Cisco Nexus Dashboard Fabric Controller Installation and Upgrade Guide, Release 12.1.3

First Published: 2023-08-23

Last Modified: 2023-08-23

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

CONTENTS

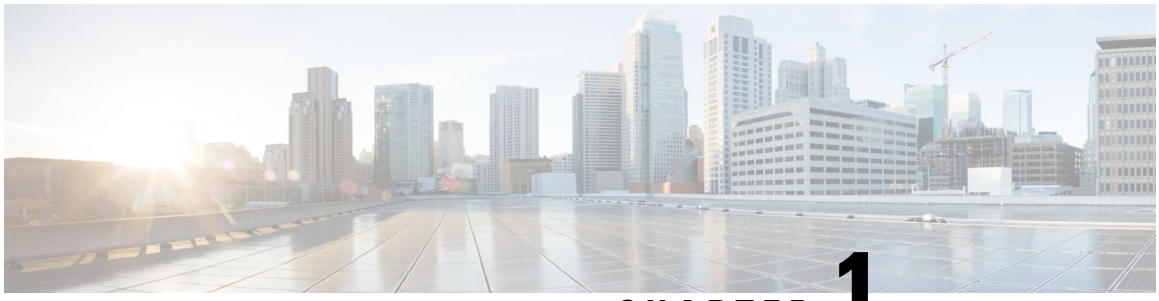

CHAPTER 1

Overview

Overview

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 Cisco Nexus Dashboard Deployment Guide.

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 12.1.3 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.

From Cisco NDFC Release 12.1.2, you can run NDFC on top of virtual Nexus Dashboard (vND) instance with promiscuous mode **disabled** on port groups that are associated with Nexus Dashboard interfaces where External Service IP addresses are specified. Recall that vND comprises a management interface and a data interface. By default, for LAN deployments, two external service IP addresses are required for the Nexus Dashboard management interface subnet. Similarly, by default, for SAN deployments, two external service IP addresses are required for the Nexus Dashboard data interface subnet.

Before the NDFC Release 12.1.2, if in-band management or Endpoint Locator or POAP feature was enabled on NDFC, you were required to enable promiscuous mode for the Nexus Dashboard data or fabric interface port-groups. This setting was mandatory for these features to work correctly. Again, as mentioned earlier, enabling promiscuous mode is no longer required for any port-groups associated with the vND. In fact, it is recommended to disable promiscuous mode for the port-groups post upgrade to ND 2.3.1/NDFC 12.1.2, in case customers are coming from previous versions.

Note

- Disabling promiscuous mode is supported from Cisco Nexus Dashboard Release 2.3.1c.
- You can disable promiscuous mode when Nexus Dashboard nodes are layer-3 adjacent on the Data network, BGP is configured, and fabric switches are reachable through the data interface.
- You can now disable promiscuous mode even when Nexus Dashboard interfaces are Layer-2 adjacent on the Management and Data networks.

Note Default option for promiscuous mode on VMware ESXi environments is **Reject**, meaning promiscuous mode is disabled.

This release of NDFC 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.

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

Deployment Options

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

- NDFC on single-node (non-HA Cluster)
  - Fabric Discovery for production environments (<= 50 switches)
  - Fabric Controller for production environments (<= 50 switches)
  - Fabric Controller in IP Fabric for Media controller mode for production environments
  - SAN Controller for production environments (<= 80 switches)
- NDFC on a 3-node Cluster (active-active HA mode)
  - 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)
  - Fabric Discovery
  - Fabric Controller
- NDFC on a 3-node physical Nexus Dashboard (pND) Cluster (active-active HA mode)
  - Nexus Dashboard Insights and NDFC in Fabric Controller persona (NDFC-Managed mode) – 3 pND nodes (<= 50 switches)
- NDFC on a Nexus Dashboard running on top of Red Hat Enterprise Linux (RHEL)
  - 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:
  - 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.

Cohosting of NDFC Managed mode with Nexus Dashboard Insights

From Release 12.1.1e, you can host NDFC Fabric Controller persona and Nexus Dashboard Insights on the same Nexus Dashboard Cluster in Managed mode to manage fabrics and Nexus Dashboard Insights to monitor the same fabrics. Note that NDFC in Fabric discovery mode, that is, monitored mode with NDI on the same Nexus Dashboard cluster is supported with NDFC Release 12.0.2f. Cohosting requires 4 physical Nexus Dashboard nodes for a maximum scale of up to 50 switches. This functionality is also supported on NDFC Release 12.1.1e with the corresponding paired Nexus Dashboard Insights release.

Note Nexus Dashboard deployed on KVM doesn't support cohosting NDFC and Insights service on the same Nexus Dashboard cluster.

Note For cohosting NDFC and Insights on the same Nexus Dashboard cluster, the Nexus Dashboard nodes must be Layer 2 adjacent. Support for Layer 3 adjacency for cohosting deployments will be introduced in future releases.

The following table shows the compatible versions for Nexus Dashboard and services.

Services	Compatible Version
Nexus Dashboard	3.0.1
Nexus Dashboard Insights	6.3.1
Nexus Dashboard Fabric Controller	12.1.3

The following table shows the system requirements for Nexus Dashboard.

Specification	Supported Scale
Number of physical Nexus Dashboard nodes	3

Specification	Supported Scale
Number of switches supported	50
Number of flows supported in Nexus Dashboard Insights	1000

Installation of NDFC and NDI on the same Nexus Dashboard

Cisco NDFC can be cohosted with Nexus Dashboard Insights on the same Nexus Dashboard.

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 you meet the requirements and guidelines described in *Prerequisites* section in *Cisco NDFC Installation 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 *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 Services Manually* section in *Cisco NDFC Installation 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*.

Installing Nexus Dashboard

Install the required form factor of Cisco Nexus Dashboard. For instructions, refer to Cisco Nexus Dashboard Deployment Guide.

Installing NDFC

Refer to *Cisco NDFC Installation Guide*.

Configure NDFC sites on Nexus Dashboard. Refer to the *Adding Sites* section in the Cisco Nexus Dashboard Deployment Guide.

Installing NDI

On the same Nexus Dashboard set up, install the Nexus Dashboard Insights service. Refer to Cisco Nexus Dashboard Insights Deployment Guide, for more information.

Post Installation

After installing compatible versions of NDFC and NDI on the 5-node physical Nexus Dashboard, launch NDFC as Fabric (LAN) Controller. Create Fabric, discover and import switches on NDFC fabric. Nexus Dashboard automatically identifies the NDFC fabric and lists on the Sites page as entities.

Note You must provide the password for each of the sites in the Nexus Dashboard site manager.

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 2.2.1, 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. In Nexus Dashboard, navigate to **Admin** > **System Settings**.
2. In the **Network Scale** tile, 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. For validated scale numbers, refer to Cisco NDFC Verified Scalability, Release 12.1.1e.

When you upgrade to this release, the individual containers are restarted and the newly spawned containers start with new resource requests and limit values.

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.

