virtual Node (vND) – app OVA (without SAN Insights)	16vCPUs	64GB	550GB SSD
	Data Node (vND) – Data OVA (with SAN Insights)	32vCPUs	128GB	3TB SSD
	Physical Node (pND) (PID: SE-NODE-G2)	2x 10-core 2.2G Intel Xeon Silver CPU	256 GB of RAM	4x 2.4TB HDDs 400GB SSD 1.2TB NVME drive
	Virtual Node (vND) Virtual Node (Default Profile on Linux RHEL)	16vCPUs	64 GB	550GB SSD 500GB HDD **Note** SSD+HDD = 550GB
	Virtual Node (vND) Virtual Node (Large Profile on Linux RHEL)	32vCPUs	128 GB	3TB

Nexus Dashboard Networks

When first configuring Nexus Dashboard, on every node, you must provide two IP addresses for the two Nexus Dashboard interfaces—one connected to the Data Network and the other to the Management Network. The data network is typically used for the nodes' clustering and north-south connectivity to the physical network. The management network typically connects to the Cisco Nexus Dashboard Web UI, CLI, or API.

For enabling the Nexus Dashboard Fabric Controller, the Management and Data Interfaces on a Nexus Dashboard node must be in different subnets. Different nodes that belong to the same Nexus Dashboard cluster can either be Layer-2 adjacent or Layer-3 adjacent. Refer to Layer 3 Reachability Between Cluster Nodes, on page 5 for more information.

Connectivity between the Nexus Dashboard nodes is required on both networks with the round trip time (RTT) not exceeding 50ms. Other applications running on the same Nexus Dashboard cluster may have lower RTT requirements and you must always use the lowest RTT requirement when deploying multiple applications in the same Nexus Dashboard cluster. Refer to for more information.

Management Interface	Data Interface	Persistent IPs
Layer 2 adjacent	Layer 2 adjacent	One of the following for LAN: • If using default LAN Device Management Connectivity (set to Management): • 2 IPs in management network for SNMP/Syslog and SCP services • Plus one IP per fabric for EPL (if enabled) in data network • Plus one IP for Telemetry receiver in management network if IP Fabric for Media is enabled • If LAN Device Management Connectivity is set to Data: • 2 IPs in data network for SNMP/Syslog and SCP services • Plus one IP per fabric for EPL (if enabled) in data network • Plus one IP for Telemetry receiver in data network if IP Fabric for Media is enabled
Layer 3 adjacent	Layer 3 adjacent	For LAN: • LAN Device Management Connectivity on NDFC must be set to Data • 2 IPs for SNMP/Syslog and SCP/POAP services • Plus one IP per fabric for EPL These IPs must be part of a subnet that is different from Nexus Dashboard management and Nexus Dashboard data subnets associated with any of Nexus Dashboard nodes. These IPs must belong to the Layer-3 External Persistent Service Pool.

Virtual Nexus Dashboard (vND) Prerequisites

For virtual Nexus Dashboard deployments, each vND node has 2 interfaces or vNICs. The Data vNIC maps to bond0 (also known as bond0br) interface and Management vNIC maps to bond1 (also known as bond1br) interface. The requirement is to enable/accept promiscuous mode on the port groups that are associated with the Nexus Dashboard Management and/or Data vNICs where IP stickiness is required. The Persistent IP addresses are given to the pods (for example, SNMP Trap or Syslog receiver, Endpoint Locator instance per Fabric, and so on). Every POD in Kubernetes can have multiple virtual interfaces. Specifically for IP stickiness, an extra virtual interface is associated with the POD that is allocated an appropriate free IP from the external service IP pool. The vNIC has its own unique MAC address that is different from the MAC addresses associated with the vND virtual vNICs. Moreover, all North-to-South communication to and from these pods go out of the same bond interface. By default, the VMware ESXi systems check if the traffic flows out of a particular VM vNIC that matches the Source-MAC that is associated with that vNIC. If NDFC pods with an external service IP, the traffic flows are sourced with the Persistent IP addresses of the given pods that map to the individual POD MAC associated with the virtual POD interface. Therefore, enable the required settings on the VMware side to allow this traffic to flow seamlessly in and out of the vND node.

