



Installing Nexus Dashboard Fabric Controller

- [Installation Requirements and Guidelines, on page 1](#)
- [Installing NDFC Using App Store, on page 16](#)
- [Installing NDFC Manually, on page 17](#)

Installation Requirements and Guidelines

The following sections describe the various requirements for deploying Nexus Dashboard Fabric Controller.

Network Time Protocol (NTP)

Nexus Dashboard nodes must be in synchronization with the NTP Server; however, there can be latency of up to 1 second between the Nexus Dashboard nodes. If the latency is greater than or equal to 1 second between the Nexus Dashboard nodes, this may result in unreliable operations on the NDFC cluster.

IPv4 and IPv6 Support

Prior releases of Nexus Dashboard supported either pure IPv4 or dual stack IPv4/IPv6 (for management network only) configurations for the cluster nodes. Beginning with release 3.0(1), Nexus Dashboard supports pure IPv4, pure IPv6, or dual stack IPv4/IPv6 configurations for the cluster nodes and services.

When defining IP configuration, the following guidelines apply:

- All nodes and networks in the cluster must have a uniform IP configuration – either pure IPv4, or pure IPv6, or dual stack IPv4/IPv6.
- If you deploy the cluster in pure IPv4 mode and want to switch to dual stack IPv4/IPv6 or pure IPv6, you must redeploy the cluster.
- For dual stack configurations:
 - Both external (data and management) and internal (app and services) networks must be in dual stack mode.
 - Partial configurations, such as IPv4 data network and dual stack management network, are not supported.
 - IPv6 addresses are also required for physical servers' CIMCs.
 - You can configure either IPv4 or IPv6 addresses for the nodes' management network during initial node bring up, but you must provide both types of IPs during the cluster bootstrap workflow.

Management IPs are used to log in to the nodes for the first time to initiate cluster bootstrap process.

- All internal certificates will be generated to include both IPv4 and IPv6 Subject Alternative Names (SANs).
- Kubernetes internal core services will start in IPv4 mode.
- DNS will serve and forward to both IPv4 and IPv6 and server both types of records.
- VxLAN overlay for peer connectivity will use data network's IPv4 addresses.
Both IPv4 and IPv6 packets are encapsulated within the VxLAN's IPv4 packets.
- The UI will be accessible on both IPv4 and IPv6 management network addresses.

- For pure IPv6 configurations:

- Pure IPv6 mode is supported for physical and virtual form factors only.
Clusters deployed in AWS, Azure, or an existing Red Hat Enterprise Linux (RHEL) system do not support pure IPv6 mode.
- You must provide IPv6 management network addresses when initially configuring the nodes.
After the nodes (physical, virtual, or cloud) are up, these IPs are used to log in to the UI and continue cluster bootstrap process.
- You must provide IPv6 CIDRs for the internal App and Service networks described above.
- You must provide IPv6 addresses and gateways for the data and management networks described above.
- All internal certificates will be generated to include IPv6 Subject Alternative Names (SANs).
- All internal services will start in IPv6 mode.
- VxLAN overlay for peer connectivity will use data network's IPv6 addresses.
IPv6 packets are encapsulated within the VxLAN's IPv6 packets.
- All internal services will use IPv6 addresses.

Nexus Dashboard

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



Note The Fabric Controller service cannot recover from a two `master` node failure of the Nexus Dashboard cluster where it is deployed. As a result, we recommend that you maintain at least one `standby` node in your Nexus Dashboard cluster and create regular backups of your NDFC configuration, as described in the **Operations > Backup and Restore** chapter of the [Cisco NDFC-Fabric Controller Configuration Guide](#) for your release.

If you run into a situation where two `master` nodes of your Nexus Dashboard cluster fail, you can follow the instructions described in the **Troubleshooting > Replacing Two Master Nodes with Standby Nodes** section of the [Cisco Nexus Dashboard User Guide](#) for your release to recover the cluster and NDFC configuration.