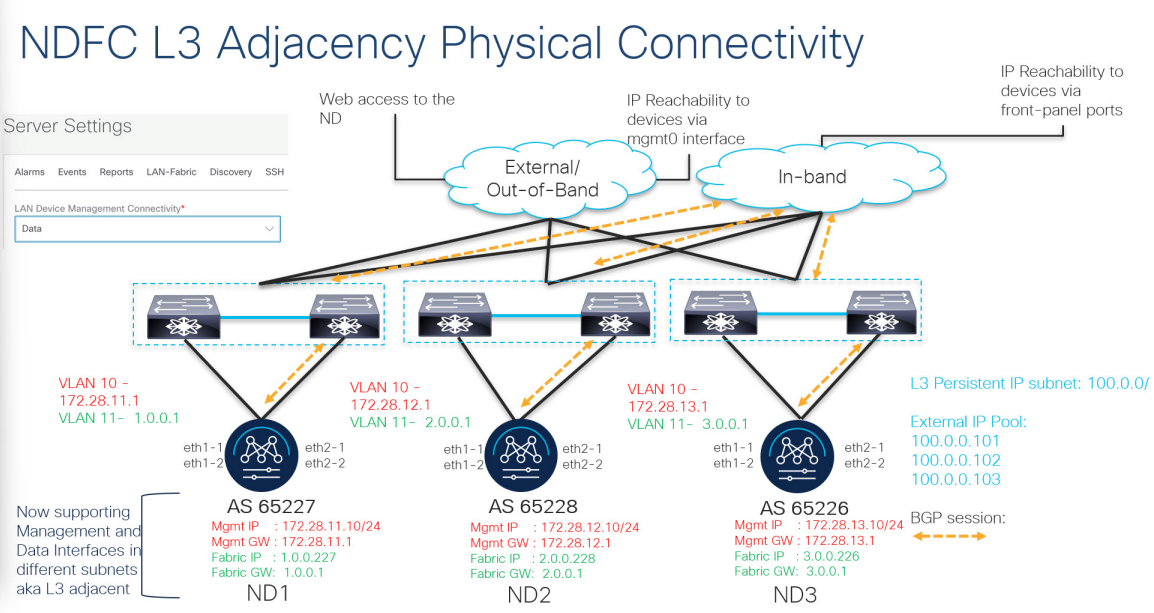

Upgrade or modification from an existing Layer-2 adjacent Nexus Dashboard cluster to a Layer-3 adjacent cluster is not supported. When using Layer 3 adjacency, NDFC service is supported only when the switch connectivity is through the Nexus Dashboard Data interface. Choose NDFC **UI > Settings > Admin**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

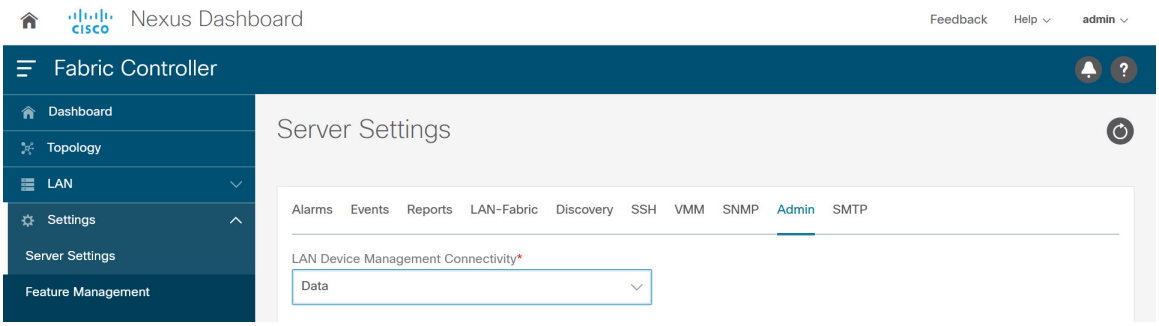

Nexus Dashboard uses eBGP to publish up-to-date reachability of /32 routes for reaching NDFC features using external service IPs obtained from the Persistent IP subnet. If a node or network fails, the external IPs are not reachable until recovery is complete (if the network can recover itself). After 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* 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

