Cisco Nexus Dashboard Fabric Controller (NDFC) Deployment Guide

Introduction

Cisco Nexus Dashboard Fabric Controller aka NDFC (formerly known as Data Center Network Manager aka DCNM) runs exclusively as an application service on top of the Cisco Nexus Dashboard Cluster. Nexus Dashboard cluster uses Kubernetes at its core with customized extensions, thereby realizing secure and scaled-out platform for deployment of microservices-based applications. Nexus Dashboard Cluster provides Active/Active HA (High Availability) for all applications running on top of that cluster.

NDFC 12.0.1a release introduces an architecture revamp from the prior pseudo monolithic/microservices architecture of the DCNM 11 release to a complete cloud-native microservices-based implementation. Everything is based on well-known K8s constructs of services, pods, namespaces, and containers which are instantiated across the Nexus Dashboard cluster. NDFC also adopts the Nexus Dashboard Blueprint UI/UX framework and components that in turn provide a seamless and symmetric look and feel across the entire Cloud Networking Product Portfolio.

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

Various personas of DCNM namely LAN, SAN, IP Fabric for Media are implemented as a feature-set that can be dynamically enabled at run time post enablement of the NDFC application. Recall that in DCNM 11, this persona selection is an install-time choice that you must make. In addition, with NDFC, as users enable different feature sets and sub-features via a newly introduced feature management workflow, appropriate sets of microservices are spawned with dynamic menu reload options. You can add extra nodes to the Nexus Dashboard cluster for realizing higher scale. As a result, appropriate microservices in NDFC spawns more replicas thereby allowing the controller to support the additional load.

You can deploy NDFC on either a Virtual Nexus Dashboard cluster, also known as vND, or Physical Nexus Dashboard cluster, also known as pND. In either case, NDFC, being a native microservices-based application, supports true scale-out where system scale can be increased simply by adding extra nodes to the Nexus Dashboard cluster. The system requirements and qualified scale support depends on the Nexus Dashboard deployment model. Refer to Networking Requirements to validate NDFC verified scale information.

Networking with Nexus Dashboard

NDFC being an application that runs on top of Cisco Nexus Dashboard cluster, uses the networking interfaces of the Nexus Dashboard to manage, automate, configure, maintain, and monitor the Cisco Nexus and MDS family of switches. In this section, we will briefly review the networking for the Nexus Dashboard cluster.

Each Nexus Dashboard node in a cluster has two interfaces:

- Management Interface and
- Data (also known as Fabric) Interface.

Hence, during deployment of the Nexus Dashboard cluster, you must provide two IP addresses for each node that is going to be part of the Nexus Dashboard cluster. At the time of deployment of the Nexus

Dashboard cluster, a user can make a choice on whether they want to deploy a single node Nexus Dashboard cluster or a 3 node Nexus Dashboard cluster. Single node Nexus Dashboard cluster deployments support NDFC IP Fabric for Media and SAN Controller production deployments and a LAN Controller lab-deployment (<=25 switches). A minimum of three Nexus Dashboard nodes is required for all production NDFC LAN Controller deployments.

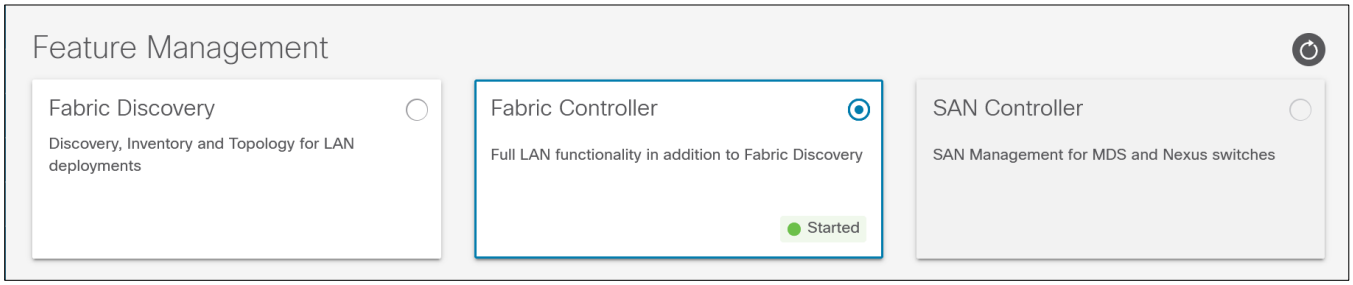

Figure 1. Feature Manager

As the name implies, Nexus Dashboard management interface connects to the management network that typically is used for web/API access to the Nexus Dashboard cluster. Nexus Dashboard data interface typically provides IP reachability to the physical data center network infrastructure.

This section describes the purpose and functionality of the networks as they are used by the Nexus Dashboard services.

Management Network

The management network is used for:

- Accessing the Nexus Dashboard GUI (Graphical User Interface)
- Accessing the Nexus Dashboard CLI (command-line interface) via SSH
- DNS (Domain Name System) and NTP (Network Time Protocol) communication
- Nexus Dashboard firmware upload
- Cisco DC App Center (AppStore) If you want to use the Nexus Dashboard App Store to install applications
- Intersight device connector

Data Network

The data network is used for:

- Nexus Dashboard clustering
- Application to application communication (SMTP and SNMP forwarding)

Networking Requirements

- Two Logical Interfaces per Nexus Dashboard Node.
  - bond1br (also known as Nexus Dashboard Management Interface)
  - bond0br (also known as Nexus Dashboard Data Interface)

- For enabling NDFC on a Nexus Dashboard cluster, the Nexus Dashboard Management, and Data Interfaces must be in different subnets. Therefore, a minimum of two IP subnets are required for deployment of such a cluster.
- Configure Nexus Dashboard Data Interface by clicking on "**Edit**" button.

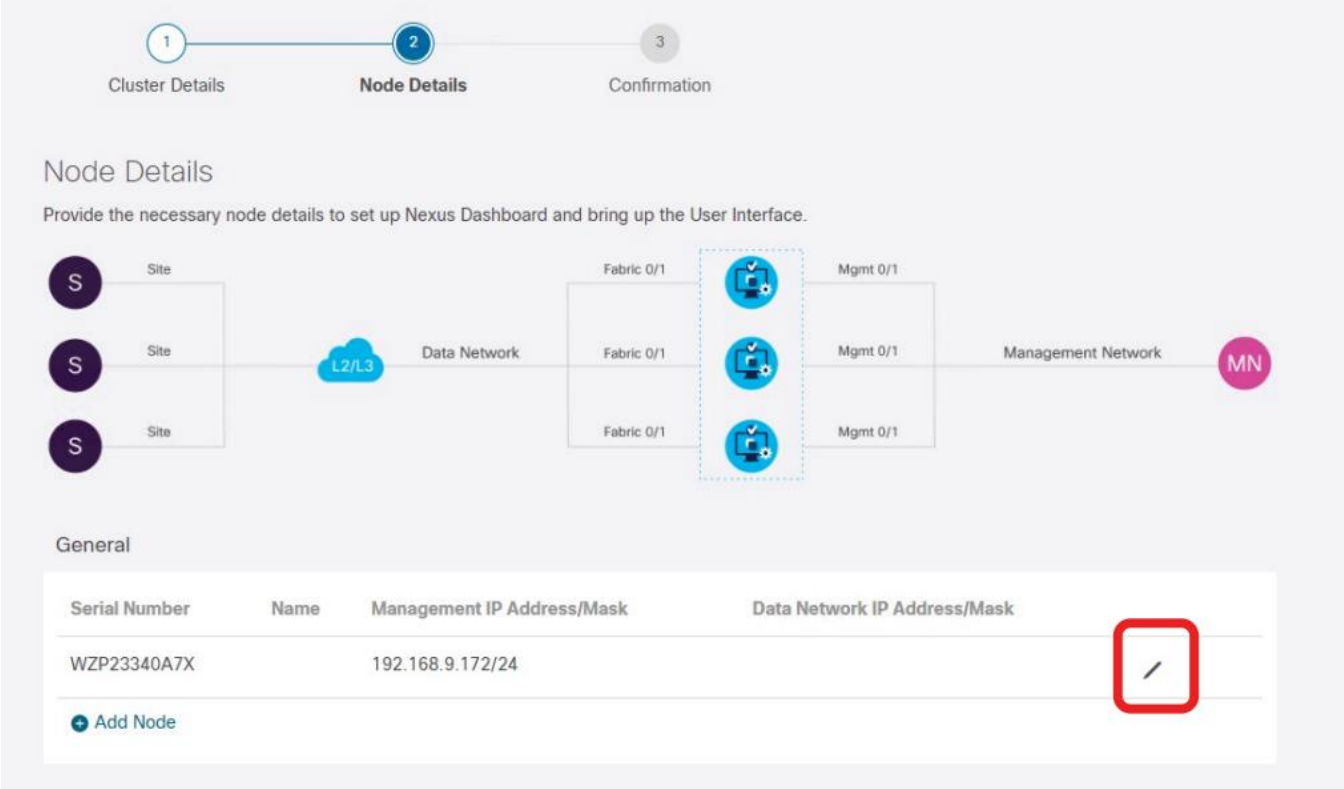

Figure 2. Nexus Dashboard Web Installer

- Enter Nexus Dashboard Data network and default GW for IP access to NDFC in-band and VLAN-ID if connected via a trunk interface.

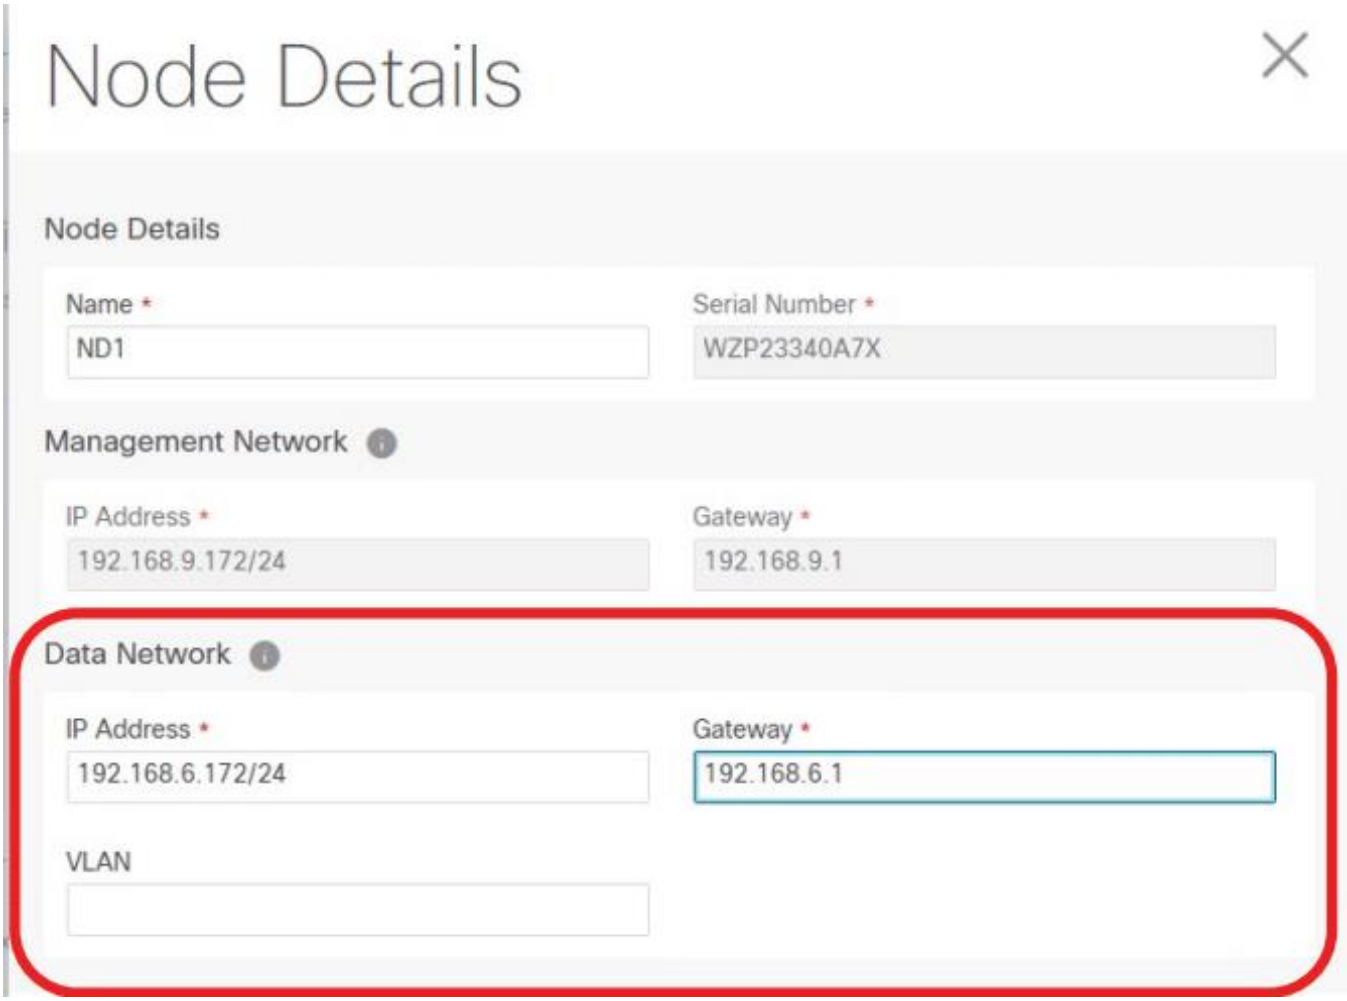

Figure 3. Nexus Dashboard Data Interface

- In NDFC Release 12.0.x, the Nexus Dashboard nodes hosting the NDFC are required to be Layer-2 adjacent. Beginning with Release 12.1.1e, NDFC on Nexus Dashboard Release 2.2.1h allows you to configure nodes within the cluster with Layer-2 or Layer-3 adjacency. For more information on Layer-3 Reachability between Cluster Nodes, see Layer-3 Reachability between Cluster Nodes.
- Typically, NDFC manages and automates Nexus switches via IP reachability to the Management 0 (also known as mgmt0) interface of the switches. This is referred to as Out-of-Band or OOB management of the switches. Switch OOB reachability from NDFC can be via the Nexus Dashboard management interface (default*) or via the Nexus Dashboard data interface. As discussed further in the following section, the user can specify, via configuration, what interface to use for this type of communication. If switches are managed by NDFC via the switch front-panel port (SVI, loopback or equivalent), it is referred to as In-band management. Switch In-band reachability from NDFC must be via the Nexus Dashboard data interface.
- In general, the default route of Nexus Dashboard services is via the Data Interface. All the NDFC application pods use Nexus Data Interface as the default route. An operator must add static routes

on the Nexus Dashboard for forcing connectivity out of Management Interface. This is done via the Nexus Dashboard Cluster Configuration workflow available on the Nexus Dashboard Admin Console.

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

Application	Connectivity	Maximum RTT
Nexus Dashboard Fabric Controller	Between Nodes	50 ms
	To Switches	50 ms

NDFC to Device Connectivity

As stated in the Networking requirements section, a switch can be managed via its mgmt0 interface, also known as OOB management, with NDFC to switch reachability, either via the Nexus Dashboard Management Interface or Data interface. This section includes different use-cases when a switch is discovered and managed by NDFC.

When the Nexus Dashboard Cluster is online for the first time, you'll see the following routing table. Note that the IP addresses may be different based on the specific cluster deployment.

```
[rescue-user@ndfc-12]# ip rule show
0:from all lookup local
32763:from 172.17.0.0/16 lookup 100
32764:from 172.25.74.144 lookup 1001
32765:from 192.168.100.5 lookup 1000
32766:from all lookup main
32767:from all lookup default
```

Note: By default, only **rescue-user** is enabled for local SSH to Nexus Dashboard CLI. Root access is denied and can be used only by Cisco TAC (Technical Assistance Center).

Table 100 - used for Inter-POD communication on Nexus Dashboard Kubernetes Architecture.

Table 1001 - used for Nexus Dashboard Management Network specified during the system installation.

Table 1000: This is used for Nexus Dashboard Fabric (also known as Data) Network which is specified during the system UI bootstrap.

```
[rescue-user @ndfc-12]# ip route show table 100
default via 192.168.100.254 dev bond0br
172.17.0.0/16 dev k8br0 scope link
172.25.74.0/23 dev bond1br scope link
192.168.100.0/24 dev bond0br scope link
```

As shown in Table 100, by default, PODs use bond0br (also known as Nexus Dashboard Fabric/Data Interface) as the Next-Hop for any Default Routing. If Nexus Dashboard to Switch mgmt0 communication is desired via the Nexus Dashboard Management interface, an operator must add those specific routes so that the switches can be reached via bond1br (also known as Nexus Dashboard Management Interface).