When vND nodes are deployed with the new Layer-3 HA feature, you need not enable Promiscuous mode on the vND vNIC interfaces. Promiscuous mode is required only for vND deployments when the vNDs are layer-2 adjacent from each other.

For more information, refer to .

Nexus Dashboard Fabric Controller Ports

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

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

Table 3: Nexus Dashboard Fabric Controller Ports

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
SSH	22	TCP	Out	SSH is a basic mechanism for accessing devices.
SCP	22	TCP	Out	SCP clients archiving NDFC backup files to remote server.
SMTP	25	TCP	Out	SMTP port is configurable through NDFC's **Server Settings** menu. This is an optional feature.

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

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

Table 4: Nexus Dashboard Fabric Controller Persistent IP Ports

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
SCP	22	TCP	In	SCP is used by various features to transfer files between devices and the NDFC service. The NDFC SCP service serves as the SCP server for both downloads and uploads. SCP is also used by the POAP client on the devices to download POAP-related files. The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the NDFC Server Settings.
TFTP (POAP)	69	TCP	In	Only used for device zero-touch provisioning via POAP, where devices can send (limited jailed write-only access to NDFC) basic inventory information to NDFC to start secure POAP communication. NDFC Bootstrap or POAP can be configured for TFTP or HTTP/HTTPS. The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the NDFC Server Settings.

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
HTTP (POAP)	80	TCP	In	Only used for device zero-touch provisioning via POAP, where devices can send (limited jailed write-only access to NDFC) basic inventory information to NDFC to start secure POAP communication. NDFC Bootstrap or POAP can be configured for TFTP or HTTP/HTTPS. The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the NDFC Server Settings.
BGP	179	TCP	In/Out	For Endpoint Locator, per fabric where it is enabled, an EPL service is spawned with its own persistent IP. This service is always associated with the Nexus Dashboard data interface. NDFC EPL service peers with the appropriate BGP entity (typically BGP Route-Reflectors) on the fabric to get BGP updates needed to track endpoint information. This feature is only applicable for VXLAN BGP EVPN fabric deployments.
HTTPS (POAP)	443	TCP	In	Secure POAP is accomplished via the NDFC HTTPS Server on port 443. The HTTPS server is bound to the SCP-POAP service and uses the same persistent IP assigned to that pod. The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the NDFC Server Settings.

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
Syslog	514	UDP	In	When NDFC is configured as a Syslog server, Syslogs from the devices are sent out toward the persistent IP associated with the SNMP-Trap/Syslog service pod The SNMP-Trap-Syslog service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the NDFC Server Settings
SCP	2022	TCP	Out	Transport tech-support file from persistent IP of NDFC POAP-SCP pod to a separate ND cluster running Nexus Dashboard Insights. The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the NDFC Server Settings
SNMP Trap	2162	UDP	In	SNMP traps from devices to NDFC are sent out toward the persistent IP associated with the SNMP-Trap/Syslog service pod. The SNMP-Trap-Syslog service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the NDFC Server Settings
GRPC (Telemetry)	50051	TCP	In	Information related to multicast flows for IP Fabric for Media deployments as well as PTP for general LAN deployments is streamed out via software telemetry to a persistent IP associated with a NDFC GRPC receiver service pod.

Supported Latency

As Cisco Nexus Dashboard Fabric Controller is deployed atop Cisco Nexus Dashboard, the latency factor is dependent on Cisco Nexus Dashboard. Refer to for information about latency.

Supported Web Browsers

Cisco Nexus Dashboard Fabric Controller is supported on the following web browsers:

- Google Chrome version 101.0.4951.64
- Microsoft Edge version 101.0.1210.47 (64-bit)
- Mozilla Firefox version 100.0.1 (64-bit)

Other Supported Software

The following table lists the other software that is supported by Cisco Nexus Dashboard Fabric Controller Release 12.1.1p.

Component	Features
Security	• ACS versions 4.0, 5.1, 5.5, and 5.8 • ISE version 2.6 • ISE version 3.0 • Telnet Disabled: SSH Version 1, SSH Version 2, Global Enforce SNMP Privacy Encryption. • Web Client: HTTPS with TLS 1, 1.1, 1.2, and 1.3

CHAPTER 3

Prerequisites

This chapter provides release-specific prerequisites information for your deployment of *Cisco Nexus Dashboard Fabric Controller*.