NDFC Release	Minimum Nexus Dashboard Release
Release 12.1.3	Cisco Nexus Dashboard, Release 3.0.1

The following Nexus Dashboard form factors are supported with NDFC deployments:

- Cisco Nexus Dashboard physical appliance (.iso)
- VMware ESX (.ova)
This release supports ESXi 7.0.
- Linux KVM (.qcow2)
This release supports CentOS 7.9 and RHEL 8.6.
- Existing Red Hat Enterprise Linux (SAN Controller persona only)
This release supports Red Hat Enterprise Linux (RHEL) 8.6

Nexus Dashboard Cluster Sizing

Refer to your release-specific [Verified Scalability Guide for NDFC](#) for information about the number of Nexus Dashboard cluster nodes required for the desired scale.

Nexus Dashboard supports co-hosting 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.

Nexus Dashboard System Resources

The following table provides information about Server Resource 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.

Cisco Nexus Dashboard can be deployed using number of different form factors. NDFC can be deployed on the following form factors:

- pND - Physical Nexus Dashboard
- vND - Virtual Nexus Dashboard
- rND - RHEL Nexus Dashboard

Table 1: Server Resource Requirements to run NDFC on top of Nexus Dashboard

Deployment Type	Node Type	CPUs	Memory	Storage (Throughput: 40-50 MB/s)
Fabric Discovery	Virtual Node (vND) – app node	16 vCPUs	64 GB	550 GB SSD
	Physical Node (pND) (PID: SE-NODE-G2)	2 x 10-core 2.2GHz Intel Xeon Silver CPU	256 GB of RAM	4 x 2.4 TB HDDs 400 GB SSD 1.2 TB NVME drive
	Physical Node (pND) (PID: ND-NODE-L4)	2.8GHz AMD CPU	256 GB of RAM	4 x 2.4 TB HDDs 960 GB SSD 1.6 TB NVME drive
Fabric Controller	Virtual Node (vND) – app node	16 vCPUs	64 GB	550 GB SSD
	Physical Node (pND) (PID: SE-NODE-G2)	2 x 10-core 2.2GHz Intel Xeon Silver CPU	256 GB of RAM	4 x 2.4 TB HDDs 400 GB SSD 1.2 TB NVME drive
	Physical Node (pND) (PID: ND-NODE-L4)	2.8GHz AMD CPU	256 GB of RAM	4 x 2.4 TB HDDs 960 GB SSD 1.6 TB NVME drive

Deployment Type	Node Type	CPUs	Memory	Storage (Throughput: 40-50 MB/s)
SAN Controller	Virtual Node (vND) – app node (with SAN Insights)	16 vCPUs (with physical reservation)	64 GB (with physical reservation)	550 GB SSD
	App Node (rND) (with SAN Insights)	16 vCPUs (with physical reservation)	64 GB (with physical reservation)	550 GB SSD
	Data Node (vND) – Data node (with SAN Insights)	32 vCPUs (with physical reservation)	128GB (with physical reservation)	3 TB SSD
	Data Node (rND) (with SAN Insights)	32 vCPUs (with physical reservation)	128 GB (with physical reservation)	3 TB SSD
	Physical Node (pND) (PID: SE-NODE-G2)	2 x 10-core 2.2GHz Intel Xeon Silver CPU	256 GB of RAM	4 x 2.4 TB HDDs 400 GB SSD 1.2 TB NVME drive
	Physical Node (pND) (PID: ND-NODE-L4)	2.8GHz AMD CPU	256 GB of RAM	4 x 2.4 TB HDDs 960 GB SSD 1.6 TB NVME drive

Nexus Dashboard System Resources

The following table provides information about Server Resource 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.