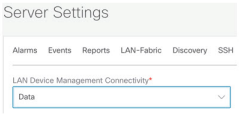

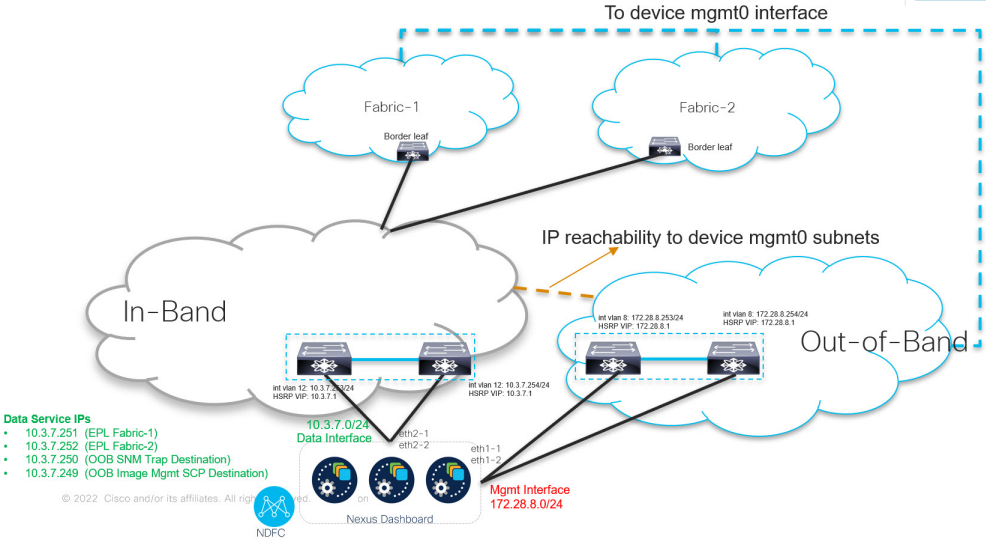

NDFC Connectivity - II

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

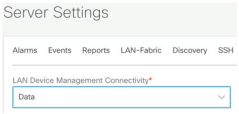

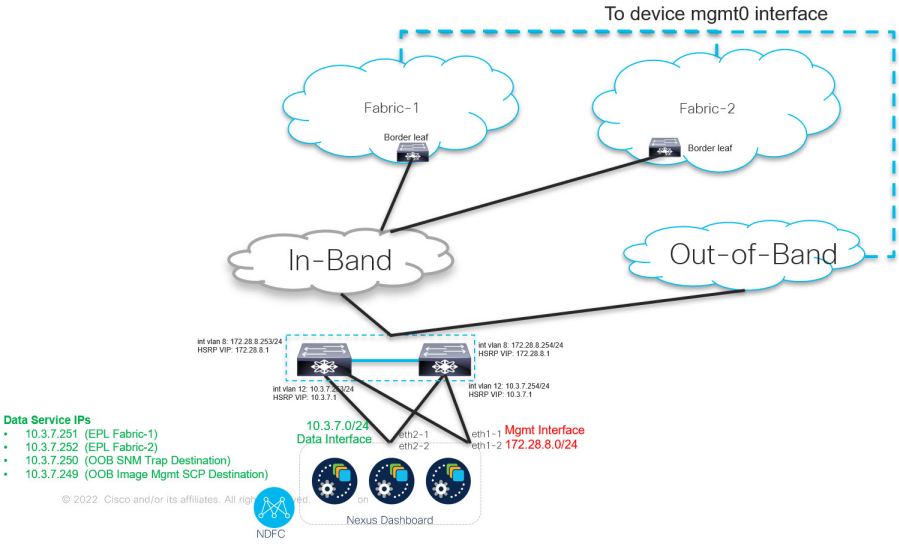

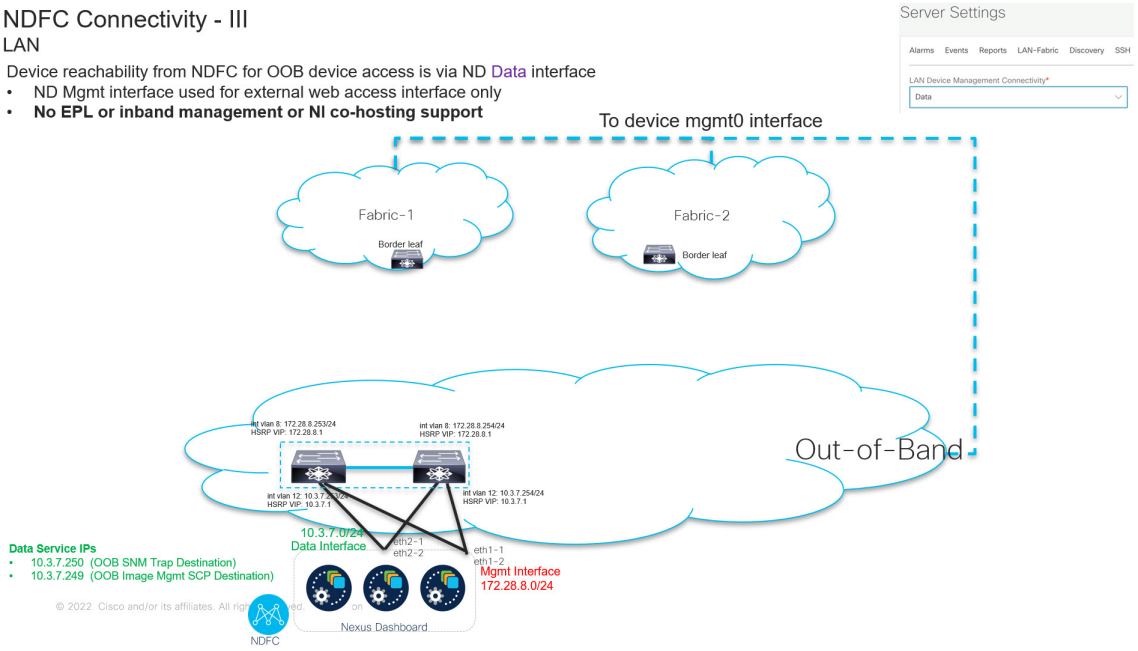
NDFC Connectivity - III
LAN
Device reachability from NDFC for OOB device access is via ND Data interface
• ND Mgmt interface used for external web access interface only
• No EPL or inband management or NI co-hosting support
Server Settings
Alarms Events Reports LAN-Fabric Discovery SSH
LAN Device Management Connectivity*
Data
To device mgmt0 interface
Fabric-1
Border leaf
Fabric-2
Border leaf
Out-of-Band
int vlan 8: 172.28.8.253/24
HSRP VIP: 172.28.8.1
int vlan 8: 172.28.8.254/24
HSRP VIP: 172.28.8.1
int vlan 12: 10.3.7.254/24
HSRP VIP: 10.3.7.1
10.3.7.0/24
Data Interface
eth2-1
eth2-2
eth1-1
eth1-2
Mgmt Interface
172.28.8.0/24
Data Service IPs
• 10.3.7.250 (OOB SNM Trap Destination)
• 10.3.7.249 (OOB Image Mgmt SCP Destination)
© 2022 Cisco and/or its affiliates. All rights reserved.
Nexus Dashboard
NDFC

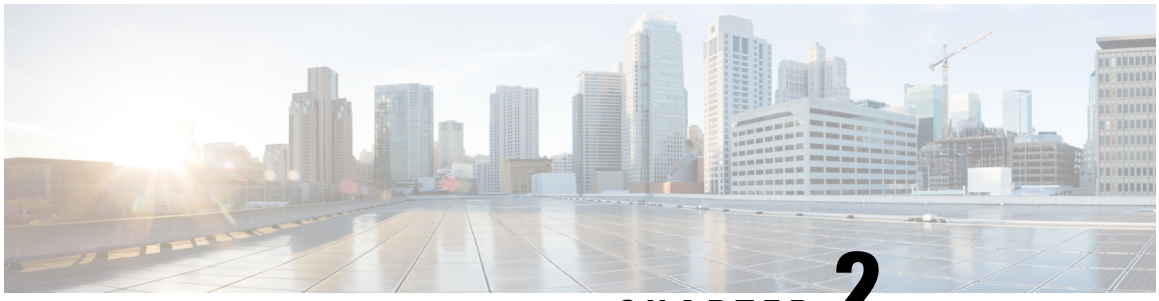

CHAPTER 2

Installing Nexus Dashboard Fabric Controller

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 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, on page 7 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 2: 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.

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

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 3: 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	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. This applies to LAN deployments only.

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)
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. This applies to LAN deployments only.
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. This applies to LAN deployments only.
HTTPS (POAP)	443	TCP	In	Secure POAP is accomplished via the NDFC HTTPS Server on port 443. The HTTPS server is bound to the SCP-POAP service and uses the same persistent IP assigned to that pod. The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the NDFC Server Settings. This applies to LAN deployments only.

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	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)	33000	TCP	In	SAN Insights Telemetry Server which receives SAN data (such as storage, hosts, flows, and so on) over GRPC transport tied to NDFC Persistent IP. This is enabled on SAN deployments only.

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)
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 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	When operating in Layer 2 mode with LAN deployment type and **LAN Device Management Connectivity** set to `Management` (default) • 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 When operating in Layer 2 mode with LAN deployment type and **LAN Device Management Connectivity** set to `Data`: • 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 For SAN Controller deployment type: • 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	When operating in Layer 3 mode with LAN deployment type: • **LAN Device Management Connectivity** must be set to `Data` • 2 IPs for SNMP/Syslog and SCP services • If **EPL** is enabled, 1 additional IP in the data network for each fabric • All persistent IPs must be part of a separate pool that must not overlap with the management or data subnets For more information about Layer 3 mode for persistent IPs, see the Persistent IPs section in the User's Guide For SAN Controller deployment type: • 1 IP for SSH • 1 IP for SNMP/Syslog • 1 IP per Nexus Dashboard cluster node for SAN Insights functionality IP Fabric for Media is not supported in Layer 3 mode

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	• 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 25.
- 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 Controllerwith 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**.

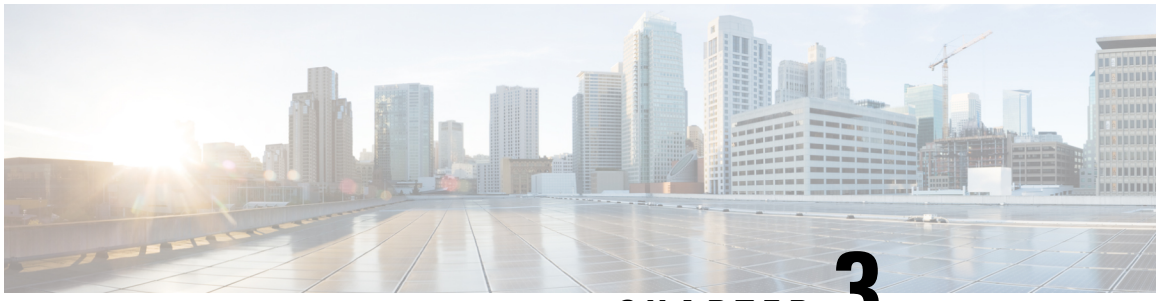

CHAPTER 3

Upgrading Nexus Dashboard Fabric Controller

Upgrade Prerequisites and Guidelines

The following sections describe the various requirements for upgrading an existing Nexus Dashboard Fabric Controller (NDFC) or the earlier DCNM software to this release.

- Ensure that your current Nexus Dashboard cluster is healthy.

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

- If you are upgrading a virtual Nexus Dashboard cluster deployed in VMware ESX, ensure that the ESX version is still supported by the target release.