Below are some specific examples of how to enable communication between Nexus Dashboard and the Switch mgmt0 interface:

Use-Case 1

Nexus Dashboard Interfaces are Layer-2 Adjacent to Switch Mgmt0 Interface

Case 1

Nexus Dashboard management IP = 172.25.74.144/24, Nexus Dashboard data IP = 192.168.100.10/24. Switch mgmt0 IP = 172.25.74.10/24

In this case, the switch management 0 IP address and Nexus Dashboard management IP address are part of the same subnet. Therefore, NDFC onboards the switch using the management interface. As it is Layer-2 in nature, it picks the interface with the matching subnet.

Nexus Dashboard Management Interface is used for switch discovery, monitoring, configuration deployments, image management, out-of-band POAP, and PMN/PTP telemetry. Nexus Dashboard Data Interface is used for inband related features, such as Endpoint Locator.

To enable Layer-2 adjacency, define the persistent IP addresses used for SCP-POAP and SNMP-trap services in the Nexus Dashboard Management subnet. For more information, see Persistent IP Requirements for NDFC.

Case 2

Nexus Dashboard management IP = 172.25.74.144/24, Nexus Dashboard data IP = 192.168.100.10/24. Switch mgmt0 IP = 192.168.100.100/24

In this case, the switch management 0 IP and Nexus Dashboard Data IP are part of the same subnet. Therefore, NDFC onboards the switch using the data interface. As it is Layer-2 in nature, it will pick the interface with the matching subnet.

Nexus Dashboard Data Interface is used for switches out-of-band discovery, monitoring, configuration deployments, image management, and POAP (out-of-band). Nexus Dashboard Data Interface is also used for Endpoint Locator and PMN/PTP Telemetry.

To enable this use-case, do the following tasks:

1. Define the persistent IP addresses used for SCP-POAP and SNMP-trap services in the Nexus Dashboard Management subnet. For more information, see Persistent IP Requirements for NDFC.
2. Change the NDFC global Server Settings **LAN Device Management Connectivity** from “Management” to “**Data**,” as shown in the following image.

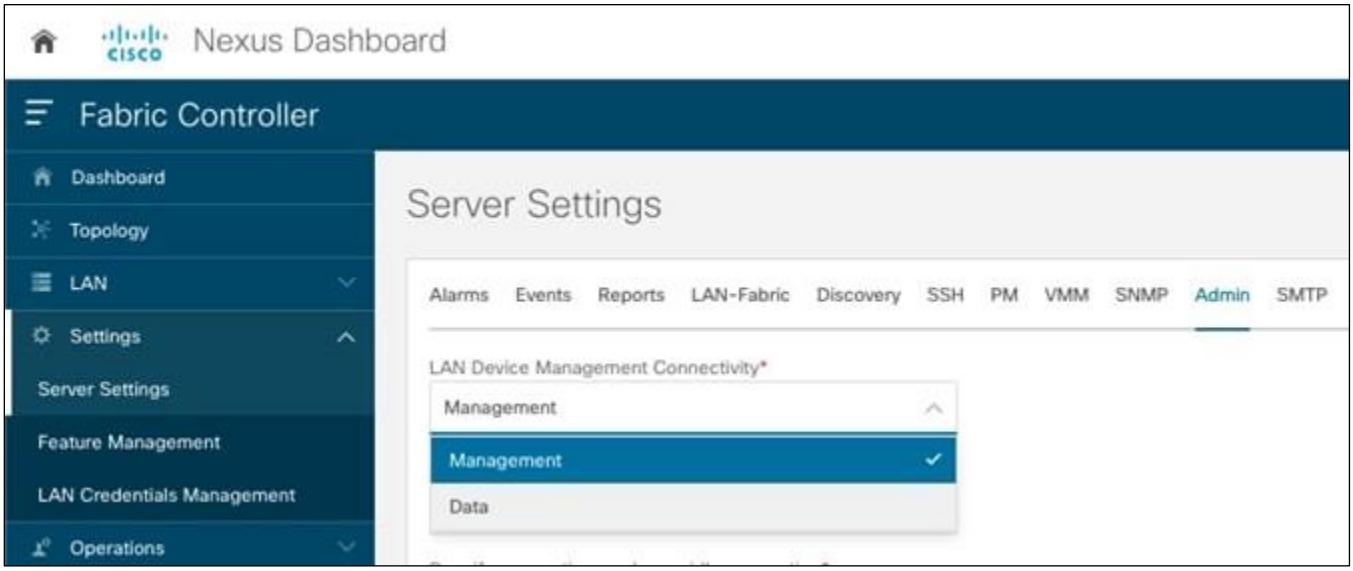

Note: There is a special configuration required for performing Image management on switches mgmt0 using the Nexus Dashboard Data interface.

3. By default, when NDFC LAN Device Management Connectivity is set to Data, the SCP-POAP, and SNMP-trap service pods are spawned with persistent IPs associated with the Nexus Dashboard data interface subnet. For operations related to any file copies to/from the NDFC from/to the switch, NDFC uses default VRF. The switch has IP reachability to the Nexus Dashboard data interface over the default VRF. This includes image copy operations required for Image and package management. In this case, NDFC utilizes the mgmt0 interface VRF, also known as management VRF for such operations. Therefore, the following options are supported:
4. If switches are not added into the NDFC, then perform Step-3a but if switches are added/already discovered and present in the fabric then perform Step-3b

   3a. Enable NDFC server settings. On the NDFC Web UI, choose **LAN > Settings > Server Settings > Discovery** tab. Enable the setting as shown in the following image.

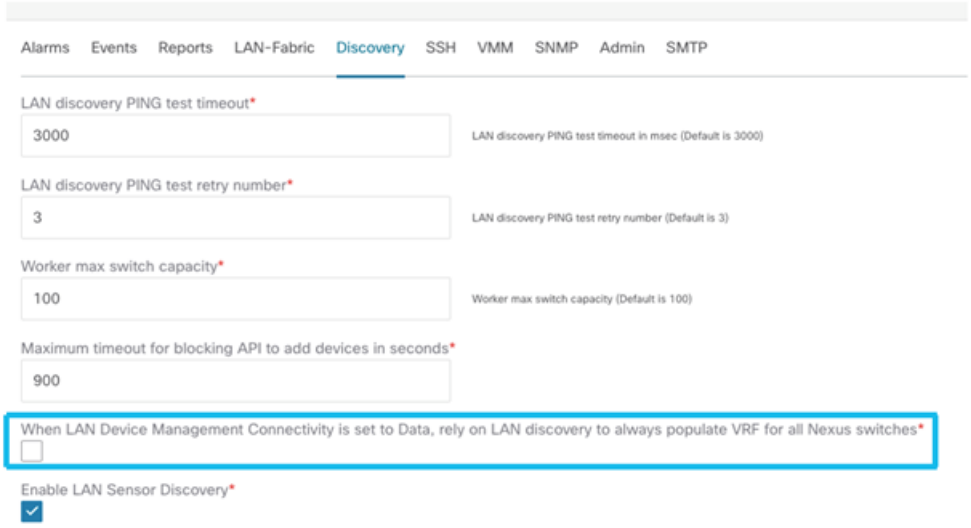

3b. Update discovery VRF per switch. Choose **LAN > Switches**. Double click on the switch to open the **Switch Overview** screen, From the **Actions** drop-down list, choose **Discovery > Update VRF**.

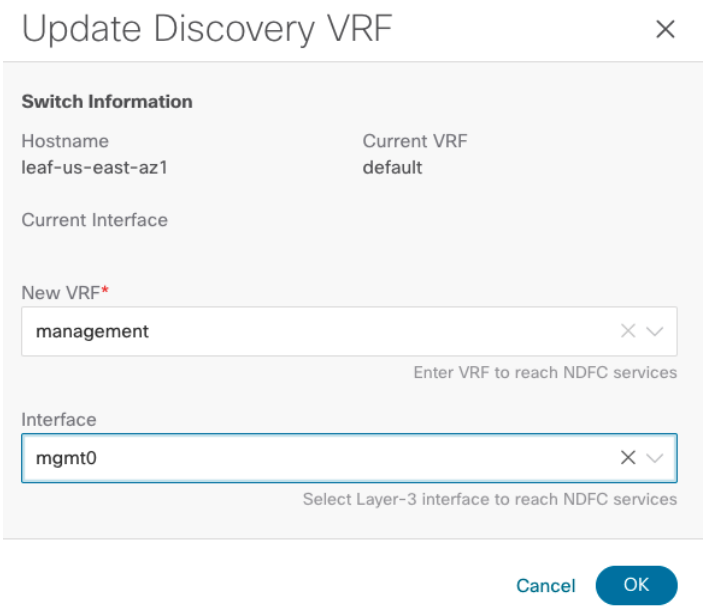

By enabling the above options, the VRF associated with the interface that has the discovery IP for a switch will be auto-discovered by NDFC during switch import. The option to override the VRF for a particular switch is always available and if set by the user, it is always honored.

Use-Case 2

Nexus Dashboard Interfaces are Layer-3 Adjacent to Switch Mgmt0 Interface

Case 3

Nexus Dashboard management IP = 172.25.74.144/24, Nexus Dashboard data IP = 192.168.100.10/24. Switch mgmt0 IP = 10.23.234.192/16

Now, Switch management 0, also known as mgmt0, IP address is not part of subnets associated with either Nexus Dashboard interface. For this typical layer-3 reachability scenario, users must decide which Nexus Dashboard Interface will provide IP reachability to the switches, and therefore, will be used for switch onboarding and discovery.

In this case, the switch mgmt0 interface is reachable via Nexus Dashboard Management Interface.

Nexus Dashboard Management Interface is used for Switch discovery, monitoring, configuration deployments, image management, and out-of-band POAP. Nexus Dashboard Data Interface is used for Endpoint Locator.

To enable this use-case, do the following tasks:

1. Add static route(s) associated with the ND Mgmt interface pointing to the switch mgmt0 subnet(s).

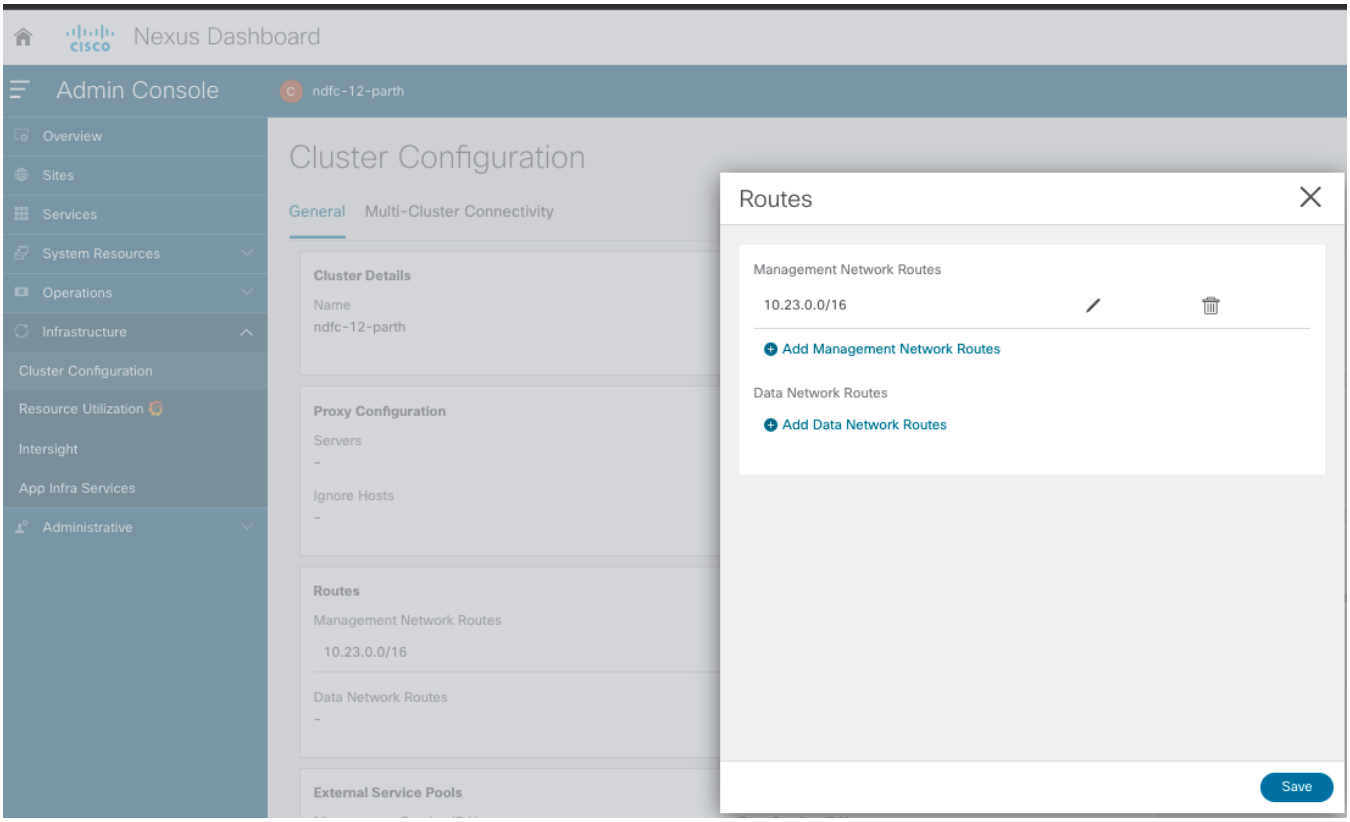

Figure 4. Adding Nexus Dashboard static routes for IP connectivity

Note that using this static route configuration option, the operator can decide whether to use Nexus Dashboard management or data interface for reachability to a particular destination even if the destination is reachable over both interfaces. Now, if we check the updated routing table then we see that a new entry has been added to Table 100.

```
[rescue-user@ndfc-12]# ip route show table 100
default via 192.168.100.254 dev bond0br
10.23.0.0/16 via 172.25.74.1 dev bond1br
172.17.0.0/16 dev k8br0 scope link
172.25.74.0/23 dev bond1br scope link
192.168.100.0/24 dev bond0br scope link
```

Therefore, switch with the IP address of 10.23.234.192 can now be reached via Nexus Dashboard Management Interface and can be successfully imported into the NDFC.

2. Define the persistent IP addresses used for SCP-POAP and SNMP-trap services in the Nexus Dashboard Management subnet. For more information, see Persistent IP Requirements for NDFC.

Recall that the NDFC LAN Device Management Connectivity in the Server Settings is set to Management by default. Consequently, the SCP-POAP and SNMP-trap service pods will be spawned with persistent IPs associated with the Nexus Dashboard management subnet.

Case 4

Nexus Dashboard management IP = 172.25.74.144/24, Nexus Dashboard data IP = 192.168.100.10/24. Switch mgmt0 IP = 10.35.35.35/24

Now, switch management 0, also known as mgmt0, IP address is not part of subnets associated with either Nexus Dashboard interface. For this typical layer-3 reachability scenario, users must decide which Nexus Dashboard Interface to use for switching onboarding.

For this use case, switch Mgmt0 interface is accessed via the ND Data Interface. Switches managed or discovered via the front-panel interfaces such as SVIs, loopbacks, and so on, must always be reachable via the ND data interface. NDFC supports simultaneous support for switches when they are reachable over the switch mgmt0 interface aka Out-of-Band access or reachable via the front-panel ports aka In-Band access.

Nexus Dashboard Data Interface is used for Switches out-of-band and inband discovery, monitoring, configuration deployments, image management, and POAP (both out-of-band and inband). Nexus Dashboard Data Interface is also used for Endpoint Locator if enabled.

To enable this use-case, the configuration steps are similar to that of Case 2. The only addition is to Add static route(s) to the ND Data interface pointing to the switches mgmt0 subnet(s). Notice that this is not required for routing reachability purposes, as the default route associated with the ND Data interface previously shown above would take care of that. However, it is necessary to ensure NDFC POAP can work adequately for touchless Day-0 switch bring-up.

Case 5

Nexus Dashboard management IP = 172.25.74.144/24, Nexus Dashboard data IP = 192.168.100.10/24. Switch mgmt0 IP = 10.55.55.55/24

Now, Switch management 0, also known as mgmt0, IP address is not part of subnets associated with either Nexus Dashboard interface. For this typical layer-3 reachability scenario, users must decide which Nexus Dashboard Interface to use for switching onboarding.