Prerequisites

Before you install the Cisco Nexus Dashboard Fabric Controller on Cisco Nexus Dashboard, you must need to meet the following prerequisites:

Nexus Dashboard

You must have Cisco Nexus Dashboard cluster deployed and its fabric connectivity configured, as described in before proceeding with any additional requirements and the Nexus Dashboard Fabric Controller service installation described here.

Nexus Dashboard Fabric Controller Release	Minimum Nexus Dashboard Release
Release 12.1.1p	Cisco Nexus Dashboard, Release or later **Note** Cisco Nexus Dashboard cluster in Linux KVM does not support Nexus Dashboard Fabric Controller Release 12.1.1p.

Nexus Dashboard Networks

When first configuring Nexus Dashboard, on every node, you must provide two IP addresses for the two Nexus Dashboard interfaces—one connected to the Data Network and the other to the Management Network. The data network is typically used for the nodes' clustering and north-south connectivity to the physical network. The management network typically connects to the Cisco Nexus Dashboard Web UI, CLI, or API.

For enabling the Nexus Dashboard Fabric Controller, the Management and Data Interfaces on a Nexus Dashboard node must be in different subnets. Different nodes that belong to the same Nexus Dashboard cluster can either be Layer-2 adjacent or Layer-3 adjacent. Refer to Layer 3 Reachability Between Cluster Nodes, on page 5 for more information.

Connectivity between the Nexus Dashboard nodes is required on both networks with the round trip time (RTT) not exceeding 50ms. Other applications running on the same Nexus Dashboard cluster may have lower RTT requirements and you must always use the lowest RTT requirement when deploying multiple applications in the same Nexus Dashboard cluster. Refer to for more information.

Management Interface	Data Interface	Persistent IPs
Layer 2 adjacent	Layer 2 adjacent	One of the following for LAN: • If using default LAN Device Management Connectivity (set to Management): • 2 IPs in management network for SNMP/Syslog and SCP services • Plus one IP per fabric for EPL (if enabled) in data network • Plus one IP for Telemetry receiver in management network if IP Fabric for Media is enabled • If LAN Device Management Connectivity is set to Data: • 2 IPs in data network for SNMP/Syslog and SCP services • Plus one IP per fabric for EPL (if enabled) in data network • Plus one IP for Telemetry receiver in data network if IP Fabric for Media is enabled
Layer 3 adjacent	Layer 3 adjacent	For LAN: • LAN Device Management Connectivity on NDFC must be set to Data • 2 IPs for SNMP/Syslog and SCP/POAP services • Plus one IP per fabric for EPL These IPs must be part of a subnet that is different from Nexus Dashboard management and Nexus Dashboard data subnets associated with any of Nexus Dashboard nodes. These IPs must belong to the Layer-3 External Persistent Service Pool.

Virtual Nexus Dashboard (vND) Prerequisites

For virtual Nexus Dashboard deployments, each vND node has 2 interfaces or vNICs. The Data vNIC maps to bond0 (also known as bond0br) interface and Management vNIC maps to bond1 (also known as bond1br) interface. The requirement is to enable/accept promiscuous mode on the port groups that are associated with the Nexus Dashboard Management and/or Data vNICs where IP stickiness is required. The Persistent IP addresses are given to the pods (for example, SNMP Trap or Syslog receiver, Endpoint Locator instance per

Fabric, and so on). Every POD in Kubernetes can have multiple virtual interfaces. Specifically for IP stickiness, an extra virtual interface is associated with the POD that is allocated an appropriate free IP from the external service IP pool. The vNIC has its own unique MAC address that is different from the MAC addresses associated with the vND virtual vNICs. Moreover, all North-to-South communication to and from these pods go out of the same bond interface. By default, the VMware ESXi systems check if the traffic flows out of a particular VM vNIC that matches the Source-MAC that is associated with that vNIC. If NDFC pods with an external service IP, the traffic flows are sourced with the Persistent IP addresses of the given pods that map to the individual POD MAC associated with the virtual POD interface. Therefore, enable the required settings on the VMware side to allow this traffic to flow seamlessly in and out of the vND node.