Note If you need to upgrade the ESX server, you must do that before upgrading your Nexus Dashboard to the target release:

1. Upgrade one of the ESX hosts as you typically would with your existing Nexus Dashboard node VM running.
2. After the host is upgraded, ensure that the Nexus Dashboard cluster is still operational and healthy.
3. Repeat the upgrade on the other ESX hosts one at a time.
4. After all ESX hosts are upgraded and the existing Nexus Dashboard cluster is healthy, proceed with upgrading your Nexus Dashboard to the target release as described in the following section.

- If you are upgrading from release 12.1.1, ensure that the VRF is correctly set on all existing switches.

  For additional information and context, see the Discovery VRF on Existing Switches, on page 31 section below.

- If you take a backup of the NDFC configuration when upgrading from DCNM release 11.5.4, and then you restore on another release of NDFC (such as NDFC 12.1.3), in some situations during the initial setup, you might see an error message such as `Failed to allocate VLAN or Loopback ID/IP due to range exhausted`. This might happen if multiple networks or VRFs, or a combination of networks and VRFs, have the same default VLAN specified and are attached to the same switch, which will result in the default VLAN being set to -1 for the second attached entity. To fix this issue, you must specify an override VLAN to be used on that switch.
- You must perform configuration backups of your Nexus Dashboard and Nexus Dashboard Fabric Controller before the upgrade to safeguard data and minimize any potential risk before proceeding with the upgrade.
- If you are upgrading from DCNM release 11.5.4, the `trap-ip` used in DCNM must not be the Nexus Dashboard management or data IP address. Instead, the IP must be configured in the Nexus Dashboard persistent IP address pool.

  For more information about configuring persistent IPs, see the "Persistent IP addresses" section of the *Cisco Nexus Dashboard Infrastructure Management* article.
- You must disable all preview/beta features and delete the associated data.

  In the **Settings** > **Feature Management** page, disable all `BETA` features.

  In the **Settings** > **Server Settings** > **LAN Fabric** page, ensure that **Enable Preview Features** is disabled.

Note Disabling a beta feature does not remove any created configurations. For example, if you had enabled the ECL preview feature and created/discovered an ECL fabric, disabling the feature does not remove the existing configuration. In this case, in addition to disabling the feature, you must also delete the configuration before the upgrade.

- Ensure that you have deleted all NDFC versions except the currently enabled one from your Nexus Dashboard.

  You must not have multiple NDFC release images coexisting in the same cluster during the upgrade. Before you upgrade to the next release, navigate to the Nexus Dashboard's **Services** page, click the **...** menu in the NDFC service tile, choose **Versions** and remove all non-active versions:

- Ensure that you have enough external IP addresses configured in your Nexus Dashboard for the Fabric Controller service.

  If there are not enough IPs available, such if the IPs are consumed by another service, the NDFC service would not be able to start.

- You must be running Nexus Dashboard release 2.2(1) or later to upgrade directly to release 3.0(1).

Note If you are upgrading from NDFC release 12.0.2 or earlier, direct upgrades are not supported. You must first upgrade to release 12.1.1 and then to release 12.1.3. See the upgrade workflows table below for additional information.

- You must disable all services running in the cluster before upgrading to release 3.0(1) or later.

Note After you have upgraded your Nexus Dashboard to release 3.0(1), you must not re-enable the existing NDFC versions that you had disabled prior to the Nexus Dashboard upgrade.

You must also not delete the existing versions of the service, which may now display an `App Version not compatible. Upgrade required.` error:

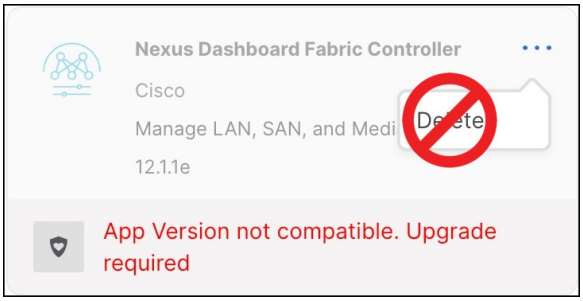

- Ensure that no configuration changes are made to the cluster, such as adding worker or standby nodes, while the upgrade is in progress.
- Nexus Dashboard and Nexus Dashboard Fabric Controller do not support software downgrades.

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

- 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 may cause functionality issues.

Discovery VRF on Existing Switches

Note The following information applies to NDFC upgrades from release 12.1.1; if you are upgrading from release 12.1.2, you can skip this subsection.

Reachability from NDFC to the switches is controlled by the routes configuration in the Nexus Dashboard cluster. By default, all NDFC pods have their default gateway set to the Nexus Dashboard's data interface gateway. However, for pods with persistent IPs (External Service IPs), the default gateway is set based on the specific management or data interface associated with the persistent IP pool:

- If a pod is using an external service IP from the management subnet, its default gateway is set to the Nexus Dashboard's management interface gateway.
- If a pod is using an external service IP from the data subnet pool, its default gateway is set to the Nexus Dashboard's data interface gateway.

NDFC's discovery pod requires IP reachability to the switches and does not have a persistent IP, so the pod's reachability to the switches depends on how routing is configured from the switches to the Nexus Dashboard cluster that is hosting the NDFC service. However, some of the pod's operations require reverse reachability from the switches to NDFC. For example, when NDFC is set as a trap-host destination for a switch, you must specify the VRF over which this IP is reachable in addition to the NDFC syslog-trap External Service IP configuration on that switch. Similarly, when a switch executes an SCP copy command to copy images for Image Management purposes, the VRF must be specified for the switch to reach the SCP destination IP (External Service IP associated with NDFC's poap-scp pod). As a result, for every switch, you must associate a VRF that can be used for connectivity from the switch to NDFC.

When NDFC's **LAN Device Management Connectivity** is set to `Management` (default setting), the POAP-SCP and SNMP-Trap pods get persistent IPs from the Management External Service IP pool. In this case, switches must be imported into NDFC using the switches' `mgmt0` IP address, which must be reachable from NDFC over the Nexus Dashboard's management interface. For switches that are reachable via Layer 3, this requires static routes to be configured in the Nexus Dashboard's **Admin** > **System Settings** > **Routes** > **Management Network Routes** (previously, **Infrastructure** > **Cluster Configuration**). For these switches, the VRF associated with the switch gets set to management.

On the other hand, when NDFC's **LAN Device Management Connectivity** is set to `Data`, the POAP-SCP and SNMP-Trap pods get persistent IPs from the Data External Service IP pool and are reachable over the Nexus Dashboard data interface. In this case, you must also enable the **When LAN Device Management Connectivity is set to Data, rely on LAN discovery to always populate VRF for all Nexus Switches** option for the VRF to be properly populated for the switch.

In NDFC release 12.1.1, the above option was disabled by default, which means that any switch imported into NDFC had its VRF set to "`default`". In this case, you had to enable reachability from the switch to NDFC via the front-panel interface or update the discovery VRF for the switch for the persistent IPs associated with the POAP-SCP and SNMP-Trap pods to be reachable. However, if you did not use Image Management or trap-related functionality and did not enable reachability from the switch to NDFC, you would not notice that the reverse communication from the switch to NDFC was not working because NDFC-to-switch communication would continue to function via either Nexus Dashboard's management or data interface.

Beginning with release 12.1.2, the **When LAN Device Management Connectivity is set to Data, rely on LAN discovery to always populate VRF for all Nexus Switches** option is enabled by default. So NDFC would properly populate the discovery VRF from the switch configuration during switch discovery.

As a result of the default setting change, if you are upgrading from release 12.1.1, you may run into an issue where the VRF is not set on the switch and you must manually update it using the GUI or API to establish proper reachability from your existing switches to NDFC. Note that if you had already explicitly configured the VRF for your existing switches, it is preserved during the upgrade but any switch whose VRF is currently set to "default" is not updated regardless of the change in default value of the **When LAN Device Management Connectivity is set to Data, rely on LAN discovery to always populate VRF for all Nexus Switches** setting unless the switch is removed and re-added.