For this use case, you must deploy an NDFC cluster with Layer-3 reachability between the ND nodes. The NDFC **LAN Device Management Connectivity** in **Server Settings** must be set to **Data**. This implies that the SCP-POAP and SNMP-trap service pods have persistent IPs associated with a subnet pool that is associated with the ND data interface. Note that the persistent IPs are not part of either the ND data or management subnet of any of the nodes, as is typically the case for the Layer-2 adjacent ND cluster deployment. They belong to a different pool so the current persistent IP reachability for the pods is dynamically advertised to the physical network via eBGP. For more information, see Layer-3 Reachability between Cluster Nodes.

In this scenario, the Switch Mgmt0 interface has IP reachability from the NDFC via ND Management Interface. However, for image management and SNMP trap purposes, the switch also must have reachability to the ND Data Interface via the front-panel ports typically part of the default VRF. Therefore, NDFC has IP reachability to Switch Mgmt0 interface via ND Management Interface and Switches Front Panel ports via the ND Data Interface. This may be a common scenario when you have physically separate Out-of-Band and Inband networks.

Nexus Dashboard Management Interface is used for switch out-of-band discovery, monitoring, and configuration deployments. Nexus Dashboard Data Interface is used for Image management, Inband management/Inband POAP, and Endpoint Locator.

Note: In this scenario, switch mgmt0 out-of-band POAP is not supported, as SCP/POAP Persistent IP address is associated with ND Data Interface and Switch mgmt0 reachability is only via ND Management Interface.

To enable this use-case, do the following tasks:

1. Add static route(s) associated with the ND Mgmt interface pointing to the switch mgmt0 subnet(s).
2. Define the persistent IP addresses used for SCP-POAP and SNMP-trap services and associate them with the ND Data interface. For more information on the use of persistent IP addresses, see Persistent IP Requirements for NDFC.
3. Change the global server setting for **LAN Device Management Connectivity** to **Data**.

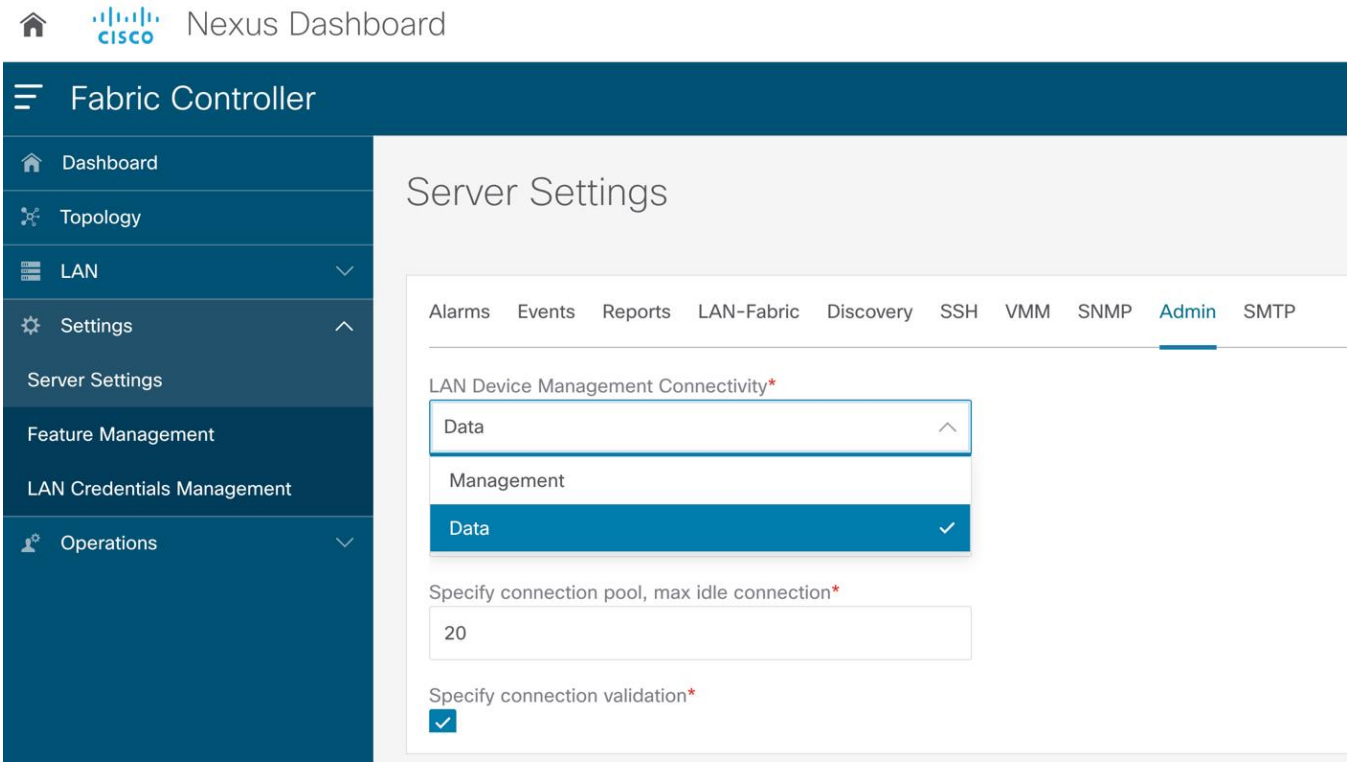

Summary

In summary, the following are the guidelines for managing the switches using NDFC.

- To discover a switch, there must be IP reachability from the NDFC to that switch. The switch can be imported over its mgmt0 IP (called Out-of-Band management) or over another IP associated with Layer-3 interface (called Inband management).
- For Out-of-Band or switch mgmt0 access, NDFC may have reachability over the ND management or ND data interface. To setup reachability over the ND management interface, appropriate static routes must be added on the ND management interface. Static routes are required if reachability is over the ND data interface, when NDFC POAP functionality is desired for touchless Day-0 switch bring up.
- For Inband management, NDFC to switch reachability must be setup over the ND data interface. NDFC Inband POAP functionality is supported only over the ND data interface.

- The **LAN Device Management Connectivity** option in the NDFC **Server Settings** controls whether the SCP-POAP and SNMP trap services are spawned with persistent IPs that are associated with either the ND management or the ND data interface. The default value for this setting is **Management**.
- For Inband management, the NDFC **LAN Device Management Connectivity** option in the NDFC **Server Settings** must be set to **Data**. This means to support both Out-of-Band and In-band management simultaneously, all switches must have IP reachability to the ND data interface.
- When configuring ND cluster with Layer-3 reachability between the ND nodes, the NDFC **LAN Device Management Connectivity** option in the NDFC **Server Settings** must be set to **Data**. This means that even to support Out-of-Band management only, switches must have IP reachability to the ND data interface. The 2 options here are
  - IP reachability to switch mgmt0 interface is over the ND data interface
  - IP reachability to switch mgmt0 interface is over ND management interface

  Note that there is also IP reachability from that switch over the front-panel ports to the ND data interface.
- If POAP is desired over Layer-3, i.e., NDFC is Layer-3 adjacent to switches, provide Nexus Dashboard Node IPs as DHCP Relay Address.
  - If the NDFC **LAN Device Management Connectivity** option in the NDFC **Server Settings** is set to **Management** (default value), set the DHCP Relay Address to the management interface IP (bond1br) of all the Nexus Dashboard Nodes.
  - If the NDFC **LAN Device Management Connectivity** option in the NDFC **Server Settings** is set to **Data**, set the DHCP Relay Address to the data interface IP (bond0br) of all the Nexus Dashboard Nodes.

Deployment Modes and Design for LAN Fabrics

The following sections provide information about deployment modes and design for VXLAN Fabrics. While the example assumes a ND cluster deployment where the ND nodes are layer-2 adjacent to each other, the general ideas are applicable even for a Layer-3 ND cluster deployment.

ND node IP assignment

Figure 5. Nexus Dashboard Interface IP Addresses

Deployment Profiles

This section is applicable for NDFC/ND deployments prior to version 12.1.1e/2.2.1h only. Beginning with NDFC Release 12.1.1e, the profiles are auto-selected by ND, based on user inputs related to the scale of sites, switches, flows, and services that they plan to deploy for a given ND deployment.

For completeness, this section has been presented here with a note that this is only applicable for NDFC 12.0.x deployments. NDFC has multiple profile options. You can choose the profile as shown in the following image.

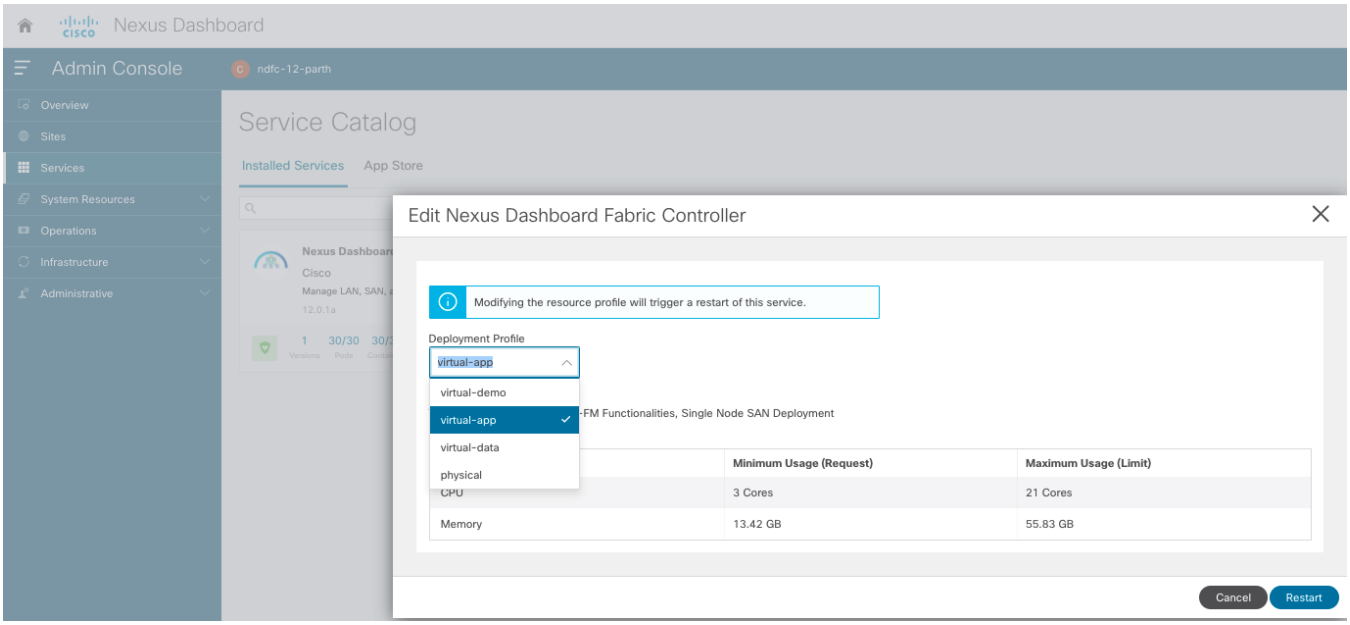

Figure 6. NDFC Profile options

virtual-demo

An operator must Select this profile when deploying Fabric Discovery/Fabric Controller with a Single Master Node.

virtual-app

This profile is auto-selected when you deploy **nd-xx-app.ova**. An operator must select this profile for deploying Fabric Controller with Three Master Node Cluster.

virtual-data

This profile is auto-selected when you deploy **nd-xx-data.ova**. An operator must select this profile for deploying SAN Controller Insights with Single or Three Master Node Cluster.

physical

This profile is auto selected when Nexus Dashboard is deployed as a Physical Appliance. An operator must select this profile for deploying Fabric Controller with Three Master Node Cluster or SAN Controller Insights with Single or 3 Master Node Cluster.

Deploying NDFC on pND

The following figure shows the Nexus Dashboard physical node interfaces.

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

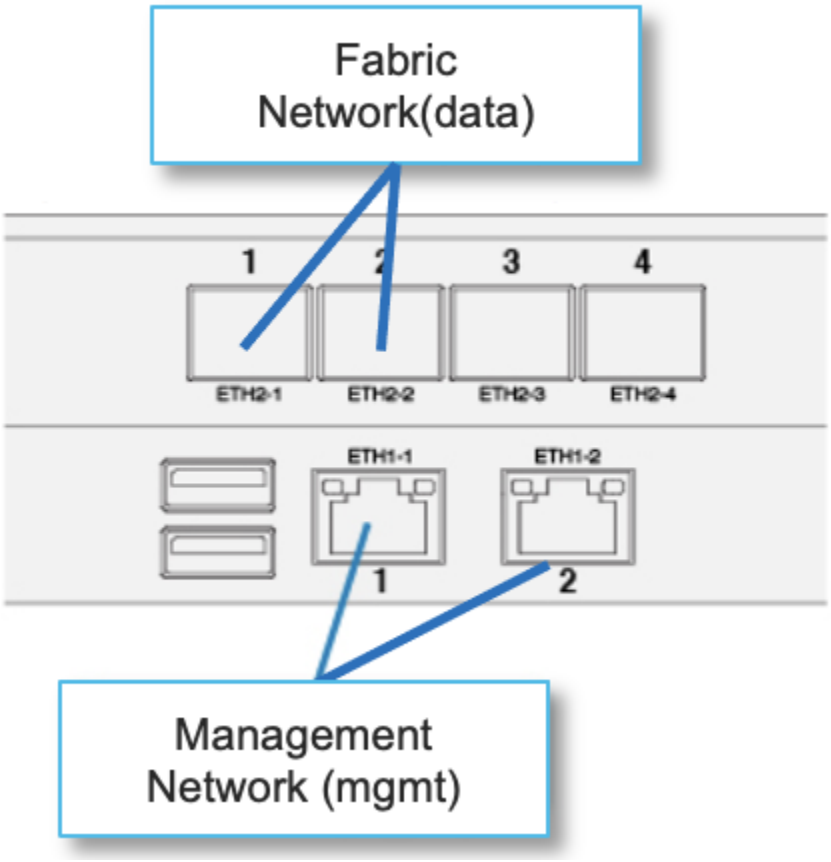

Figure 7. Physical Nexus Dashboard Interface Mapping

The interfaces are configured as Linux bonds: one for the data interfaces and one for the management interfaces. All interfaces must be connected to individual host ports. Port- Channel or vPC are not supported.

Deployment Option 1

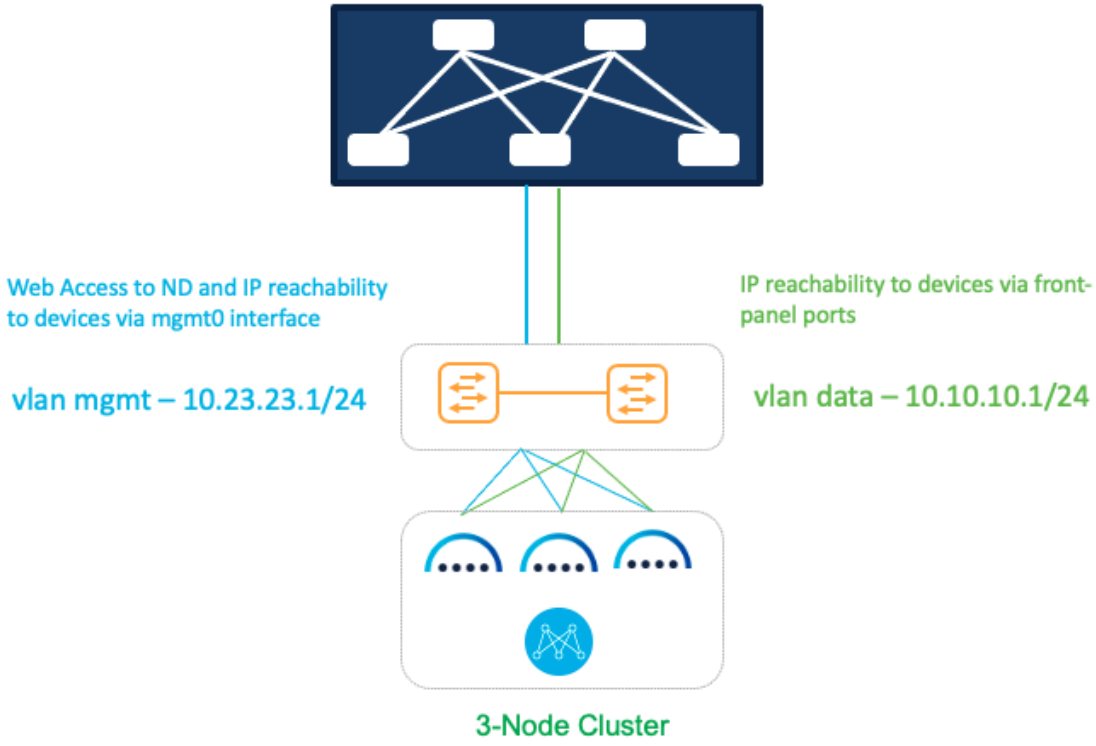

Figure 8. Deploying NDFC on pND Deployment Option 1

In this option, Nexus Dashboard Management and Data interfaces are connected to a network infrastructure that provides reachability to the switches mgmt0 interfaces and the front-panel ports. The ND interfaces are connected to a pair of upstream switches part of such infrastructure.