Cisco Nexus Dashboard can be deployed using number of different form factors. NDFC can be deployed on the following form factors:

- pND - Physical Nexus Dashboard
- vND - Virtual Nexus Dashboard

Table 2: Server Resource Requirements to run NDFC on top of Nexus Dashboard

Deployment Type	Node Type	CPUs	Memory	Storage (Throughput: 40-50 MB/s)
Fabric Discovery	Virtual Node (vND) – app node	16 vCPUs	64 GB	550 GB SSD
	Physical Node (pND) (PID: SE-NODE-G2)	2 x 10-core 2.2GHz Intel Xeon Silver CPU	256 GB of RAM	4 x 2.4 TB HDDs 400 GB SSD 1.2 TB NVME drive
	Physical Node (pND) (PID: ND-NODE-L4)	2.8GHz AMD CPU	256 GB of RAM	4 x 2.4 TB HDDs 960 GB SSD 1.6 TB NVME drive
Fabric Controller	Virtual Node (vND) – app node	16 vCPUs	64 GB	550 GB SSD
	Physical Node (pND) (PID: SE-NODE-G2)	2 x 10-core 2.2GHz Intel Xeon Silver CPU	256 GB of RAM	4 x 2.4 TB HDDs 400 GB SSD 1.2 TB NVME drive
	Physical Node (pND) (PID: ND-NODE-L4)	2.8GHz AMD CPU	256 GB of RAM	4 x 2.4 TB HDDs 960 GB SSD 1.6 TB NVME drive
SAN Controller	Virtual Node (vND) – app node (with SAN Insights)	16 vCPUs (with physical reservation)	64 GB (with physical reservation)	550 GB SSD
	Data Node (vND) – Data node (with SAN Insights)	32 vCPUs (with physical reservation)	128GB (with physical reservation)	3 TB SSD
	Physical Node (pND) (PID: SE-NODE-G2)	2 x 10-core 2.2GHz Intel Xeon Silver CPU	256 GB of RAM	4 x 2.4 TB HDDs 400 GB SSD 1.2 TB NVME drive
	Physical Node (pND) (PID: ND-NODE-L4)	2.8GHz AMD CPU	256 GB of RAM	4 x 2.4 TB HDDs 960 GB SSD 1.6 TB NVME drive

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 number of subnets needed for the Management and Data Interfaces on a Nexus Dashboard node varies, depending on the release:

- For releases prior to NDFC release 12.1.3, the management and data interfaces on a Nexus Dashboard node must be in different subnets. The External Service Pool IP addresses may come from certain subnet pools, depending on the type of deployment:
 - For LAN deployments, the External Service IPs may come from the Nexus Dashboard management subnet pool or the data subnet pool, depending on the configured settings.
 - For SAN deployments, the External Service IPs come from the Nexus Dashboard data subnet pool.
- For NDFC release 12.1.3, the LAN deployment requirements do not change but SAN deployments now support the management and data interfaces in the same subnet. When the management and data interfaces on a Nexus Dashboard node are in the same subnet, then the External Service Pool IP addresses also come from that same, single subnet.

Note that separate subnets are still supported as before for SAN deployments as well.

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](#) 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 [Nexus Dashboard Fabric Controller Deployment Guide](#) for more information.

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 (Applies to both LAN and SAN deployments, unless stated otherwise)
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.
DHCP	67	UDP	In	If NDFC local DHCP server is configured for Bootstrap/POAP purposes. This applies to LAN deployments only. Note When using NDFC as a local DHCP server for POAP purposes, all ND master node IPs must be configured as DHCP relays. Whether the ND nodes' management or data IPs are bound to the DHCP server is determined by the LAN Device Management Connectivity in the NDFC Server Settings.
DHCP	68	UDP	Out	
SNMP	161	TCP/UDP	Out	SNMP traffic from NDFC to devices.

Service	Port	Protocol	In—towards the cluster	Connection (Applies to both LAN and SAN deployments, unless stated otherwise)
			Out—from the cluster towards the fabric or outside world	
HTTPS/HTTP (NX-API)	443/80	TCP	Out	<p>NX-API HTTPS/HTTP client connects to device NX-API server on port 443/80, which is also configurable. This applies to LAN deployments only.</p> <p>NX-API is an optional feature, used by limited set of NDFC functions, such as:</p> <ul style="list-style-type: none"> Endpoint locator (EPL), as described in EPL Behavior with NX-API Configuration Layer 4 to Layer 7 services, as described in Guidelines and Limitations for L4-L7 Services <p>You must enable the NX-API feature in the Advanced tab for the appropriate fabrics, such as the VXLAN, Enhanced Classic LAN, eBGP, and Campus fabrics.</p>
HTTPS (vCenter, Kubernetes, OpenStack, Discovery)	443	TCP	Out	<p>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.</p> <p>This is an optional feature</p>