When vND nodes are deployed with the new Layer-3 HA feature, you need not enable Promiscuous mode on the vND vNIC interfaces. Promiscuous mode is required only for vND deployments when the vNDs are layer-2 adjacent from each other.

For more information, refer to .

Nexus Dashboard Cluster Sizing

Nexus Dashboard supports cohosting of services. Depending on the type and number of services you choose to run, you may be required to deploy extra worker nodes in your cluster. For cluster sizing information and recommended number of nodes based on specific use cases, see the Cisco Nexus Dashboard Capacity Planning tool.

If you plan to host other applications in addition to the Nexus Dashboard Fabric Controller, ensure that you deploy and configure additional Nexus Dashboard nodes based on the cluster sizing tool recommendation, as described in the , which is also available directly from the Nexus Dashboard Web UI.

Network Time Protocol (NTP)

Nexus Dashboard Fabric Controller uses NTP for clock synchronization, so you must have an NTP server configured in your environment.

Clocks on all nodes must be synchronized within the same second. Any delta between two nodes that exceeds more than 1 second could affect database consistency mechanism between the nodes.

CHAPTER 4

Installing Cisco Nexus Dashboard Fabric Controller

This chapter contains the following sections:

Note Only fresh installation of NDFC Release 12.1.1p is supported. You cannot upgrade to Release 12.1.1p from older NDFC or DCNM releases.

Installing Nexus Dashboard Fabric Controller Service Using App Store

To install Cisco Nexus Dashboard Fabric Controller Release 12.1.1p in an existing Cisco Nexus Dashboard cluster, perform the following steps:

Before you begin

- Ensure that you’ve installed the required form factor of Cisco Nexus Dashboard. For instructions, refer to .
- Ensure that you meet the requirements and guidelines described in Prerequisites, on page 21.
- If you choose to deploy NDFC on Nexus Dashboard on KVM, you must create bridge interfaces on Linux before installing Nexus Dashboard on KVM with Centos7. Ensure that you use bridge interfaces and do not allow other interfaces during Nexus Dashboard installation.
- The Cisco DC App Center must be reachable from the Nexus Dashboard via the Management Network directly or using a proxy configuration. Nexus Dashboard proxy configuration is described in the .

  If you are unable to establish the connection to the DC App Center, skip this section and follow the steps described in Installing Nexus Dashboard Fabric Controller Service Manually, on page 26.
- Ensure that the services are allocated with IP pool addresses on the Cisco Nexus Dashboard. For more information, refer to *Cluster Configuration* section in .

Procedure

Step 1 Launch the Cisco **Nexus Dashboard** Web UI using appropriate credentials.

Step 2 Click on **Admin Console > Services** menu in the left navigation pane to open the Services Catalog window.

Step 3 On the **App Store** tab, identify the Nexus Dashboard Fabric Controller Release 12.1.1p card and click **Install**.

Step 4 On the License Agreement screen, read the CISCO APP CENTER AGREEMENT and click on **Agree and Download**.

Wait for the application to be downloaded to the Nexus Dashboard and deployed.

It may take up to 30 minutes for the application to replicate to all nodes and all services to fully deploy.

Nexus Dashboard Fabric Controller application appears in the **Services Catalog**. The status is shown as **Initializing**.

Step 5 Click **Enable**.

After the services are enabled, the button on the Nexus Dashboard Fabric Controller card shows **Open**.

Wait until all the pods and containers are up and running.

Step 6 Click on **Open** to launch Cisco Nexus Dashboard Fabric Controller Web UI.

Note The single sign-on (SSO) feature allows you to log in to the application using the same credentials as you used for the Nexus Dashboard.

The **Nexus Dashboard Fabric Controller Web UI** opens in a new browser. The **Feature Management** window appears.

Note If External Service Pool IP addresses are not configured, an error message appears. Go to **Nexus Dashboard Web UI > Infrastructure > Cluster Configuration**. Configure the Management Service and Data Service IP addresses in the External Service Pools section. For more information, refer to *Cluster Configuration* section in .