Sample Configurations

On both uplink switches (marked as yellow) for Nexus Dashboard management-

```
Interface eth1/1, eth1/3, eth1/5
  switchport mode access
  switchport access 23
```

On both uplink switches (marked as yellow) for Nexus Dashboard fabric-

```
Interface eth1/2,eth1/4,eth1/6
  switchport mode access
  switchport access vlan 10
```

OR

```
Interface eth1/2,eth1/4,eth1/6
  switchport mode trunk
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10
```

OR

```
Interface eth1/2,eth1/4,eth1/6
```

```
  switchport mode trunk
  switchport trunk allowed vlan 10
```

For the last option without the trunk native VLAN, an operator provides VLAN ID 10 as VLAN tag during Nexus Dashboard installation and Interface bootstrap as shown in Figure 3 in Networking Requirements section.

Deployment Option 2

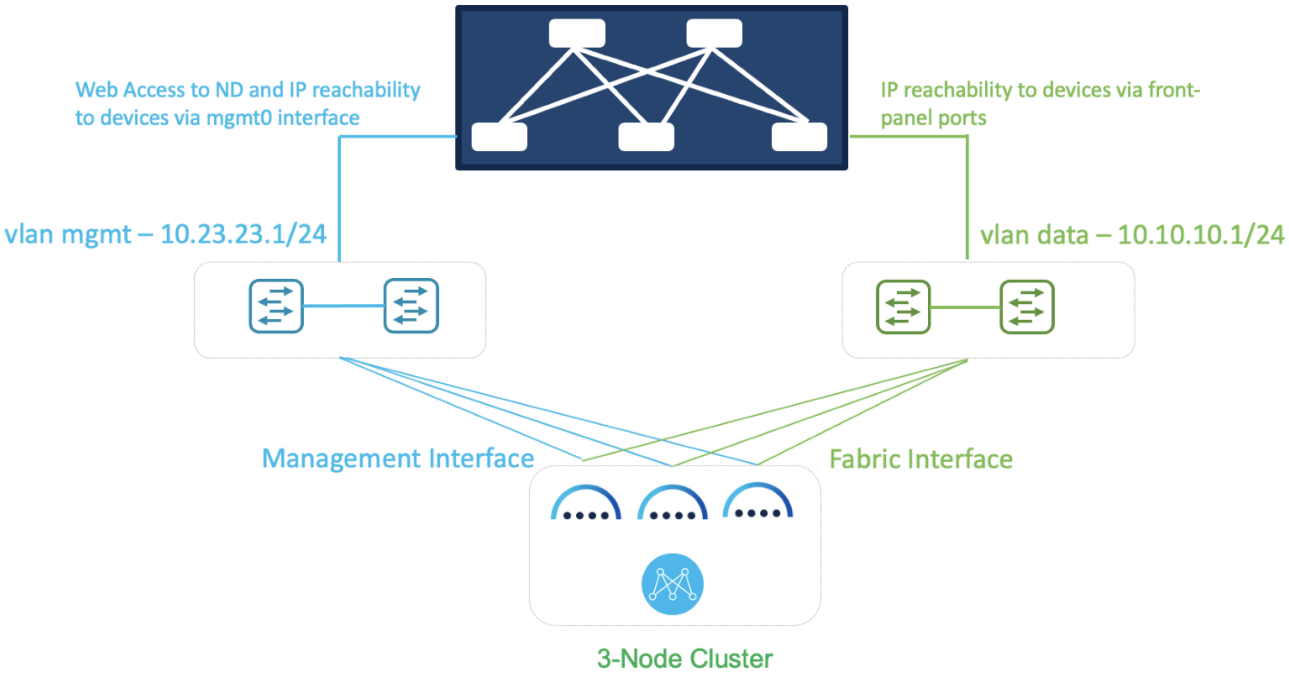

Figure 9. Deploying NDFC on pND Deployment Option 2

In this option, two separate network infrastructures provide access to the switch mgmt0 interfaces and the front-panel ports in this option. Consequently, the ND Management and Data interfaces are connected to those separate networks.

Sample Configurations

On both uplink switches (marked as blue) for Nexus Dashboard management-

```
Interface eth1/1-3
  switchport mode access
  switchport access 23
```

On both uplink switches (marked as green) for Nexus Dashboard fabric-

```
Interface eth1/1-3
  switchport mode access
  switchport access vlan 10
```

OR

```
Interface eth1/1-3
  switchport mode trunk
  switchport trunk native vlan 10
```

```
  switchport trunk allowed vlan 10
```

OR

```
Interface eth1/1-3
  switchport mode trunk
  switchport trunk allowed vlan 10
```

Note: For the last option without the trunk native vlan, an operator provides VLAN ID 10 as VLAN tag during Nexus Dashboard installation and Interface bootstrap as shown in Figure 3 from the Networking Requirements Section.

Deploying NDFC on vND

A vND node can be deployed as an OVA in ESXi with or without a vCenter.

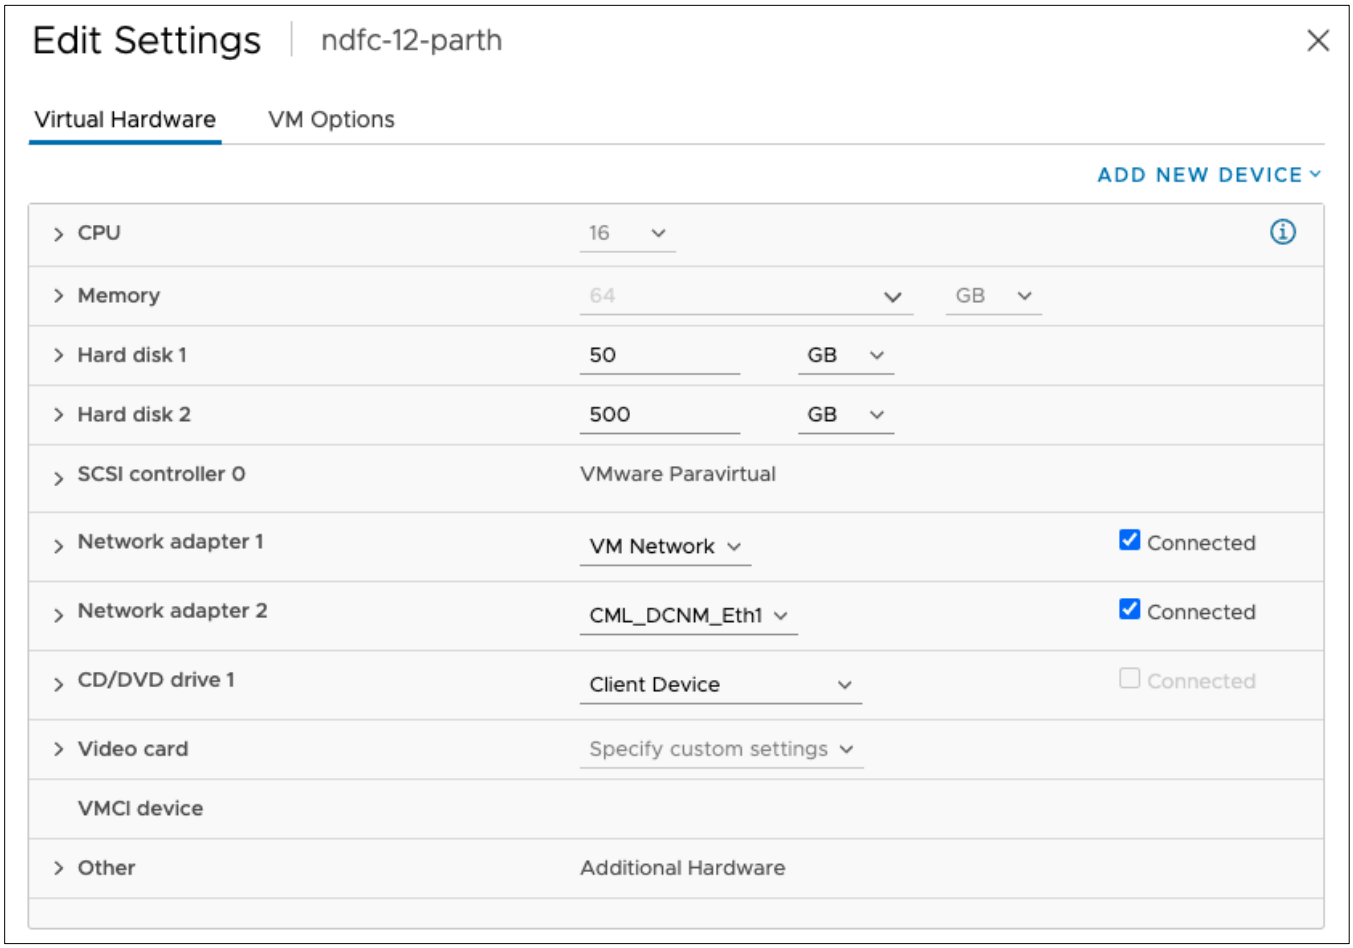

Figure 10. vND VM settings

Deployment Option 1

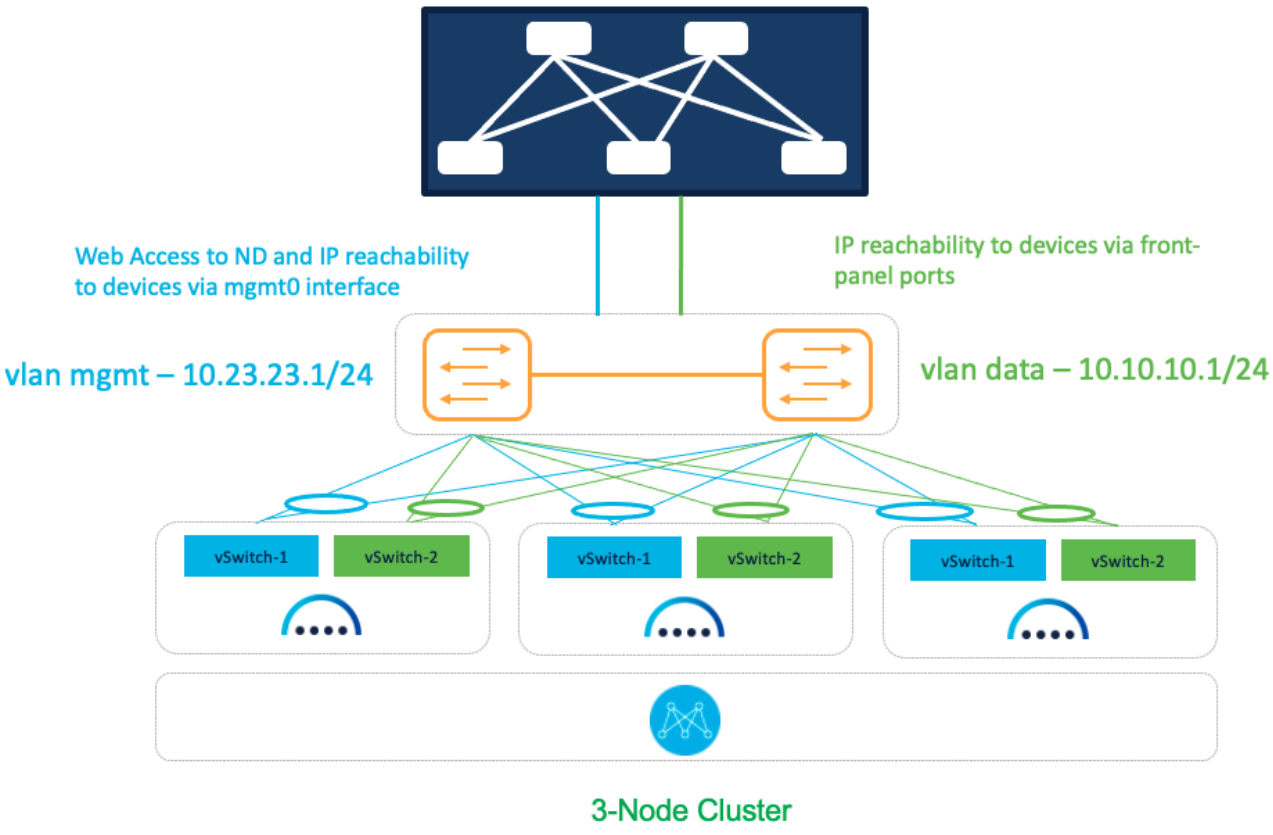

Figure 11. NDFC on vND Deployment Option 1

In this option, we are using a common set of switches that can provide IP reachability to Fabric switches via Nexus Dashboard Management and Data Interfaces. It also uses separate uplinks for Management and Data traffic.

Sample Configurations

On both uplink switches (marked as yellow) for Nexus Dashboard management-

```
interface port-channel1
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 23
  spanning-tree port type edge trunk
  mtu 9216
  vpc 1
interface Ethernet1/1
  description To-ESXi-vND1-mgmt
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 23
  mtu 9216
```

```
  channel-group 1 mode on
  no shutdown
```

An operator must repeat the configuration for the remaining interfaces that are attached to servers hosting vND.

On both uplink switches (marked as yellow) for Nexus Dashboard fabric-

```
interface port-channel2
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10
  spanning-tree port type edge trunk
  mtu 9216
  vpc 2
interface Ethernet1/2
  description To-ESXi-vND1-fabric
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10
  mtu 9216
  channel-group 2 mode on
  no shutdown
```

An operator must repeat the configuration for the remaining interfaces that are attached to servers hosting vND.

Deployment Option 2

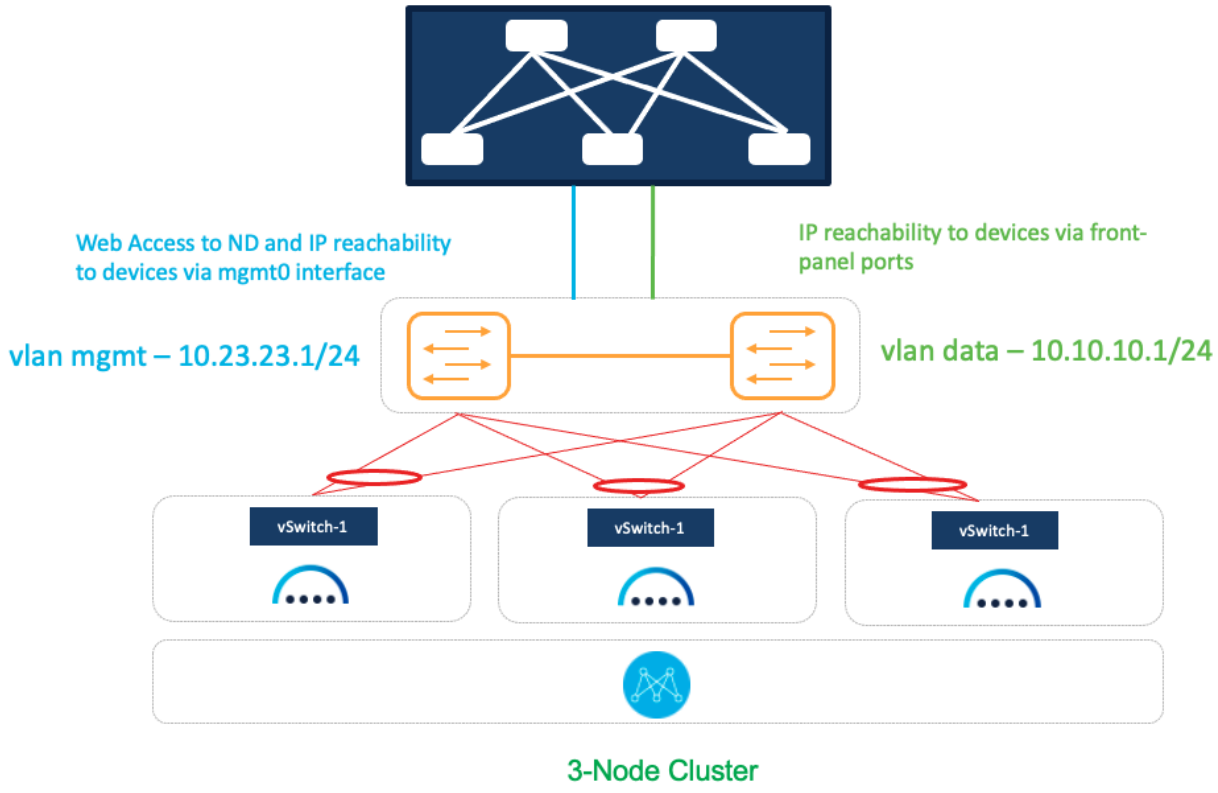

Figure 12. NDFC on vND Deployment Option 2

In this option, we are using a common set of switches that can provide IP reachability to Fabric switches via Nexus Dashboard Management and Data Interfaces. It also uses shared uplinks for both Management and Data traffic.