Upgrade Workflow Overview

The following table summarizes the upgrade workflows required to upgrade from your current NDFC release to release 12.1.3.

Current Release	Nexus Dashboard Version for the Current Release	Upgrade workflow when upgrading to release 12.1.3
12.1.2	2.3.x	1. Create a configuration backup from the **Operations** > **Backup & Restore** page. 2. Disable **Nexus Dashboard Fabric Controller** service on **Nexus Dashboard** 3. Upgrade your **Nexus Dashboard** to release 3.0.1. 4. Upgrade the **NDFC** service to release 12.1.3. For detailed instructions, see Upgrading From NDFC Release 12.1.x, on page 34.
12.1.1e	2.2.x	1. Create a configuration backup from the **Operations** > **Backup & Restore** page. 2. Disable **Nexus Dashboard Fabric Controller** service on **Nexus Dashboard** 3. Upgrade your **Nexus Dashboard** to release 3.0.1. 4. Upgrade the **NDFC** service to release 12.1.3. For detailed instructions, see Upgrading From NDFC Release 12.1.x, on page 34.
12.0.2f	2.1.2d	Direct upgrade is not supported. We recommend upgrading your NDFC to release 12.1.2 as described in the *Nexus Dashboard Fabric Controller Installation and Upgrade Guide, Release 12.1.2e* before returning to this document to upgrade to release 12.1.3.
12.0.1a	2.1.1e	Direct upgrade is not supported. We recommend upgrading your NDFC to release 12.1.2 as described in the *Nexus Dashboard Fabric Controller Installation and Upgrade Guide, Release 12.1.2e* before returning to this document to upgrade to release 12.1.3.

Current Release	Nexus Dashboard Version for the Current Release	Upgrade workflow when upgrading to release 12.1.3
11.5(4)	Not Applicable	1. Backup using **DCNM_To_NDFC_Upgrade_Tool_OVA_ISO** 2. Install **Nexus Dashboard** version 3.0.1 3. Install **NDFC** Release 12.1.3 4. Restore on **Web UI > Operations > Backup & Restore** For detailed instructions, see Upgrading from DCNM (Data Center Network Manager), on page 36.

Upgrading From NDFC Release 12.1.x

The following steps described how to upgrade from release 12.1.x to release 12.1.3.

Before you begin

Ensure that you have:

- Familiarized yourself with the upgrade workflow and completed the prerequisites, as described in Upgrade Prerequisites and Guidelines, on page 29.

Procedure

Step 1 Back up existing configuration.

You can back up the configuration as you typically would from the **Operations** > **Backup & Restore** page.

If you take a local backup, download the copy and store it for later use if needed; if you take a remote backup, save the backup file in a secure remote location.

For additional information about backups, see *Backup & Restore* chapter in the *Cisco NDFC Configuration Guide*.

Step 2 Disable the existing NDFC service.

You can disable the service by navigating to the Nexus Dashboard's **Services** page and choosing **Disable** from the actions menu (**...**) on the NDFC tile.

Step 3 Upgrade your Nexus Dashboard to release 3.0.1.

Note Ensure that your NDFC service is disabled as described in the previous step before starting the Nexus Dashboard upgrade.

Detailed information about upgrading your Nexus Dashboard, including any additional prerequisites and guidelines is available in *Nexus Dashboard Deployment Guide*.

Step 4 After the Nexus Dashboard upgrade is completed, navigate to the **Services** page in the ND UI.

Note You must not re-enable or delete the current version of NDFC. Simply ignore the message and proceed to the next step.

At this point, you will see the following error in the NDFC service tile:

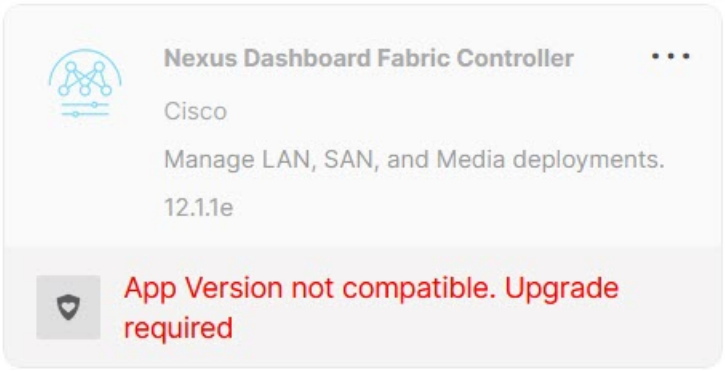

Step 5 Upgrade the NDFC service using the Cisco App Store.

Note The App Store allows you to upgrade to the latest available version of the service only. If a later release is available on the App Store when you are upgrading to this release, you must skip this step and manually upload the new service image as described in the next step.

a) Navigate to the **Nexus Dashboard** > **Services** page.
b) Choose the **App Store** tab.
c) In the App Store's NDFC tile, click **Update**.
d) Accept the **License Agreement**.

This starts the NDFC image download and installation.

e) Choose the **Installed Services** tab and wait for the image to be downloaded and installed.

Step 6 Upgrade the NDFC service by manually upload an image.

Note If you already used the App Store to upgrade as described above, skip this step.

You can choose to upgrade the service by manually uploading the new releases's image:

a) Obtain the upgrade image.

You can download this release's image from the Cisco App Center.

Optionally, you can choose to 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.

b) Navigate to the **Nexus Dashboard** > **Services** page.
c) Choose the **Installed Services** tab.
d) From the **Actions** menu, choose **Upload Service**.
e) Choose the image you downloaded in a previous substep and click **Upload**.

You can choose to either upload the image from your local machine or provide a full URL if you hosted the image on a server in your environment.

f) Wait for the image to upload and initialize.

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

Step 7 Enable the new image.

a) On the **Nexus Dashboard Fabric Controller** tile, click the actions menu (**…**) and choose **Available Versions**.

b) Click **Enable** next to the 12.1.3 version.

Note You must not delete the previous version at this time as it would result in a loss of the service's data before the upgrade completes.

Step 8 Wait for the final upgrade processes to complete.

When the process is done, the **Open** button on the NDFC's tile becomes available.

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

Step 9 Click **Open** to launch Nexus Dashboard Fabric Controller release 12.1.3 UI.

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

Caution You must not delete the previous NDFC image after the upgrade as it can cause a service disruption if the cluster is rebooted. Should you encounter this issue, contact Cisco TAC for assistance.

Upgrading from DCNM (Data Center Network Manager)

The following sections contain information specific to upgrades from an 11.x release of DCNM before the software moved to the Nexus Dashboard platform and was renamed to Nexus Dashboard Fabric Controller.

Note If you current release is already deployed in Nexus Dashboard, skip the following subsections.

Persona and Feature Compatibility After Upgrading From Cisco DCNM 11.5(4)

Persona Compatibility

By using the appropriate Upgrade Tool, you can restore data that is backed up from DCNM Release 11.5(4) on a newly deployed Cisco Nexus Dashboard Fabric Controller for the personas as mentioned in the following table:

Backup from DCNM 11.5(4)	Persona Enabled in NDFC 12.1.3 after Upgrade
DCNM 11.5(4) LAN Fabric Deployment on OVA/ISO/SE	Fabric Controller + Fabric Builder
DCNM 11.5(4) PMN Deployment on OVA/ISO/SE	Fabric Controller + IP Fabric for Media (IPFM)
DCNM 11.5(4) SAN Deployment on OVA/ISO/SE	SAN Controller
DCNM 11.5(4) SAN Deployment on Linux	SAN Controller

Backup from DCNM 11.5(4)	Persona Enabled in NDFC 12.1.3 after Upgrade
DCNM 11.5(4) SAN Deployment on Windows	SAN Controller

Feature Compatibility Post Upgrade

The following table lists caveats associated with features that are restored from DCNM 11.5(4) backup after upgrade to NDFC, Release 12.1.3.