**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 (Applies to both LAN and SAN deployments, unless stated otherwise)
SCP	22	TCP	In	<p>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.</p> <p>The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the LAN Device Management Connectivity setting in the NDFC Server Settings.</p>
TFTP (POAP)	69	TCP	In	<p>Only used for device zero-touch provisioning via POAP, where devices can send (limited jailed write-only access to NDFC) basic inventory information to NDFC to start secure POAP communication. NDFC Bootstrap or POAP can be configured for TFTP or HTTP/HTTPS.</p> <p>The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the LAN Device Management Connectivity setting in the NDFC Server Settings.</p> <p>This applies to LAN deployments only.</p>

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

Service	Port	Protocol	Direction In—towards the cluster Out—from the cluster towards the fabric or outside world	Connection (Applies to both LAN and SAN deployments, unless stated otherwise)
Syslog	514	UDP	In	<p>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</p> <p>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</p>
SCP	2022	TCP	Out	<p>Transport tech-support file from persistent IP of NDFC POAP-SCP pod to a separate ND cluster running Nexus Dashboard Insights.</p> <p>The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the LAN Device Management Connectivity setting in the NDFC Server Settings</p>
SNMP Trap	2162	UDP	In	<p>SNMP traps from devices to NDFC are sent out toward the persistent IP associated with the SNMP-Trap/Syslog service pod.</p> <p>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</p>
GRPC (Telemetry)	33000	TCP	In	<p>SAN Insights Telemetry Server which receives SAN data (such as storage, hosts, flows, and so on) over GRPC transport tied to NDFC Persistent IP.</p> <p>This is enabled on SAN deployments only.</p>

Service	Port	Protocol	In—towards the cluster	Connection (Applies to both LAN and SAN deployments, unless stated otherwise)
			Out—from the cluster towards the fabric or outside world	
GRPC (Telemetry)	50051	TCP	In	Information related to multicast flows for IP Fabric for Media deployments as well as PTP for general LAN deployments is streamed out via software telemetry to a persistent IP associated with a NDFC GRPC receiver service pod. This is enabled on LAN and Media deployments only.

NDFC Latency Requirement

As Cisco Nexus Dashboard Fabric Controller is deployed in Cisco Nexus Dashboard, the latency factor is dependent on Cisco Nexus Dashboard. Refer to [Nexus Dashboard Fabric Controller Deployment Guide](#) for information about latency.

NDFC Network Connectivity

- **LAN Device Management Connectivity** – Fabric discovery and Fabric controller features can manage Devices over both Management Network and Data Network of ND Cluster Appliances.
- When using Management network, add the routes to all subnets of the devices that NDFC needs to manage or monitor in the Management Network.
- When using Data Network, add the route towards to all subnets of all devices for which POAP is enabled, when using the pre-packaged DHCP server in NDFC for touchless Day-0 device bring-up.
- SAN controller persona requires all the devices to be reachable via the Data network of Nexus Dashboard cluster nodes.

NDFC Persistent IP address

- If Nexus Dashboard cluster is deployed over a Layer 3 separation of network, you must configure BGP on all ND nodes.
- All Persistent IPs must be configured such that they are not part of any of the Nexus Dashboard nodes' subnets. This is supported only when LAN Device Management connectivity is Data. This is not supported with a cluster that co-hosts Nexus Dashboard Insights with NDFC.
- If Nexus Dashboard cluster is deployed with all nodes in the same subnet, persistent IPs can be configured to be from the same subnet.

In this case, persistent IPs must belong to the network chosen based on LAN Device Management connectivity setting in the NDFC Server Settings.