On both uplink switches (marked as yellow) for Nexus Dashboard management and fabric-

```
interface port-channel1
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 23, 10
  spanning-tree port type edge trunk
  mtu 9216
  vpc 1
interface Ethernet1/1
  description To-ESXi-vND1
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 23, 10
  mtu 9216
  channel-group 1 mode on
  no shutdown
```

An operator must repeat the configuration for the remaining interfaces that are attached to servers hosting vND.

Deployment Option 3

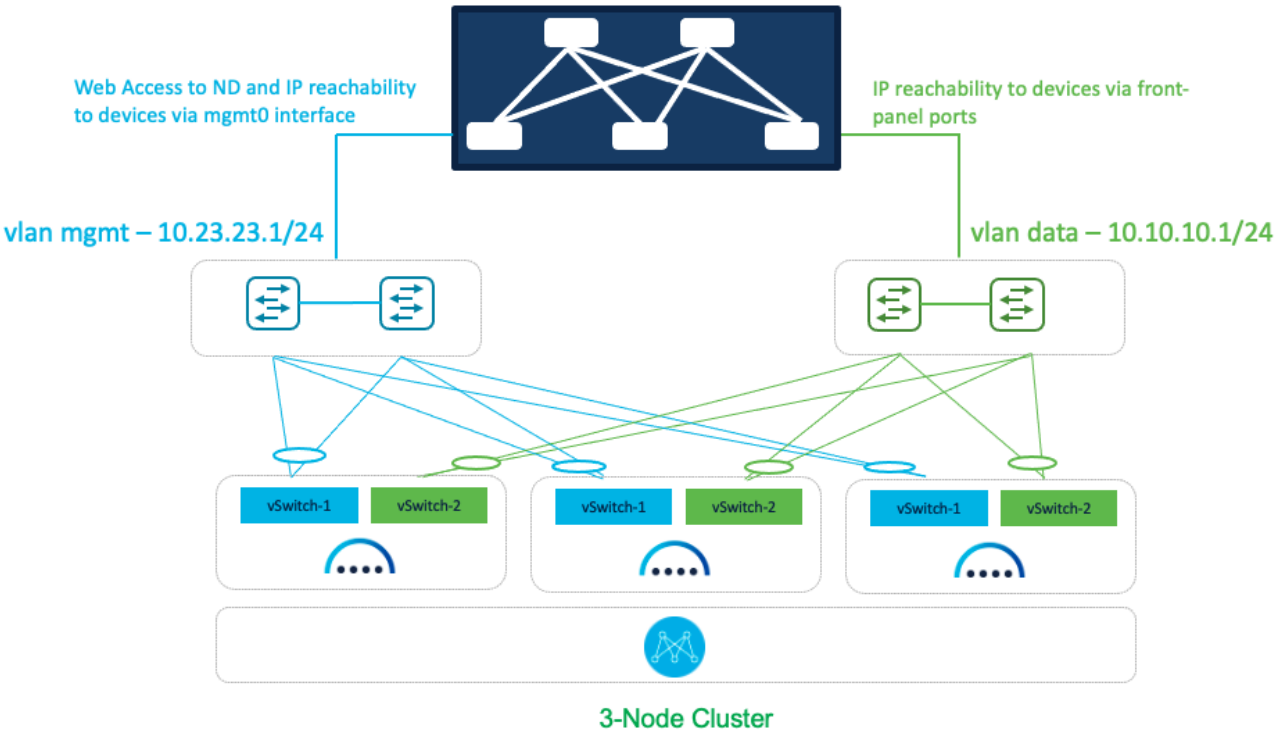

Figure 13. NDFC on vND Deployment Option 3

In this option, we are using a dedicated pair of switches that provide IP reachability to Fabric via Nexus Dashboard Management and Data Interfaces. It also uses separate uplinks for Management and Data traffic.

Sample Configurations

On both uplink switches (marked as blue) for Nexus Dashboard management-

```
interface port-channel1
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 23
  spanning-tree port type edge trunk
  mtu 9216
  vpc 1
interface Ethernet1/1
  description To-ESXi-vND1-mgmt
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 23
  mtu 9216
```

```
  channel-group 1 mode on
  no shutdown
```

An operator must repeat the configuration for the remaining interfaces that are attached to servers hosting vND.

On both uplink switches (marked as green) for Nexus Dashboard fabric-

```
interface port-channel1
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10
  spanning-tree port type edge trunk
  mtu 9216
  vpc 1
interface Ethernet1/1
  description To-ESXi-vND1-fabric
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10
  mtu 9216
  channel-group 1 mode on
  no shutdown
```

An operator must repeat the configuration for the remaining interfaces that are attached to servers hosting vND.

Deployment Modes and Design for SAN Fabrics

When NDFC is enabled with the SAN Controller persona selection, the resultant application can then be employed for managing and monitoring SAN Fabrics. This includes the ability to enable SAN Insights for deep analytics via streaming telemetry. SAN fabrics typically comprise the Cisco MDS family of switches that support SAN traffic over the Fiber Channel. Recall that for NDFC SAN Controller deployments, both a single and a 3-node vND/pND deployments are supported. Refer to NDFC Verified Scalability Guide for more details on the supported scale especially with SAN Insights.

Deploying SAN Controller on pND

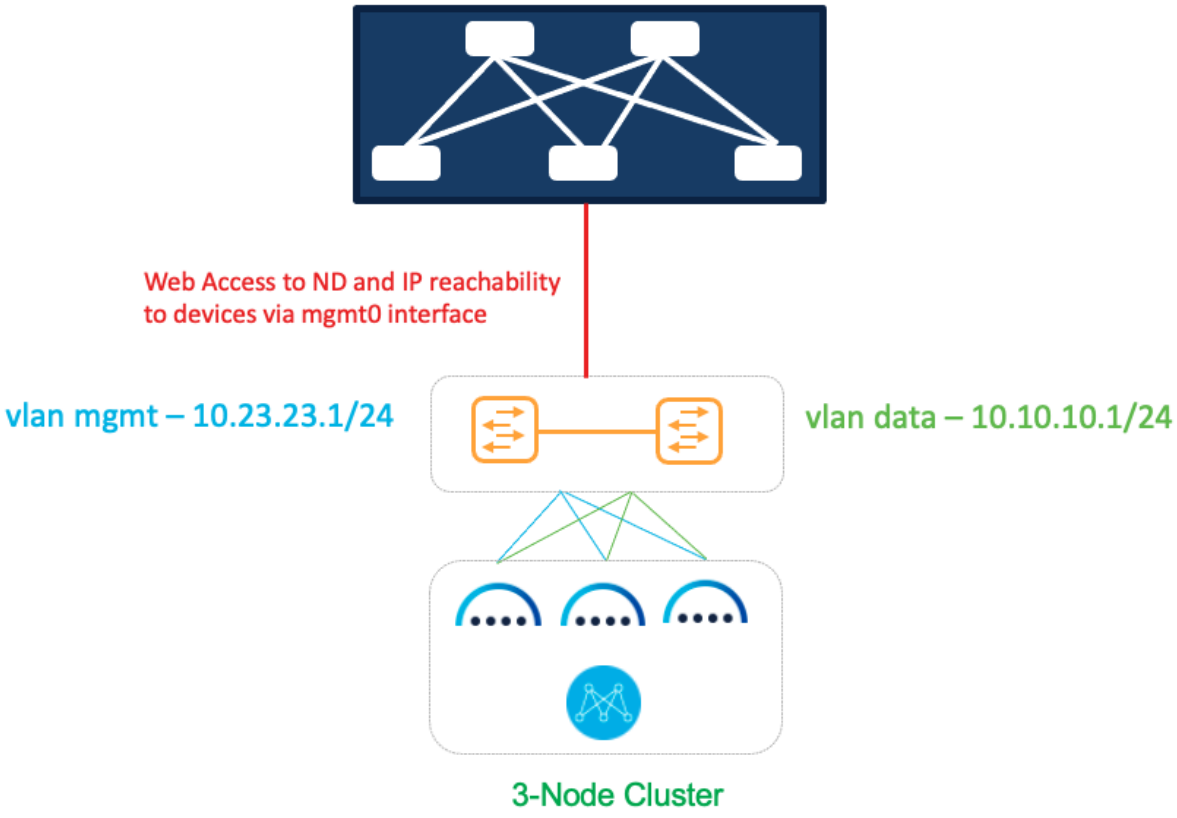

Figure 14. Deploying SAN Controller on pND

In this option, we are using a common set of switches that can provide IP reachability to Fabric switches via Nexus Dashboard Management or Data Interfaces.

Sample configurations

On both uplink switches (marked as yellow) for Nexus Dashboard management-

```
Interface eth1/1, eth1/3, eth1/5
  switchport mode access
  switchport access 23
```

On both uplink switches (marked as yellow) for Nexus Dashboard fabric-

```
Interface eth1/2,eth1/4,eth1/6
  switchport mode access
  switchport access vlan 10
```

OR

```
Interface eth1/2,eth1/4,eth1/6
  switchport mode trunk
  switchport trunk native vlan 10
  switchport trunk allowed vlan 10
```

OR

```
Interface eth1/2,eth1/4,eth1/6
```

```
  switchport mode trunk
  switchport trunk allowed vlan 10
```

For the last option without the trunk native vlan, an operator provides VLAN ID 10 as VLAN tag during Nexus Dashboard installation and Interface bootstrap as shown in Figure 3 from the Networking Requirements Section.

Deploying SAN Controller on vND

Deployment Option 1

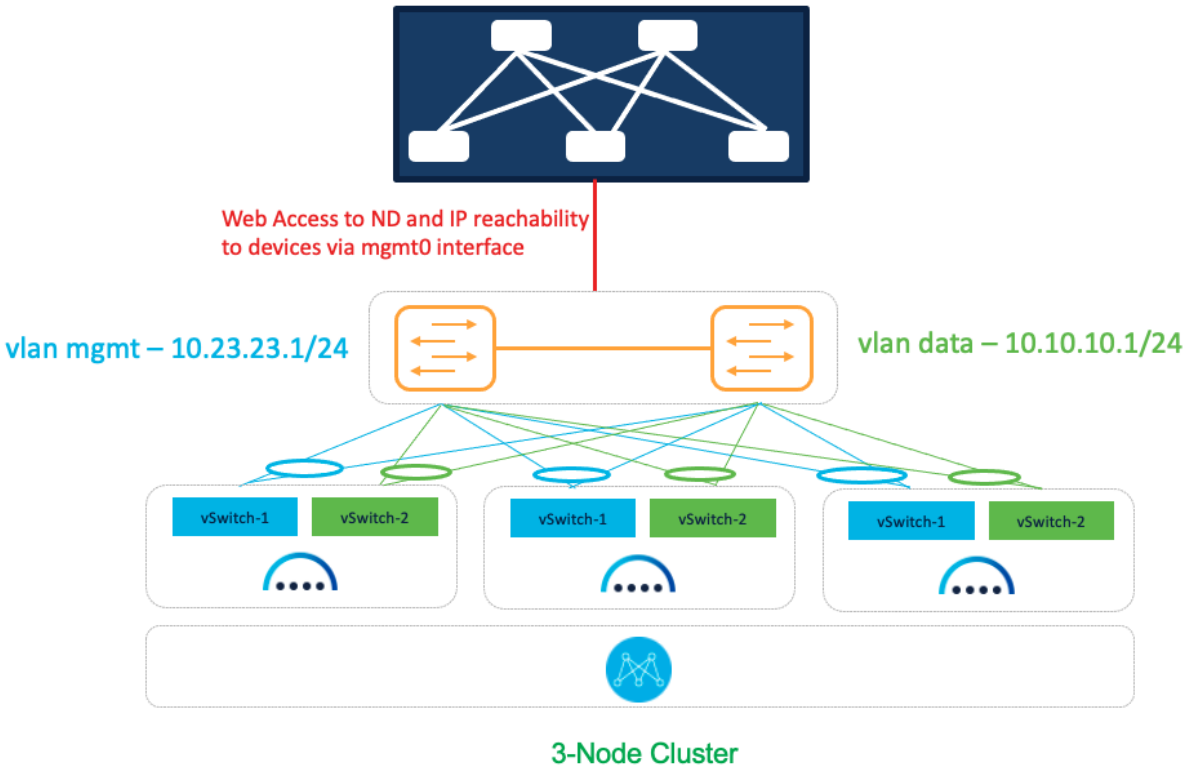

Figure 15. SAN Controller on vND Deployment Option 1

In this option, we are using a common set of switches that can provide IP reachability to Fabric switches via Nexus Dashboard Management or Data interfaces. It also uses separate uplinks for Management and Data traffic.

Sample Configurations

On both uplink switches (marked as yellow) for Nexus Dashboard management-

```
interface port-channel1
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 23
  spanning-tree port type edge trunk
  mtu 9216
  vpc 1
interface Ethernet1/1
```

```
  description To-ESXi-vND1-mgmt
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 23
  mtu 9216
  channel-group 1 mode on
  no shutdown
```

An operator must repeat the configuration for the remaining interfaces that are attached to servers hosting vND.

On both uplink switches (marked as yellow) for Nexus Dashboard fabric-

```
interface port-channel2
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10
  spanning-tree port type edge trunk
  mtu 9216
  vpc 2
interface Ethernet1/2
  description To-ESXi-vND1-fabric
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 10
  mtu 9216
  channel-group 2 mode on
  no shutdown
```

An operator must repeat the configuration for the remaining interfaces that are attached to servers hosting vND.

Deployment Option 2

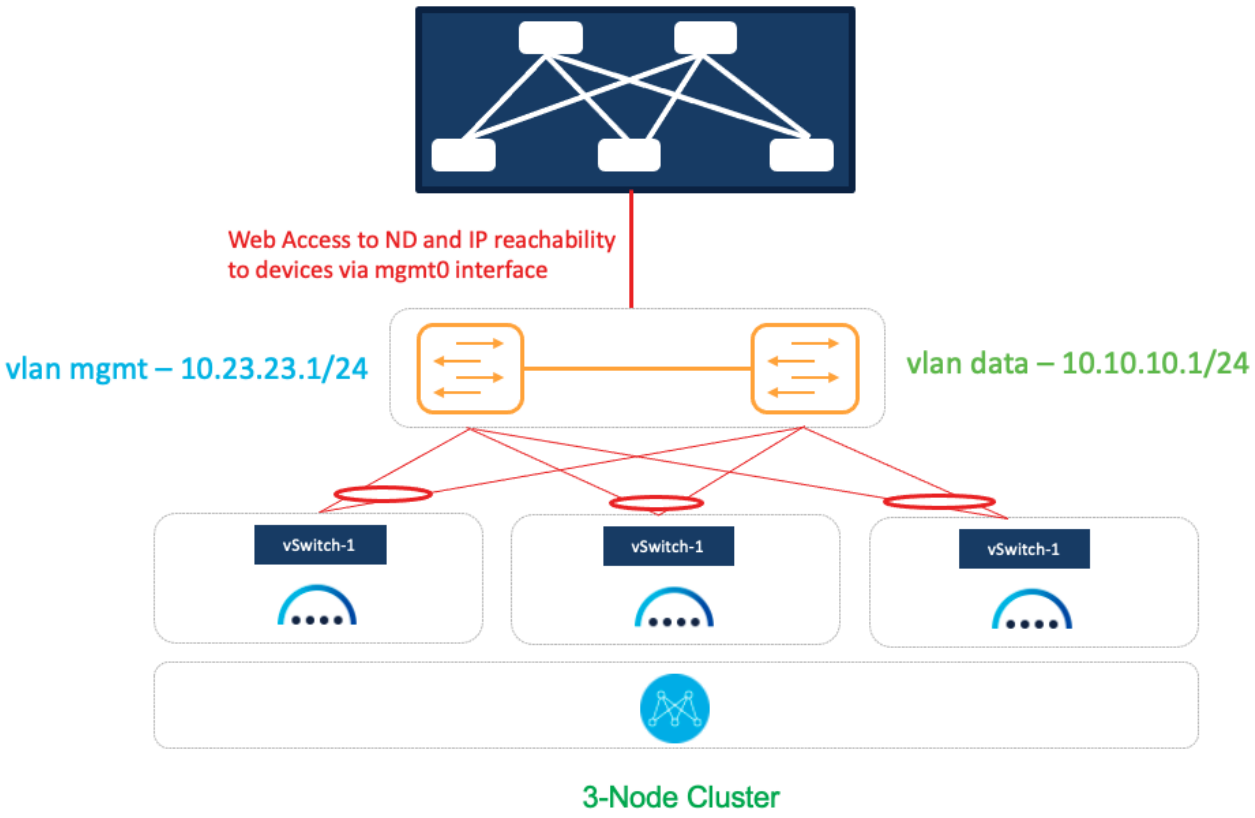

Figure 16. SAN Controller on vND Deployment Option 2

In this option, we are using a common set of switches that can provide IP reachability to Fabric switches via Nexus Dashboard Management or Data interfaces. It also uses shared uplinks for both Management and Data traffic.