Feature in DCNM 11.5(4)	Upgrade Support
Nexus Dashboard Insights configured Refer to Cisco Nexus Dashboard User Guide for more information.	Supported
Container Orchestrator (K8s) Visualizer	Supported
VMM Visibility with vCenter	Supported
Nexus Dashboard Orchestrator configured	Not Supported
Preview features configured	Not supported
LAN switches in SAN installations	Not supported
Switches discovered over IPv6	Not supported
DCNM Tracker	Not supported
Fabric Backups	Not supported
Report Definitions and Reports	Not supported
Switch images and Image Management policies	Not supported
SAN CLI templates	Not carried over from 11.5(4) to 12.1.3
Switch images/Image Management data	Not carried over from 11.5(4) to 12.1.3
Slow drain data	Not carried over from 11.5(4) to 12.1.3
Infoblox configuration	Not carried over from 11.5(4) to 12.1.3
Endpoint Locator configuration	You must reconfigure Endpoint Locator (EPL) post upgrade to Release 12.1.3. However, historical data is retained up to a maximum size of 500 MB.
Alarm Policy configuration	Not carried over from 11.5(4) to 12.1.3
Performance Management data	CPU/Memory/Interface statistics up to 90 days is restored post upgrade.

Note SAN Insights and VMM Visualizer features are not enabled after restore. You must choose check boxes on **Settings** > **Feature Management** and click **Save** to enable these features after restore.

Download NDFC Upgrade Tool

To download Upgrade tool to upgrade from Cisco DCNM to Nexus Dashboard Fabric Controller, perform the following steps:

Before you begin

- Identify the deployment type of Cisco DCNM Release 11.5(x) setup.

Procedure

Step 1 Browse to the NDFC download page: https://software.cisco.com/download/home/281722751/type/282088134/ .

A list of the latest release software for Cisco Nexus Dashboard Fabric Controller available for download is displayed.

Step 2 In the Latest Releases list, choose release 12.1.3.

Step 3 Based on your Cisco DCNM 11.5(x) deployment type, locate the **DCNM_To_NDFC_Upgrade_Tool** and click the **Download** icon.

The following table displays the DCNM 11.5(x) deployment type, and the corresponding Nexus Dashboard Fabric Controller upgrade tool that you must download.

Table 4: DCNM 11.5(x) Deployment type and Upgrade Tool Compatibility Matrix

DCNM 11.5(x) deployment type	UpgradeTool Name
ISO/OVA	`DCNM_To_NDFC_Upgrade_Tool_OVA_ISO`
Linux	`DCNM_To_NDFC_Upgrade_Tool_LIN_WIN.zip`
Windows	`DCNM_To_NDFC_Upgrade_Tool_LIN_WIN.zip`

Step 4 Save the appropriate **Upgrade Tool** to the 11.5(x) server using **sysadmin** credentials.

Back Up Configuration Using Upgrade Tool

Stop Performance Management collection before running backup script for large scaled DCNM. To stop the Performance Management collection, perform the following steps:

- Navigate to **Administration** > **DCNM Server** > **Server Status**.
- Click on **Stop Service** of **Performance Collector** and wait a few seconds.

- Click on the **refresh** icon on the top right to check the status. Make sure it shows **Stopped**.

The backup tool collects last 90 days Performance Management data.

To run the **DCNM_To_NDFC_Upgrade_Tool** to take a backup of all the applications and data on DCNM 11.5, perform the following steps:

Before you begin

- On Cisco DCNM Release 11.5(1), ensure that you validate each fabric before proceeding to take backup. Choose Cisco DCNM **Web UI > Administration > Credentials Management > SAN Credentials**. Select each fabric and click **Validate** to validate credentials before taking backup.
- Ensure that you've copied the appropriate Upgrade Tool to the server of your DCNM 11.5(x) setup.

Procedure

Step 1 Log on to the Cisco DCNM Release 11.5(x) appliance console.

Step 2 Run the following command to create a screen session.

```
dcnm# screen
```

This creates a session which allows you to execute the commands. The commands continue to run even when the window is not visible or if you get disconnected.

Step 3 Log on to the /root/ directory, by using the su command.

```
dcnm# su
Enter password: <<enter-password>>
[root@dcnm]#
```

Step 4 Execute the upgrade tool, by using the **./DCNM_To_NDFC_Upgrade_Tool** command.

Ensure that you have enabled execution permissions to the Upgrade tool. Use **chmod +x .** to enable executable permissions.

For OVA/ISO-

```
[root@dcnm]# chmod +x ./DCNM_To_NDFC_Upgrade_Tool_OVA_ISO
[root@dcnm]# ./DCNM_To_NDFC_Upgrade_Tool_OVA_ISO /* for OVA/ISO
```

For Windows/Linux-

```
[root@dcnm]# chmod +x ./DCNM_To_NDFC_Upgrade_Tool_LIN_WIN
root@dcnm]# unzip DCNM_To_NDFC_Upgrade_Tool_LIN_WIN.zip
[root@dcnm-rhel]# cd DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/
[root@dcnm-rhel DCNM_To_NDFC_Upgrade_Tool_LIN_WIN]# ls
DCNMBackup.bat  DCNMBackup.sh  jar
[root@rhel DCNM_To_NDFC_Upgrade_Tool_LIN_WIN]# ./DCNMBackup.sh        /* Enter this command
 for Linux appliance */
OR
[root@rhel DCNM_To_NDFC_Upgrade_Tool_LIN_WIN]# ./DCNMBackup.bat       /* Enter this command
 for Windows appliance */
```

The upgrade tool analysis the DCNM appliance data, and determines whether you can upgrade to Cisco Nexus Dashboard Fabric Controller Release 12.1.3 or not.

Note The backup that is generated by using this tool can be used to restore data on NDFC 12.1.3 only.

Step 5 At the prompt to continue with backup, press **y**.

```
********************************************************************************
Welcome to DCNM-to-NDFC Upgrade Tool for OVA/ISO.
This tool will analyze this system and determine whether you can move to NDFC 12.1.3 or
not.
If upgrade to NDFC 12.1.3 is possible, this tool will create files to be used for performing
 the upgrade.
NOTE: only backup files created by this tool can be used for upgrading, older backup files
 created with 'appmgr backup'
CAN NOT be used for upgrading to NDFC 12.1.3
Thank you!
********************************************************************************

Continue? [y/n]: y

Collect operational data (e.g. PM, EPL)? [y/n]: y

Does this DCNM 11.5(1) have DCNM Tracker feature enabled on any switch on any fabric? [y/n]:
 n
```

Step 6 Enter the encryption key to the backup file.

Note You must provide this encryption key when you're restoring the backup file. Ensure that you save the encryption key in a safe location. If you loose the encryption key, you cannot restore the backup.

```
Sensitive information will be encrypted using an encryption key.
This encryption key will have to be provided when restoring the backup file generated by
this tool.

Please enter the encryption key:      /* enter the encryption key for the backup file */
Enter it again for verification:      /* re-enter the encryption key for the backup file
*/

...
...
Creating backup file
Done.
Backup file: backup11_dcnm-172-23-87-224_20210928-093355.tar.gz       /* backup file name*/
[root@dcnm]#
```

The encrypted backup file is created.