For more information, see [Persistent IP Requirements for NDFC](#).



Note Because this release supports NDFC in pure IPv4, pure IPv6, or dual stack IPv4/IPv6, the following Persistent IP requirements are per IP family.

For example, if you have deployed in dual stack mode and the following table states that two IP addresses are required in the management network, that means two IPv4 addresses and two IPv6 addresses.

Management Interface	Data Interface	Persistent IPs
Layer 2 adjacent	Layer 2 adjacent	<p>When operating in Layer 2 mode with LAN deployment type and LAN Device Management Connectivity set to Management (default)</p> <ul style="list-style-type: none"> • 2 IPs in the management network for SNMP/Syslog and SCP services • If EPL is enabled, 1 additional IP in the data network for each fabric • If IP Fabric for Media is enabled, 1 additional IP in the management network for telemetry <p>When operating in Layer 2 mode with LAN deployment type and LAN Device Management Connectivity set to Data:</p> <ul style="list-style-type: none"> • 2 IPs in the data network for SNMP/Syslog and SCP services • If EPL is enabled, 1 additional IP in the data network for each fabric • If IP Fabric for Media is enabled, 1 additional IP in the data network for telemetry <p>For SAN Controller deployment type:</p> <ul style="list-style-type: none"> • 1 IP for SSH • 1 IP for SNMP/Syslog • 1 IP per Nexus Dashboard cluster node for SAN Insights functionality

Management Interface	Data Interface	Persistent IPs
Layer 3 adjacent	Layer 3 adjacent	<p>When operating in Layer 3 mode with LAN deployment type:</p> <ul style="list-style-type: none"> • LAN Device Management Connectivity must be set to <code>Data</code> • 2 IPs for SNMP/Syslog and SCP services • If EPL is enabled, 1 additional IP in the data network for each fabric • All persistent IPs must be part of a separate pool that must not overlap with the management or data subnets <p>For more information about Layer 3 mode for persistent IPs, see the Persistent IPs section in the User's Guide</p> <p>For SAN Controller deployment type:</p> <ul style="list-style-type: none"> • 1 IP for SSH • 1 IP for SNMP/Syslog • 1 IP per Nexus Dashboard cluster node for SAN Insights functionality <p>IP Fabric for Media is not supported in Layer 3 mode</p>

POAP related requirements

- Devices must support POAP.
- Device must have no start up configuration or **boot poap enable** command must be configured to bypass the start up configuration and enter the POAP mode.
- DHCP server with scope defined. For POAP purposes, either the pre-packaged NDFC DHCP server can be used or an external DHCP server.
- The script server that stores POAP script and devices' configuration files must be accessible.
- Software and Image Repository server must be used to store software images for the devices.

Web Browsers Compatibility

Cisco Nexus Dashboard Fabric Controller GUI 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 12.1.3.

Component	Features
Security	<ul style="list-style-type: none"> • 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

Installing NDFC Using App Store

To install Cisco Nexus Dashboard Fabric Controller Release 12.1.3 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 [Cisco Nexus Dashboard Deployment Guide](#).
- 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 [Cisco Nexus Dashboard User Guide](#).

If you are unable to establish the connection to the DC App Center, skip this section and follow the steps described in [Installing NDFC Manually, on page 17](#).
- Ensure that the services are allocated with IP pool addresses on the Cisco Nexus Dashboard. For more information, refer to *Cluster Configuration* section in [Cisco Nexus Dashboard User Guide](#).

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 12.1.3 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 [Cisco Nexus Dashboard User Guide](#).

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 Controller with the selected features.

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

Installing NDFC Manually

To manually upload and install Cisco Nexus Dashboard Fabric Controller Release 12.1.3 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 [Cisco Nexus Dashboard Deployment Guide](#).
- Ensure that the services are allocated with IP pool addresses on the Cisco Nexus Dashboard. For more information, refer to *Cluster Configuration* section in [Cisco Nexus Dashboard User Guide](#).

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 12.1.3 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 [Cisco Nexus Dashboard User Guide](#).

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**.