On both uplink switches (marked as yellow) for Nexus Dashboard management and fabric-

```
interface port-channel1
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 23, 10
  spanning-tree port type edge trunk
  mtu 9216
  vpc 1
interface Ethernet1/1
  description To-ESXi-vND1
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 23, 10
  mtu 9216
  channel-group 1 mode on
  no shutdown
```

An operator must repeat the configuration for the remaining interfaces that are attached to servers hosting vND.

Persistent IP Requirements for NDFC

Persistent IP addresses, also known as External Service IP addresses, are required for pods/services in NDFC that require sticky IP addresses. In other words, pods that are provisioned with a persistent IP retain their IP address even if they are re-provisioned either on the same Nexus Dashboard node or a different Nexus Dashboard node within the same Nexus Dashboard cluster. The reason persistent IP addresses are required is that switches may be configured with a certain NDFC service as a destination (e.g., SNMP trap destination) and one should not require this switch configuration to change on the failure of a Nexus Dashboard node that is hosting the corresponding service/pod. Therefore, for uninterrupted service, the associated service/pod must be respawned somewhere else in the Nexus Dashboard cluster aka another node, such that the pod/service IP remains the same. Since the Nexus Dashboard nodes are typically Layer-2 adjacent, from a network reachability point of view, there is nothing else required for traffic to be redirected to the new location of that destination service/pod. Note that with the introduction of Layer-3 reachability for a ND cluster hosting NDFC, eBGP is employed to dynamically advertise the updated location of the service post a node failure. Consequently, from a network reachability point of view, as soon as the pod is up at the new location, service resumption occurs without any user intervention.

External Service IP addresses are configured under Nexus Dashboard Cluster Configuration. The usage of Persistent IP addresses is based on what features are enabled on the NDFC, the deployment model and the way NDFC connects with the switches. Based on various use-cases you may need IP addresses on the Nexus Dashboard management pool or data pool or both.

For Virtual Nexus Dashboard Deployments, enable/accept promiscuous mode on the port-groups associated with the Nexus Dashboard Management and/or Data vNICs where IP stickiness is required. The Persistent IP addresses are given to the pods (e.g., SNMP Trap/Syslog receiver, Endpoint Locator instance per Fabric, SAN Insights receiver, etc.). Every POD in Kubernetes can have multiple virtual interfaces. Specifically for IP stickiness an additional virtual interface is associated with the POD that is allocated an appropriate free IP from the appropriate external service IP pool. The vNIC has its own unique MAC address that is different from the MAC addresses associated with the vND virtual vNICs. Moreover, all communication to and from the PODs towards an external switch goes out of the same bond interface for North-to-South traffic flows. The Data vNIC maps to bond0 (also known as bond0br) interface and Management vNIC maps to bond1 (also known as bond1br) interface. By default, the VMware system checks if the traffic flows out of a particular vNIC is matched with the Source-MAC associated with the vNIC. In the case of NDFC, the traffic flows are sourced with the Persistent IP address and associated MAC of the given PODs. Therefore, we need to enable the required settings on the VMware side.

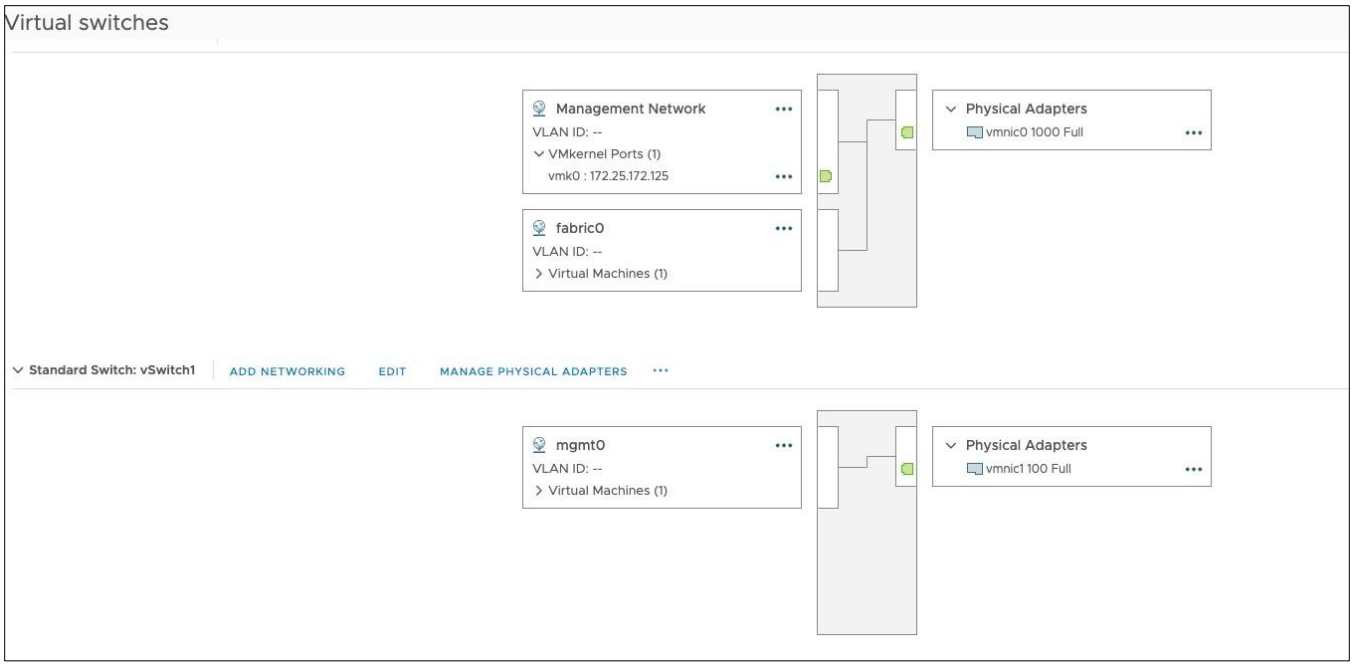

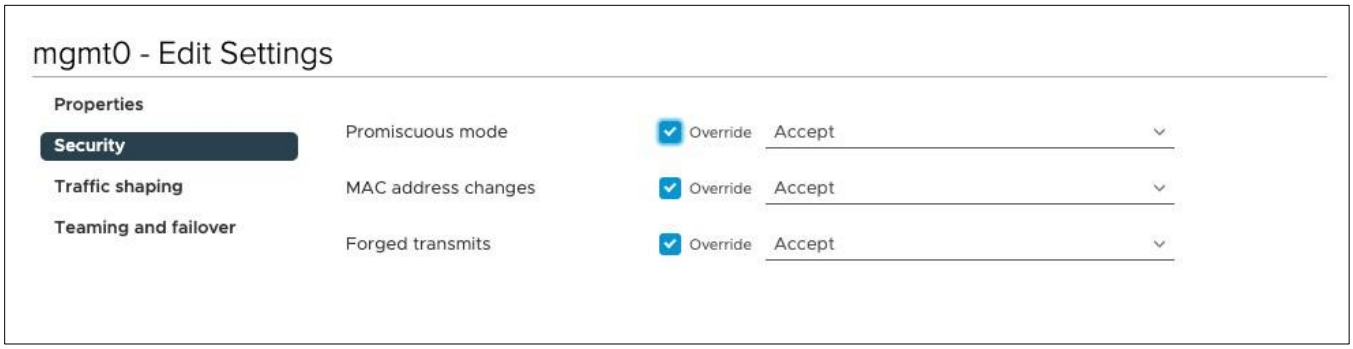

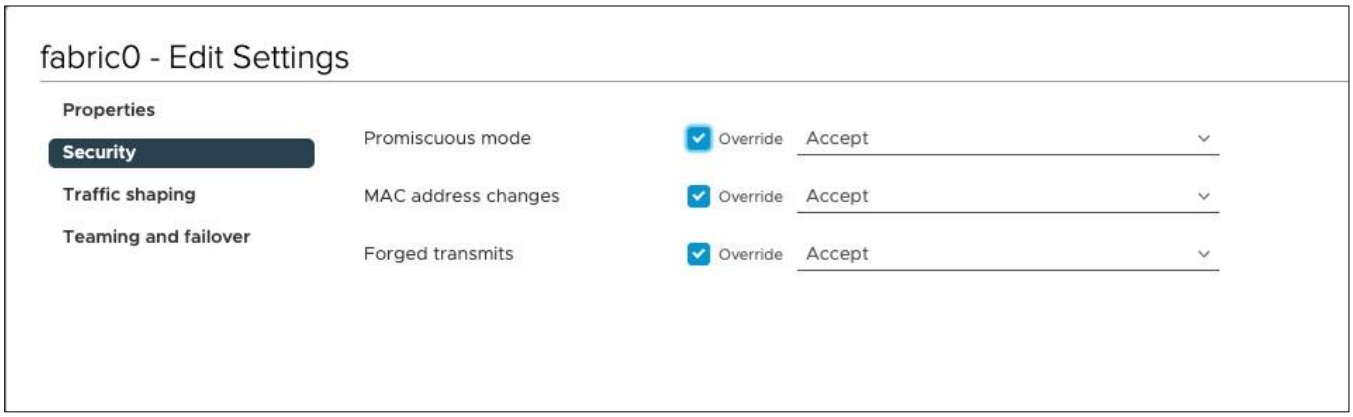

Note: An operator will not be able to activate an NDFC feature if appropriate Persistent IP addresses are not available. NDFC has a precheck that confirms that there are enough free External Service IP addresses configured on the Nexus Dashboard, in the corresponding pool before a feature that has such a requirement can be enabled.

As mentioned in NDFC to Device Connectivity section, depending on the specific use case and the selected interface for communicating with the switches mgmt0 interfaces, the persistent IP addresses will have to be associated with the ND Management interface or to the ND Data interface.

From Cisco NDFC Release 12.1.2e, 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 12.1.2e, if Inband 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.

Note: You can disable promiscuous mode when Nexus Dashboard nodes are Layer-3 adjacent to the Data network, BGP is configured, and fabric switches are reachable through the data interface.

Note: 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.

Configuring Persistent IP Addresses

To configure the Persistent IP addresses (also known as External Service IP) perform the following steps:

Step 1. Navigate to **Nexus Dashboard Admin console**.

Step 2. Click on **Infrastructure** Tab.

Step 3. Under Infrastructure navigate to **Cluster Configuration**.

Step 4. Based on the deployment model and use-case edit the **External Service Pools**, and associate the persistent IP addresses to the Management or Data interfaces.