Step 7 Copy the backup file to a safe location and shut down the application 11.5(x) DCNM appliance.

Example

Example for taking backup using the DCNM backup Tool

- **Taking backup on DCNM 11.5(x) OVA/ISO appliance**

```
[root@dcnm]# chmod +x DCNM_To_NDFC_Upgrade_Tool_OVA_ISO
[root@dcnm]# ./DCNM_To_NDFC_Upgrade_Tool_OVA_ISO
********************************************************************************

Welcome to DCNM-to-NDFC Upgrade Tool for OVA/ISO.

This tool will analyze this system and determine whether you can move to
NDFC 12.1.3 or not.
```

```
If upgrade to NDFC 12.1.3 is possible, this tool will create files
to be used for performing the upgrade.

NOTE:
only backup files created by this tool can be used for upgrading,
older backup files created with 'appmgr backup' CAN NOT be used
for upgrading to NDFC 12.1.3

Thank you!

********************************************************************************

Continue? [y/n]: y

Collect operational data (e.g. PM, EPL)? [y/n]: y

Does this DCNM 11.5(1) have DCNM Tracker feature enabled on any switch on any fabric?
[y/n]: n

Sensitive information will be encrypted using an encryption key.
This encryption key will have to be provided when restoring
the backup file generated by this tool.

Please enter the encryption key:      /* enter the encryption key for the backup file
*/
Enter it again for verification:      /* re-enter the encryption key for the backup
file */

Adding backup header
Collecting DB table data
Collecting DB sequence data
Collecting stored credentials
Collecting Custom Templates
Collecting CC files
Collecting L4-7-service data
Collecting CVisualizer data
Collecting EPL data
Collecting PM data - WARNING: this will take a while!
Collecting AFW app info
Decrypting stored credentials
Creating backup file
Done.
Backup file: backup11_dcnm-172-23-87-224_20210913-012857.tar.gz        /* backup file
name*/
[root@dcnm]#
```

- **Taking backup on DCNM 11.5(x) Windows/Linux appliance**

```
[root@dcnm]# chmod +x DCNM_To_NDFC_Upgrade_Tool_LIN_WIN
[root@dcnm]# unzip DCNM_To_NDFC_Upgrade_Tool_LIN_WIN.zip
Archive:  DCNM_To_NDFC_Upgrade_Tool_LIN_WIN.zip
    creating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/
    creating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/
   inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/bcprov-jdk15on-1.68.jar
   inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/DCNMBackup.java
   inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/sequences.info.oracle
   inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/slf4j-simple-1.7.21.jar
   inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/jnm.jar
  inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/not-going-to-be-commons-ssl-0.3.20.jar

   inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/tables.info.postgres
   inflating:
DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/jarchivelib-0.7.1-jar-with-dependencies.jar
```

```
   inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/tables.info.oracle
   inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/sequences.info.postgres
   inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/log4j.properties
   inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/DCNMBackup.sh
   inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/DCNMBackup.bat

[root@dcnm-rhel]# cd DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/
[root@dcnm-rhel DCNM_To_NDFC_Upgrade_Tool_LIN_WIN]# ls
DCNMBackup.bat  DCNMBackup.sh  jar
[root@rhel DCNM_To_NDFC_Upgrade_Tool_LIN_WIN]# ./DCNMBackup.sh        /* Enter this
command for Linux appliance */
OR
[root@rhel DCNM_To_NDFC_Upgrade_Tool_LIN_WIN]# ./DCNMBackup.bat        /* Enter this
command for Windows appliance */

Enter DCNM root directory [/usr/local/cisco/dcm]:

Initializing, please wait...

Note: ./jar/DCNMBackup.java uses unchecked or unsafe operations.
Note: Recompile with -Xlint:unchecked for details.
********************************************************************************

Welcome to DCNM-to-NDFC Upgrade Tool for Linux/Windows.

This tool will analyze this system and determine whether you can move to NDFC 12.1.3
or not.

If upgrade to NDFC 12.1.3 is possible, this tool will create files to be used for
performing the upgrade.

Thank you!

********************************************************************************

This tool will backup config data. Exporting Operational data like Performance(PM) might
 take some time.

Do you want to export operational data also? [y/N]: y
********************************************************************************

Sensitive information will be encrypted using an encryption key.
This encryption key will have to be provided when restoring the backup file generated
by this tool.

Please enter the encryption key:       /* enter the encryption key for the backup file
 */
Enter it again for verification:       /* re-enter the encryption key for the backup
file */
2021-09-13 14:36:31 INFO  DCNMBackup:223 - Inside init() method
2021-09-13 14:36:31 INFO  DCNMBackup:245 - Loading properties....
2021-09-13 14:36:31 INFO  DCNMBackup:301 - Inside checkLANSwitches...
2021-09-13 14:36:32 INFO  DCNMBackup:315 - LAN Switch count: 0
2021-09-13 14:36:32 INFO  DCNMBackup:342 - Inside exportDBTables...
2021-09-13 14:36:32 INFO  DCNMBackup:358 - Exporting ---------> statistics
2021-09-13 14:36:32 INFO  DCNMBackup:358 - Exporting ---------> sequence
...
...
...
2021-09-13 14:49:48 INFO  DCNMBackup:1760 - ###### Total time to export Hourly data:
42 seconds.


2021-09-13 14:49:48 INFO  DCNMBackup:1767 - Exporting SanPort Daily entries.
```

```
2021-09-13 14:49:48 INFO  DCNMBackup:1768 - Total number of ports: 455
2021-09-13 14:49:48 INFO  DCNMBackup:1769 - This might take a while, please wait...
2021-09-13 14:50:23 INFO  DCNMBackup:1791 - Total number of Json data entries in
backup/es/pmdb_sanportratedata_daily.data ==> 13751
2021-09-13 14:50:23 INFO  DCNMBackup:1795 - ###### Total time to export Daily data: 34
 seconds.


2021-09-13 14:50:23 INFO  DCNMBackup:1535 - ###### Total time to export PM data: 81
seconds.


2021-09-13 14:50:23 INFO  DCNMBackup:879 - Creating final tar.gz file....
2021-09-13 14:50:30 INFO  DCNMBackup:892 - Final tar.gz elapsed time: 7049 in ms
2021-09-13 14:50:30 INFO  DCNMBackup:893 - Backup done.
2021-09-13 14:50:30 INFO  DCNMBackup:894 - Log file: backup.log
2021-09-13 14:50:30 INFO  DCNMBackup:895 - Backup file:
backup11_rhel77-160_20210913-149215.tar.gz        /* backup file name*/
[root@rhel DCNM_To_NDFC_Upgrade_Tool_LIN_WIN]#
```

Upgrading from Cisco DCNM Release 11.5(4) to Cisco NDFC Release 12.1.3

To upgrade to Cisco Nexus Dashboard Fabric Controller Release 12.1.3 from DCNM Release 11.5(4), perform the following steps:

Before you begin

- Ensure that you've access to the Backup file created from 11.5(4) appliance. For instructions to take backup of all the applications and data on DCNM 11.5(4), see Back Up Configuration Using Upgrade Tool, on page 38.

Note If you do not have the encryption key, you cannot restore from the backup file.

- Ensure that you've installed the required form factor of Cisco Nexus Dashboard. For instructions, refer to Cisco Nexus Dashboard Deployment Guide.
- Ensure that you've installed a fresh installation of Cisco NDFC. For instructions to install Cisco NDFC, refer to:
  - Installing NDFC Manually, on page 25
  - Installing NDFC Using App Store, on page 24
- If your existing configuration used smart licensing with direct connectivity to Cisco Smart Software Management (CSSM), you must ensure that your Nexus Dashboard has the routes required to reach the CSSM website.