Three cards namely **Fabric Discovery**, **Fabric Controller**, and **SAN Controller** is displayed.

Step 7 Based on the requirement, select the deployment.

From the list of Features, select features that you need to enable on the Nexus Dashboard Fabric Controller deployment.

Note The list of features displayed is based on the Deployment selected on the card.

Step 8 Click **Apply** to deploy Nexus Dashboard Fabric Controllerwith the selected features.

After the installation is complete, the deployment card and all the features status show as **Started**.

Installing Nexus Dashboard Fabric Controller Service Manually

To manually upload and install Cisco Nexus Dashboard Fabric Controller Release 12.1.1p in an existing Cisco Nexus Dashboard cluster, perform the following steps:

Before you begin

- Ensure that you've installed the required form factor of Cisco Nexus Dashboard. For instructions, refer to .
- Ensure that you meet the requirements and guidelines described in Prerequisites, on page 21.
- If you choose to deploy NDFC on Nexus Dashboard on KVM, you must create bridge interfaces on Linux before installing Nexus Dashboard on KVM with Centos7. Ensure that you use bridge interfaces and do not allow other interfaces during Nexus Dashboard installation.
- Ensure that the services are allocated with IP pool addresses on the Cisco Nexus Dashboard. For more information, refer to *Cluster Configuration* section in .

Procedure

Step 1 Go to the following site: https://dcappcenter.cisco.com.

Cisco DC App Center page opens.

In the **All apps** section, all the applications supported on Cisco Nexus Dashboard.

Step 2 Locate the Cisco Nexus Dashboard Fabric Controller Release 12.1.1p application and click the **Download** icon.

Step 3 On the License Agreement screen, read the CISCO APP CENTER AGREEMENT and click on **Agree and Download**.

Save the Nexus Dashboard Fabric Controller application to your directory that is easy to find when you must import/upload to Nexus Dashboard.

Step 4 Launch the Cisco **Nexus Dashboard** using appropriate credentials.

Step 5 Choose **Admin Console** > **Services** > **Installed Services** to view the services installed on the Cisco Nexus Dashboard.

Step 6 From the **Actions** drop-down list, choose **Upload Service**.

Step 7 Choose the **Location** toggle button and select either Remote or Local.

You can choose to either upload the service from a remote or local directory.

- If you select **Remote**, in the **URL** field, provide an absolute path to the directory where the Nexus Dashboard Fabric Controller application is saved.
- If you select **Local**, click **Browse** and navigate to the location where the Nexus Dashboard Fabric Controller application is saved. Select the application and click **Open**.

Step 8 Click **Upload**.

Nexus Dashboard Fabric Controller application appears in the Services Catalog. The status is shown as Initializing.

Wait for the application to be downloaded to the Nexus Dashboard and deployed.

It may take up to 30 minutes for the application to replicate to all nodes and all services to fully deploy.

Nexus Dashboard Fabric Controller application appears in the **Services Catalog**. The status is shown as **Initializing**.

Step 9 Click **Enable**.

After the services are enabled, the button on the Nexus Dashboard Fabric Controller card shows **Open**.

Wait until all the pods and containers are up and running.

Step 10 Click on **Open** to launch Cisco Nexus Dashboard Fabric Controller Web UI.

Note The single sign-on (SSO) feature allows you to log in to the application using the same credentials as you used for the Nexus Dashboard.

The **Nexus Dashboard Fabric Controller Web UI** opens in a new browser. The **Feature Management** window appears.

Note If External Service Pool IP addresses are not configured, an error message appears. Go to **Nexus Dashboard Web UI > Infrastructure > Cluster Configuration**. Configure the Management Service and Data Service IP addresses in the External Service Pools section. For more information, refer to *Cluster Configuration* section in .

Three cards namely **Fabric Discovery**, **Fabric Controller**, and **SAN Controller** is displayed.

Step 11 Based on the requirement, select the deployment.

From the list of Features, select features that you need to enable on the Nexus Dashboard Fabric Controller deployment.

Note The list of features displayed is based on the Deployment selected on the card.

Step 12 Click **Apply** to deploy Nexus Dashboard Fabric Controller with the selected features.

After the installation is complete, the deployment card and all the features status show as **Started**.