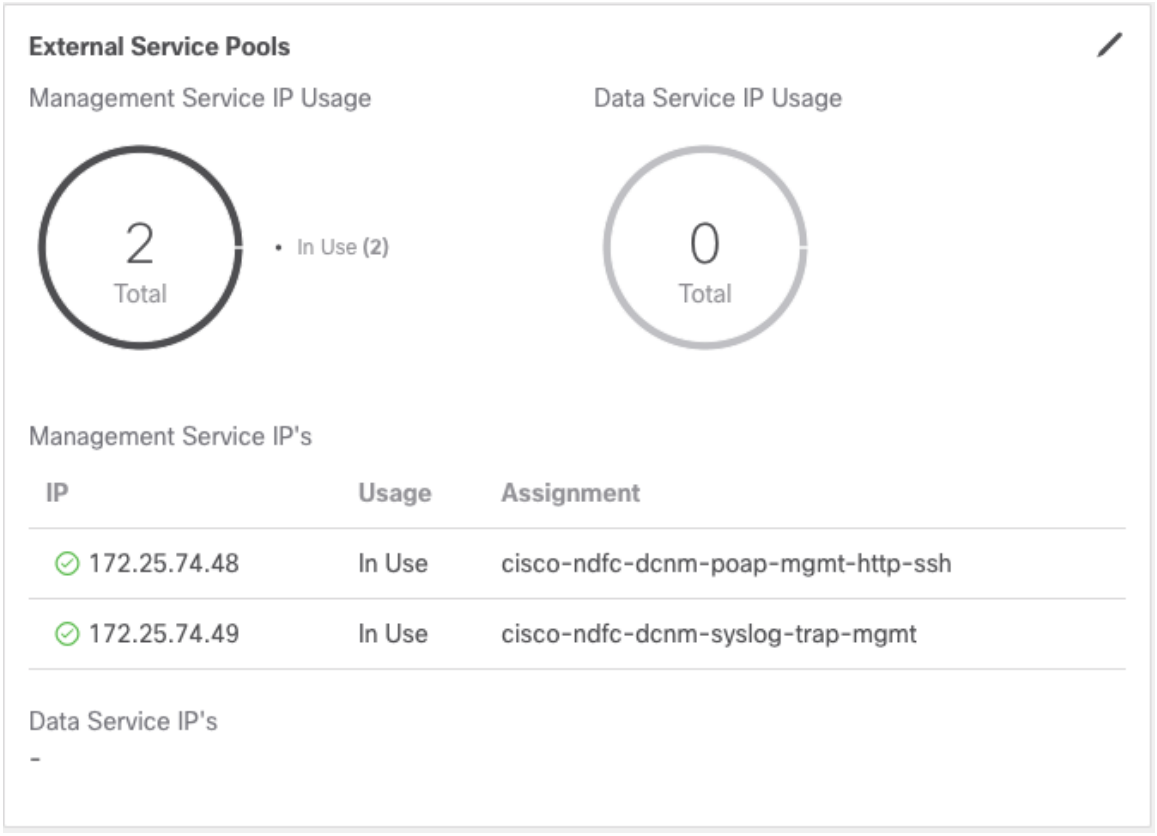

Figure 17. Nexus Dashboard Persistent IP in Management Pool for LAN deployments

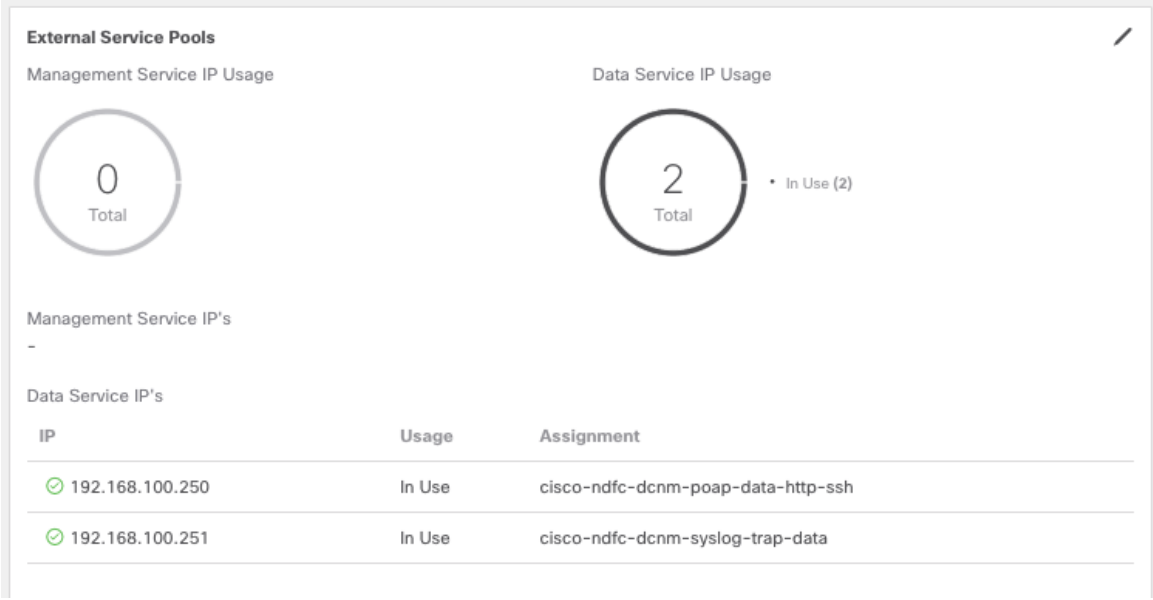

Figure 18. Nexus Dashboard Persistent IP in Data Pool for LAN deployments

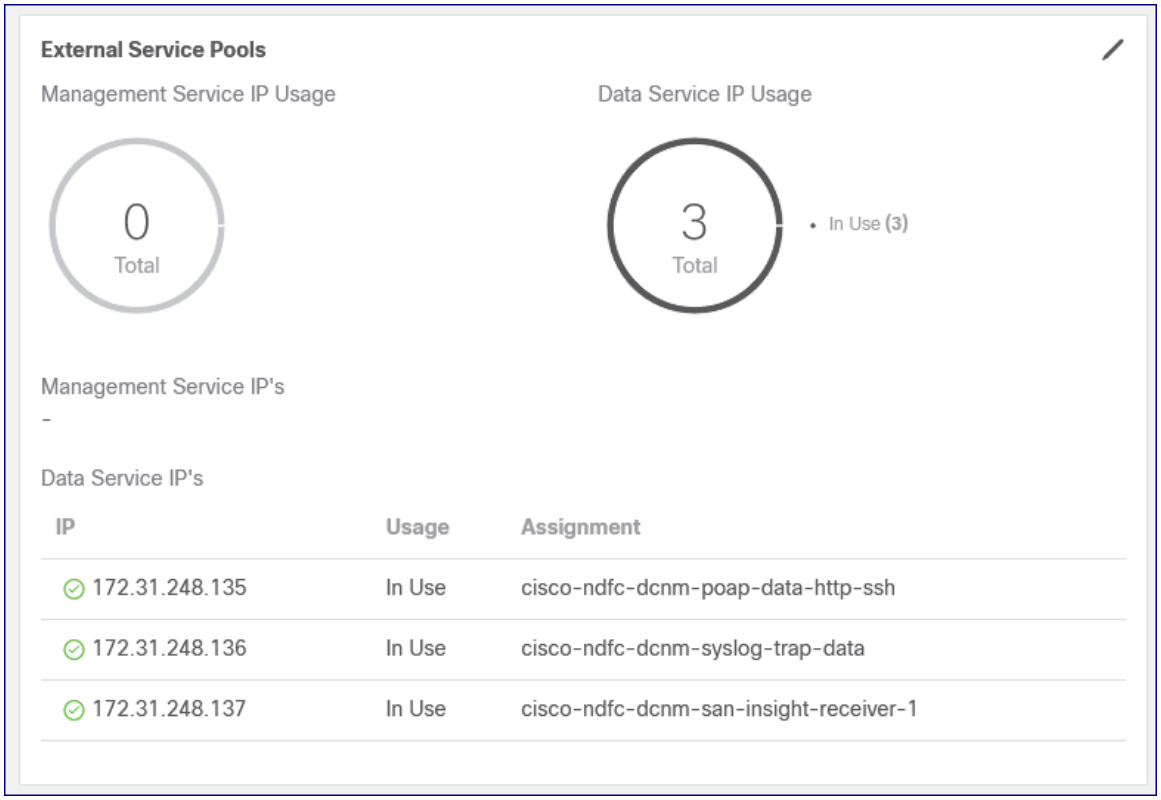

Figure 19. Nexus Dashboard Persistent IP in Data Pool for SAN deployments

As a reminder, if the user selects to use the ND Data interface to communicate with the switches mgmt0 interfaces, before assigning the persistent IP addresses to such interface, it is also required to override the default global server settings for LAN Device Management Connectivity. For this purpose, an operator must navigate to NDFC server settings, go to **Admin** tab and specify data in the **LAN Device Management Connectivity** field.

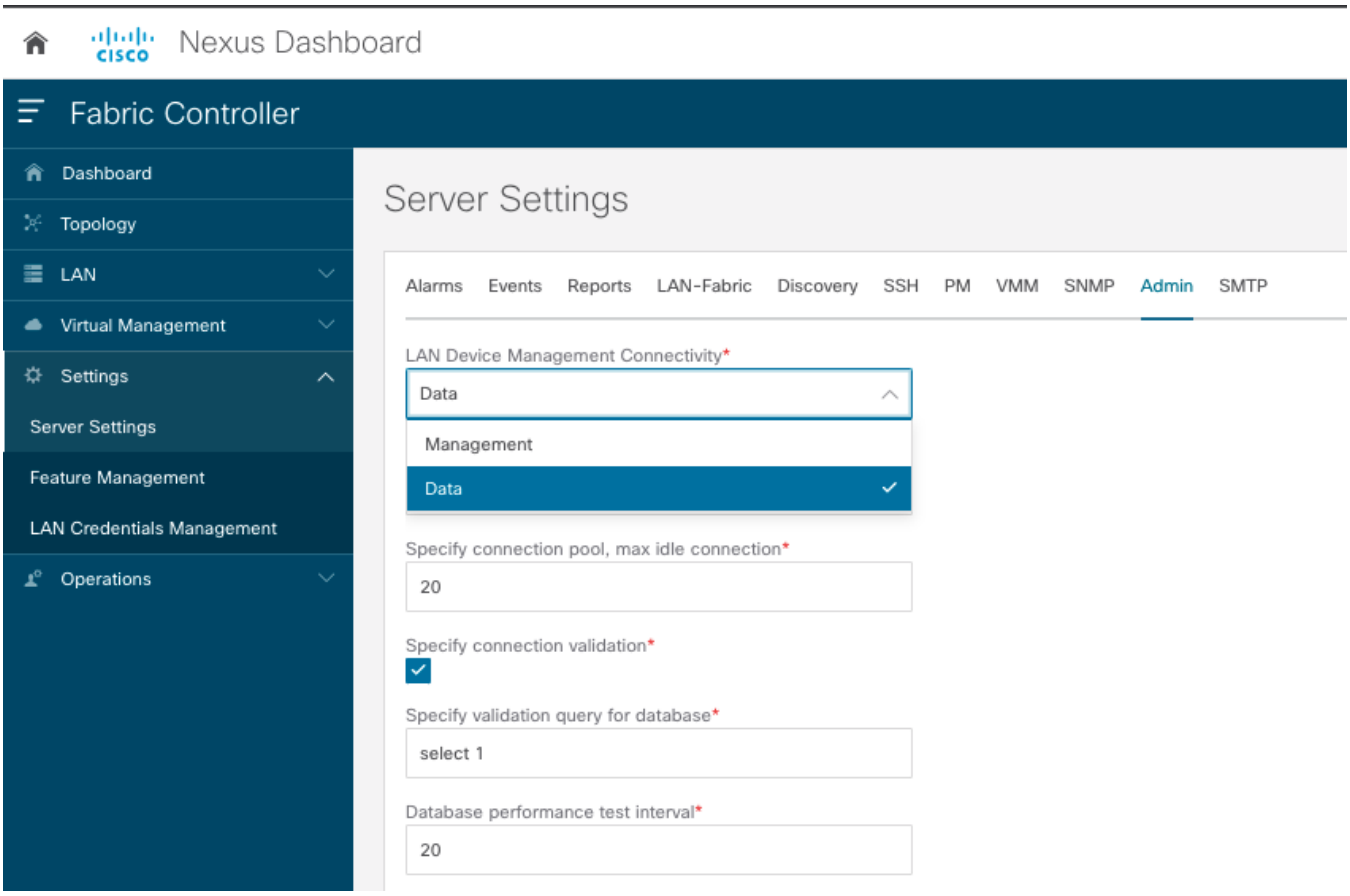

Figure 20. Server Settings for LAN Device Management

For SAN deployments, recall that all NDFC SAN controller to device reachability is over the Nexus Dashboard data interface. Hence the requirement is the same as above, two free IP addresses are required in the Nexus Dashboard External Service Data IP Pool. Additionally, one IP address per cluster node is required to receive SAN Insights streaming data.

Examples of Persistent IP addresses

As mentioned earlier, there are various features in NDFC that require a persistent IP to be allocated to one of the associated pods to provide the desired functionality. In this section, we will provide an explanation of those features.

SNMP Trap/Syslog Receiver

For SNMP trap, switches are typically configured with an appropriate SNMP trap destination IP. Similarly, an external entity is a setup for SYSLOG reception and its IP address is typically configured on the switch, so that the appropriate SYSLOGs of interest can be directed to this entity. NDFC has the capability to function as a SNMP Trap destination, as well as a SYSLOG receiver. It is important to keep the IP associated with these services as static, instead of changing the switch configuration when the associated service container/pod is moved across ND nodes.

NDFC can manage switches via Out-of-Band or Inband access. The **LAN Device Management Connectivity** option in the NDFC Web UI **Server Settings**, determines whether the SNMP Trap/SYSLOG receiver will be spawned with a persistent IP associated with the ND management or ND data interface. If an operator selects the Nexus Dashboard management interface, then persistent IP must be added to the Nexus Dashboard management Interface. The default value for the setting is **Management**. If an operator

changes this setting to **Data**, then prior to that, a persistent IP must be configured in the **Cluster Configuration** on the Nexus Dashboard **Admin Console**.

POAP/SCP

NX-OS had a constraint that does not allow to change the port (22) when setting up an SCP session with NDFC or for that matter, any SCP server. This restriction has since been lifted in the latest versions of NX-OS. Port 22 on the Nexus Dashboard node is bound to the Nexus Dashboard host itself and is not available for use by any of the application services such as NDFC. Hence, in order to use port 22 for SCP purposes, a new IP is required that is different from the ND node IP itself, which is assigned to the SCP pod. SCP services are required for image management, NX-API certificate installation as well as POAP purposes. This same IP address is shared by the POAP services (http/tftp) for touchless Day-0 switch bring-up. Note that POAP is supported on either the ND management or ND data interface. The LAN Device Management Connectivity in the NDFC Server Settings controls which interface POAP is supported. As mentioned earlier, by default the value of this field is set to Management. When an operator changes this setting to Data, both Out-of-Band and Inband POAP are supported simultaneously by NDFC albeit for different fabrics.

Hence, Persistent IP addresses are required even for Single Nexus Dashboard Node Deployments.

EPL (Endpoint Locator)

For each VXLAN (Virtual extensible Local Area Network) fabric, a specific container is spawned running a BGP instance to peer with the spines aka Route-Reflectors (RRs) of the fabric. This container must have an associated persistent IP that is then configured as a BGP neighbor on the spine RRs. A different container or pod is used for each fabric, so the number of fabrics managed by NDFC where EPL is enabled decides how many persistent IP addresses are needed for EPL. Also, the EPL establishes BGP sessions only over the Nexus Dashboard Data interface.

PMN

For IP Fabric for Media (IPFM) deployments, a telemetry receiver pod is spawned that is responsible for the reception of multicast flow and PTP-related information via streaming telemetry. Since the receiver IP address is configured on the switch, again this explains the need for a persistent or sticky IP to prevent reconfiguration of the physical network in case the node with the pod goes down. Up to a maximum of 3 PMN receivers are supported for a 3-node PMN cluster deployment.

SAN

Similar to IPFM, for SAN Insights functionality, there is a requirement to support a receiver per ND node that is responsible for processing and analyzing the SAN analytics data received from the SAN MDS switches. Both single node (1 receiver) and 3 nodes (3 receivers) deployments are supported.

Nexus Dashboard IP Allocation

In this section, we will cover different IP allocations including the Persistent IP addresses for different Deployment models.

LAN Fabrics

This section details information about Nexus Dashboard IP Allocation for LAN Fabrics.

Case 1

OOB Device Management over Nexus Dashboard Management Interface

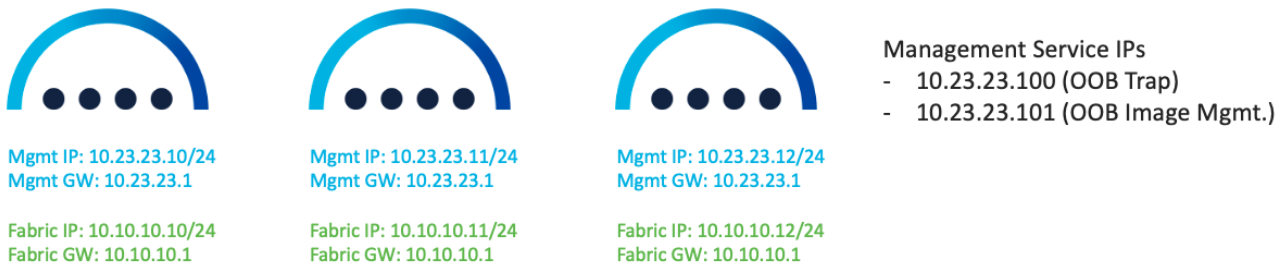

Figure 21. Nexus Dashboard IP Allocation for LAN Fabrics Case 1

Case 2

OOB Device Management over Nexus Dashboard Data Interface

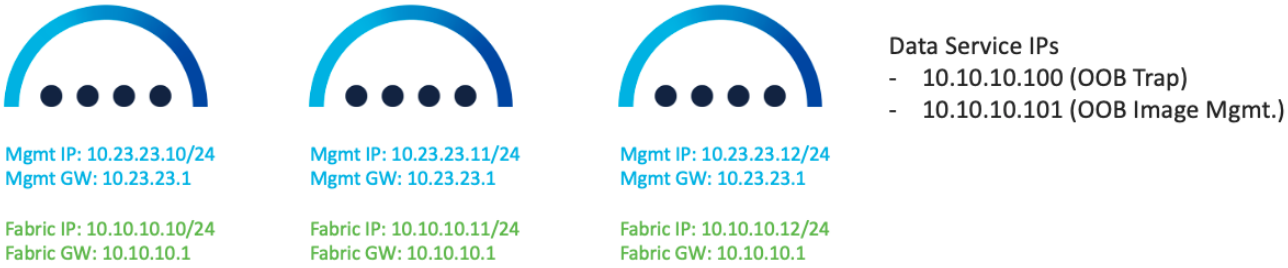

Figure 22. Nexus Dashboard IP Allocation for LAN Fabrics Case 2

Case 3

- OOB Device Management over Nexus Dashboard Management Interface
- EPL over Nexus Dashboard Data Interface

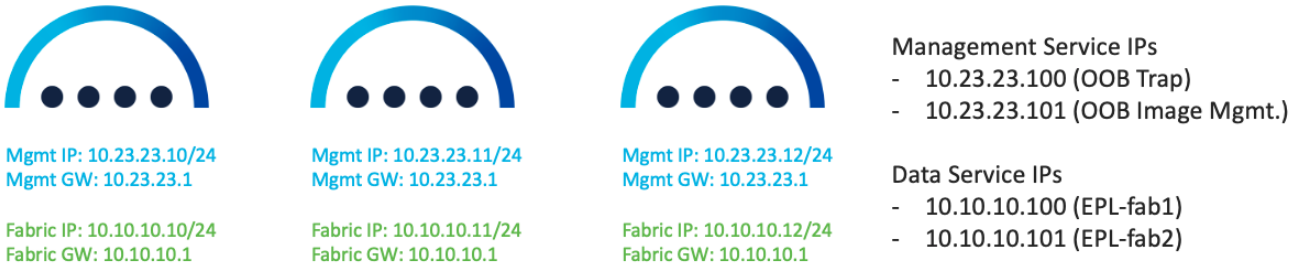

Figure 23. Nexus Dashboard IP Allocation for LAN Fabrics Case 3

Case 4a

ND Cluster is Layer-2 Adjacent

- OOB Device Management over Nexus Dashboard Data Interface
- Inband Device Management over Nexus Dashboard Data Interface
- EPL over Nexus Dashboard Data Interface

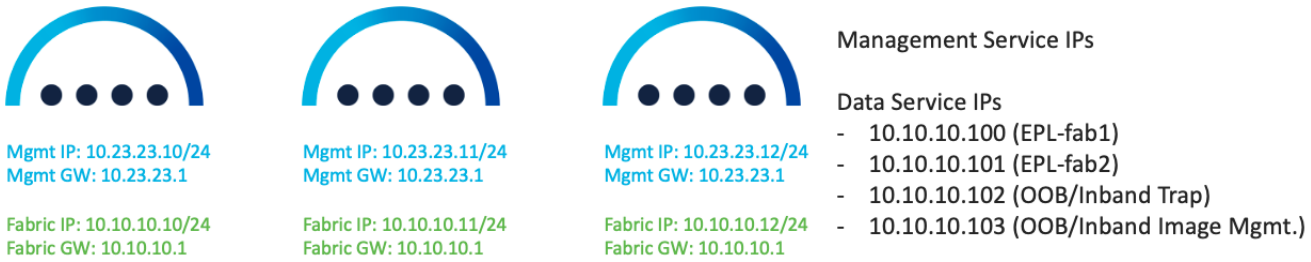

Figure 24. Nexus Dashboard IP Allocation for LAN Fabrics Case 4a

Case 4b

ND Cluster is Layer-3 Adjacent

- OOB Device Management over Nexus Dashboard Data Interface
- Inband Device Management over Nexus Dashboard Data Interface
- EPL over Nexus Dashboard Data Interface

With Layer-3 HA, the ND forms eBGP peering with uplink switches (no support for eBGP Multi-Hop) and the External Persistent/Service IP Pool belongs to a unique subnet pool that is outside the subnet range of ND Management and Data interfaces.