  Ensure that subnets for IP addresses on `https://smartreceiver.cisco.com` are added to the route table in the Nexus Dashboard running NDFC. Navigate to **Admin** > **System Settings** > **Routes**. Click **Edit** in the **Management Network Routes** area and add the necessary IP addresses/subnets, then click **Save** to confirm.

  You can ping https://smartreceiver.cisco.com to find the most recent subnet. For example:

```
$ ping smartreceiver.cisco.com
PING smartreceiver.cisco.com (146.112.59.81): 56 data bytes
64 bytes from 146.112.59.81: icmp_seq=0 ttl=52 time=48.661 ms
64 bytes from 146.112.59.81: icmp_seq=1 ttl=52 time=44.730 ms
64 bytes from 146.112.59.81: icmp_seq=2 ttl=52 time=48.188 ms
```

In addition, because NDFC is considered a new product instance, you must re-establish trust. If you took the backup with an expired Trust Token, you must manually run the Smart Licensing Configuration wizard and enter a valid token after the upgrade.

Procedure

Step 1 Log on to Cisco **Nexus Dashboard Web UI** using correct credentials.

Step 2 From the One View drop-down list, select Nexus Dashboard Fabric Controller.

On the Nexus Dashboard Fabric Controller Web UI, **Feature Management** screen is displayed.

Note that none of the personas are selected on the freshly installed Nexus Dashboard Fabric Controller.

Step 3 Click **Restore**.

The **Operations > Backup & Restore** window opens.

Step 4 Click **Restore**.

The **Restore now** window appears.

Step 5 Under **Type**, select your desired format to restore.

Note Select **Config only** or **Full** based on the backup that was created on DCNM Release 11.5(4).

- Choose **Config only** to restore only configuration data.

  You can choose either **Config only** or **Full** backup files.

- Choose **Full** to restore all previous version data to this application.

  You must choose **Full** backup files.

Step 6 Choose the appropriate destination where you have stored the backup file.

- Choose **Upload File** if the file is stored in a local directory.
  - **a.** Open the directory where you've saved the backup file.
  - **b.** Drag and drop the backup file to the **Restore now** window

    or

    Click **Browse**. Navigate to the directory where you've saved the backup file. Select the backup file and click **Open**.
  - **c.** Enter the **Encryption Key** to the backup file.

- Choose **Import from SCP server** or **Import from SFTP server**if the backup file is stored in a remote directory.
  - **a.** In the **Server** field, provide the server IP Address.

b. In the **File Path** field, provide the relative file path to the backup file.

c. In the **Username** and **Password** fields, enter appropriate details.

d. In the **Encryption Key** field, enter the Encryption Key to the backup file.

Step 7 Leave **Ignore External Service IP Configuration** unchecked.

This option is not used during the upgrade.

Step 8 Click **Restore**.

A progress bar appears showing the completed percentage and the description of the operation. The Web UI is locked while the upgrade is in progress. After the restore is complete, the backup file appears in the table on **Backup & Restore** screen. The time required to restore depends on the data in the backup file.

Note An error appears if you've not allocated with IP pool addresses on the Cisco Nexus Dashboard. For more information, refer to *Cluster Configuration* section in Cisco Nexus Dashboard User Guide.

After successful restoration, a notification banner appears as below:

```
Reload the page to see latest changes.
```

Click **Reload the page**, or refresh the browser page to complete restore and begin using you Cisco Nexus Dashboard Fabric Controller Web UI.

Post Upgrade Tasks

The following sections describe the tasks that must be performed post upgrading to Cisco NDFC, Release 12.1.3.

Post Upgrade tasks for SAN Controller

After restoring the data from backup, all the server-smart licenses are **OutofCompliance**.

To migrate to Smart Licensing using Policy, launch Nexus Dashboard Fabric Controller. On the Web UI, choose **Operations > License Management > Smart** tab. Establish trust with CCSM using SLP. For instructions, refer to *License Management* chapter in *Cisco Nexus Dashboard Fabric Controller Configuration Guides*.

Post Upgrade tasks for Fabric Controller

The following features are not carried over when you upgrade from DCNM 11.5(x) to Cisco NDFC 12.1.3:

- Endpoint Locator must be reconfigured
- IPAM Integration must be reconfigured
- Alarm Policies must be reconfigured
- Custom topologies must be recreated and saved
- PM collection must be re-enabled on fabrics

- Switch images must be uploaded

Managing Trap IP on Nexus Dashboard and Nexus Dashboard Fabric Controller

Deployment Type in Release 11.5(x)	In 11.5(x), trap IP address is collected from	LAN Device Management Connectivity	In 12.1.3, trap IP address belongs to	Result
LAN Fabric Media Controller	eth1 (or vip1 for HA systems)	Management	Belongs to Management subnet	Honored There is no configuration difference. No further action required.
LAN Fabric Media Controller	eth0 (or vip0 for HA systems)	Management	Does not belong to Management subnet	Ignored, another IP from the Management pool will be used as trap IP. Configuration difference is created. On the **Web UI > LAN > Fabrics > Fabrics**, double click on the Fabric to view **Fabric Overview**. From **Fabrics Actions** drop-down list, select **Recalculate Config**. Click **Deploy Config**.
LAN Fabric Media Controller	eth0 (or vip0 for HA systems)	Data	Belongs to Data subnet	Honored There is no configuration difference. No further action required.

Deployment Type in Release 11.5(x)	In 11.5(x), trap IP address is collected from	LAN Device Management Connectivity	In 12.1.3, trap IP address belongs to	Result
LAN Fabric Media Controller	eth0 (or vip0 for HA systems)	Data	Does not belong to Data subnet	Ignored, another IP from the Data pool will be used as trap IP. Configuration difference is created. On the **Web UI > LAN > Fabrics > Fabrics**, double click on the Fabric to view **Fabric Overview**. From **Fabrics Actions** drop-down list, select **Recalculate Config**. Click **Deploy Config**.
SAN Management	OVA/ISO – • trap.registaddress (if set) • eth0 (if trap.registaddress is not set) Windows/Linux – • trap.registaddress (if set) • Interface based on event-manager algorithm (if trap.registaddress is not set	Not applicable	Belongs to Data subnet	Honored There is no configuration difference. No further action required.
		Not applicable	Does not belong to Data subnet	Ignored, another IP from the Data pool will be used as trap IP.

Feature Management

After restoring the backup, based on the type of deployment, Nexus Dashboard Fabric Controller Release 12.1.3 is deployed with one of the following personas:

- Fabric Controller
- SAN Controller

The status on the Feature Management changes to **Starting**. Additionally, you can select the features that you want to enable. Check the **Feature** check box and click **Save & Continue**.

Changing across Feature-Set

Nexus Dashboard Fabric Controller 12 allows you to switch from one feature set to another. Choose **Settings > Feature Management**. Select the desired feature set and applications in the table below. Click **Save & Continue**. Refresh the browser to begin using Cisco Nexus Dashboard Fabric Controller with the new feature set and applications.

There are a few features/applications supported with specific deployments. When you change the feature set, some of these features are not supported in the new deployment. The following table provides details about the pre-requisites and criteria based on which you can change the feature set.

Table 5: Supported Switching between deployments

From/To	Fabric Discovery	Fabric Controller	SAN Controller
Fabric Discovery	-	Only monitor mode fabric is supported in Fabric Discovery deployment. When you change the feature set, the fabric can be used in the Fabric Controller deployment.	Not supported
Fabric Controller	You must delete the existing fabrics before changing the fabric set.	If you're changing from Easy Fabric to IPFM fabric application, you must delete the exiting fabrics.	Not supported
SAN Controller	Not supported	Not supported	-