Figure 25. Nexus Dashboard IP Allocation for LAN Fabrics Case 4b

Case 5a

ND Cluster is Layer-2 Adjacent

- OOB Device Management over Nexus Dashboard Management Interface
- In-band Device Management over Nexus Dashboard Data Interface

- EPL over Nexus Dashboard Data Interface

Figure 26. Nexus Dashboard IP Allocation for LAN Fabrics Case 5a

Case 5b

ND Cluster is Layer-3 Adjacent

- OOB Device Management over Nexus Dashboard Management Interface
- In-band Device Management over Nexus Dashboard Data Interface
- EPL over Nexus Dashboard Data Interface

With Layer-3 HA, the ND forms eBGP peering with uplink switches (no support for eBGP Multi-Hop) and the External Persistent/Service IP Pool belongs to a unique subnet pool that is outside the subnet range of ND Management and Data interfaces.

Figure 27. Nexus Dashboard IP Allocation for LAN Fabrics Case 5b

IPFM Fabrics

This section details information about Nexus Dashboard IP Allocation for IPFM Fabrics.

Case 1

OOB Device Management over Nexus Dashboard Management Interface

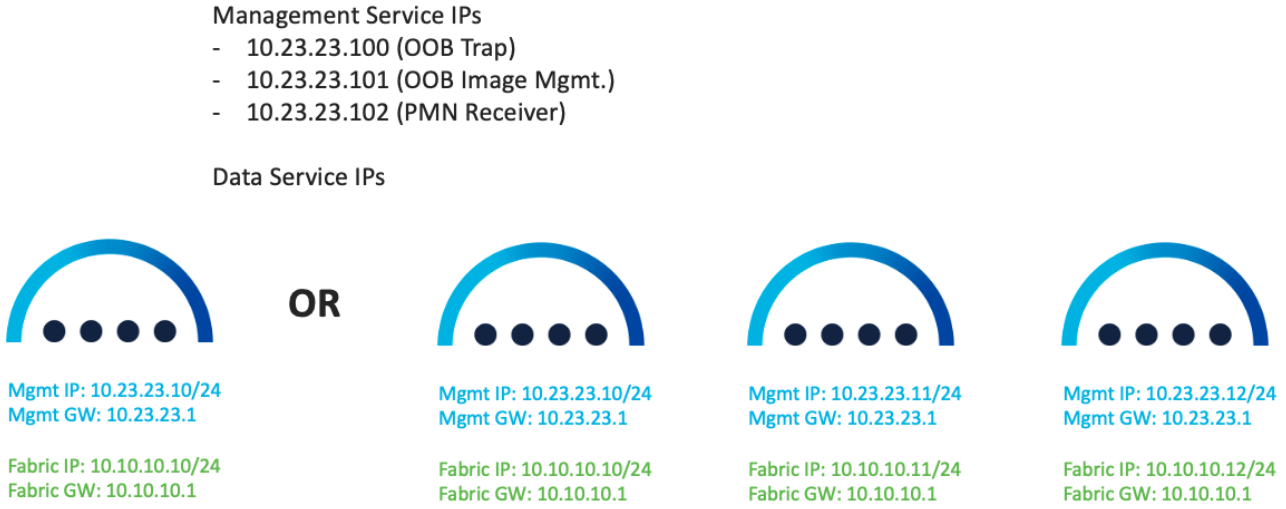

Figure 28. Nexus Dashboard IP Allocation for IPFM Fabrics Case 1

Case 2

- OOB Device Management over Nexus Dashboard Data Interface OR/AND
- In-band Device Management over Nexus Dashboard Data Interface

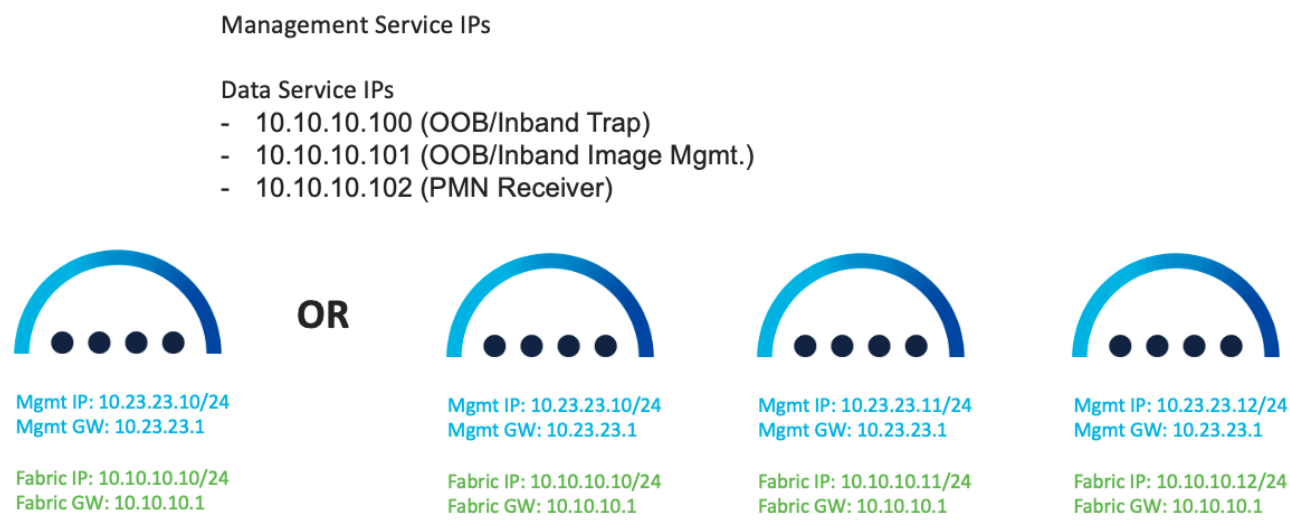

Figure 29. Nexus Dashboard IP Allocation for IPFM Fabrics Case 2

SAN Fabrics

This section details information about Nexus Dashboard IP Allocation for SAN Fabrics.

OOB Device Management Over Nexus Dashboard Data Interface

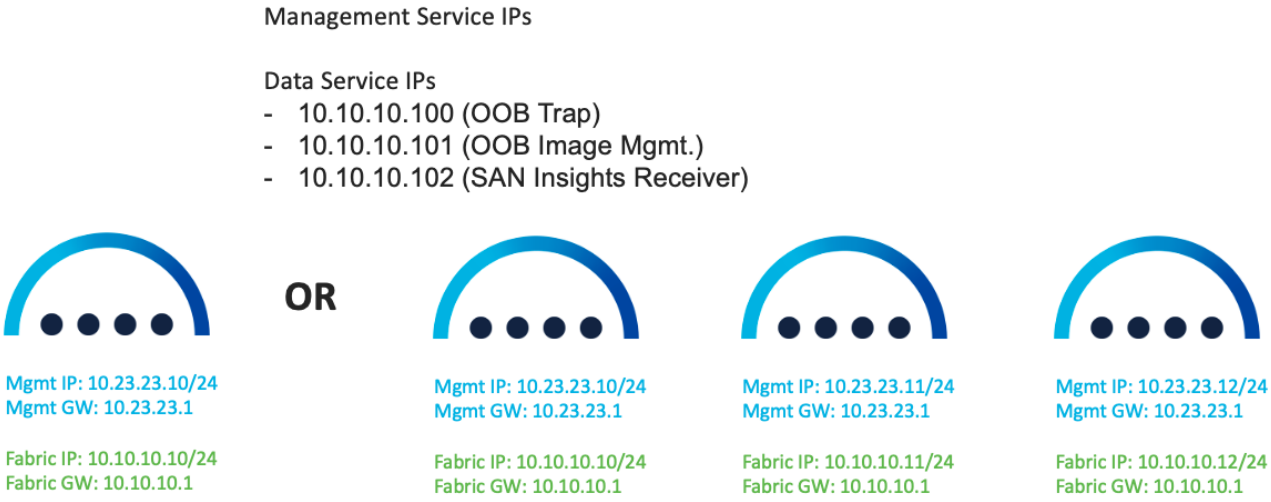

Figure 30. Nexus Dashboard IP Allocation for SAN Fabrics

DCNM to NDFC Migration

NDFC supports a seamless migration or upgrade option for LAN, SAN, IPFM customers using the DCNM 11.x release. At a high level, this requires a backup to be taken on the DCNM 11.x instance, which in turn can then be state fully restored on a new NDFC instance deployed on a Nexus Dashboard cluster. For detailed instructions on how to perform the upgrade from DCNM to NDFC, refer to NDFC 12.0.1a Installation and Upgrade guide.

Recall that for DCNM 11.x, there are various deployment options supported for customer deployments.

DCNM SAN:

- DCNM Windows Installer
- DCNM Linux Installer
- OVA/ISO for ESXi
- Appliance (DCNM on SE (Service Engine))

DCNM LAN/IP Fabric for Media:

- OVA/ISO for ESXi
- QCOW2 for KVM
- Appliance (DCNM on SE)

In addition, for LAN environments, standalone, native HA (recommended) and native HA cluster-based DCNM deployments are supported. Upgrades to NDFC 12.0 are supported for all these DCNM 11 deployments. The minimum DCNM version required for NDFC upgrade is 11.5(1).

LAN Fabrics

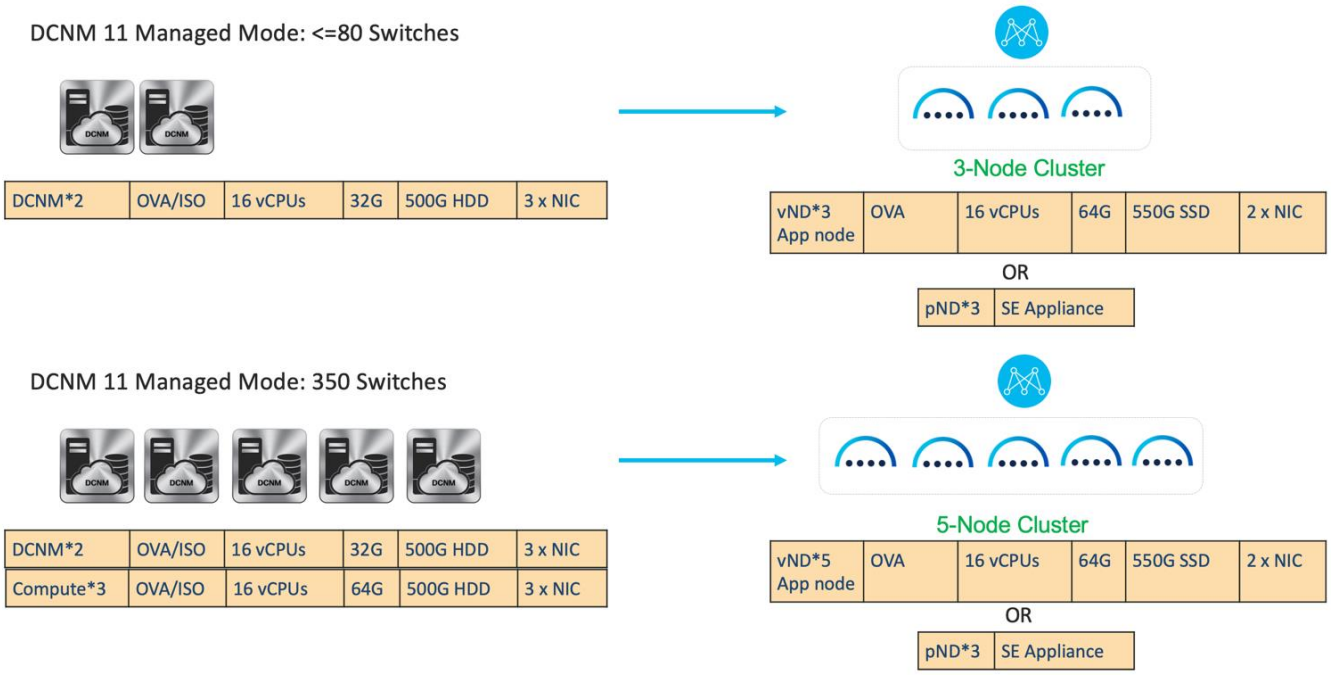

- NDFC 12.0.1 supports up to 80 switches.
- NDFC 12.0.2 supports up to 350 switches

This document does not include the system requirements for pND. See Additional Information section for links that provide access to the Cisco websites specific to Nexus Dashboard.

DCNM LAN Standalone to NDFC Migration

Case 1 - eth0/eth1 in different subnets and eth2 is not used

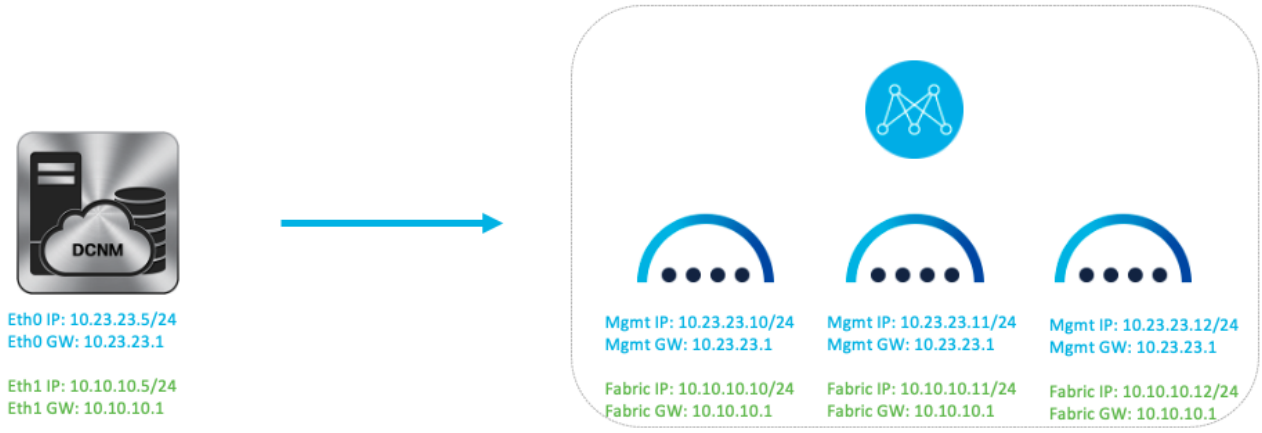

Figure 31. DCNM LAN Standalone to NDFC Migration Case 1

Case 2 - eth0/eth1 are in same subnets and eth2 is not used.

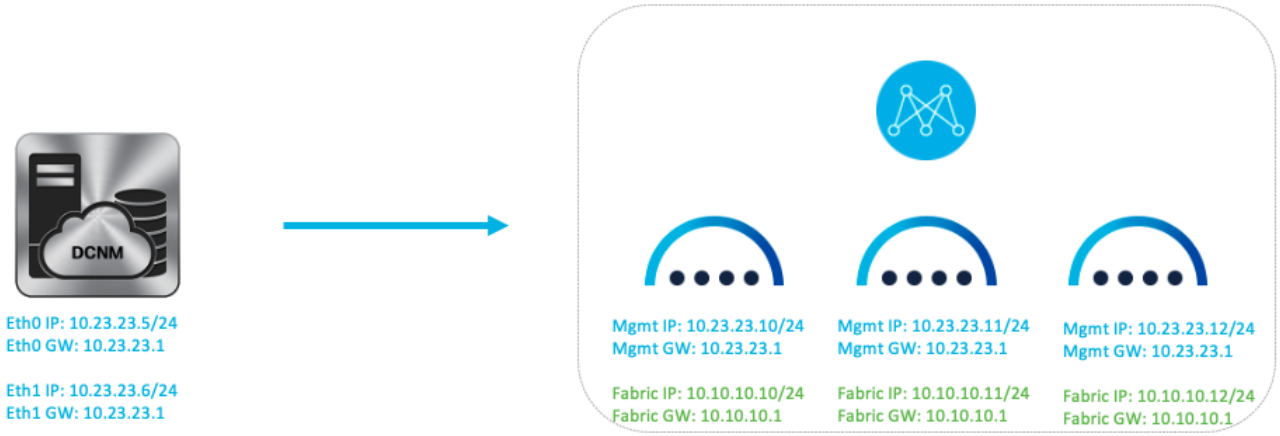

Figure 32. DCNM LAN Standalone to NDFC Migration Case 2

Case 3 - eth0/eth1 are in different subnets and eth2 is used for In-band

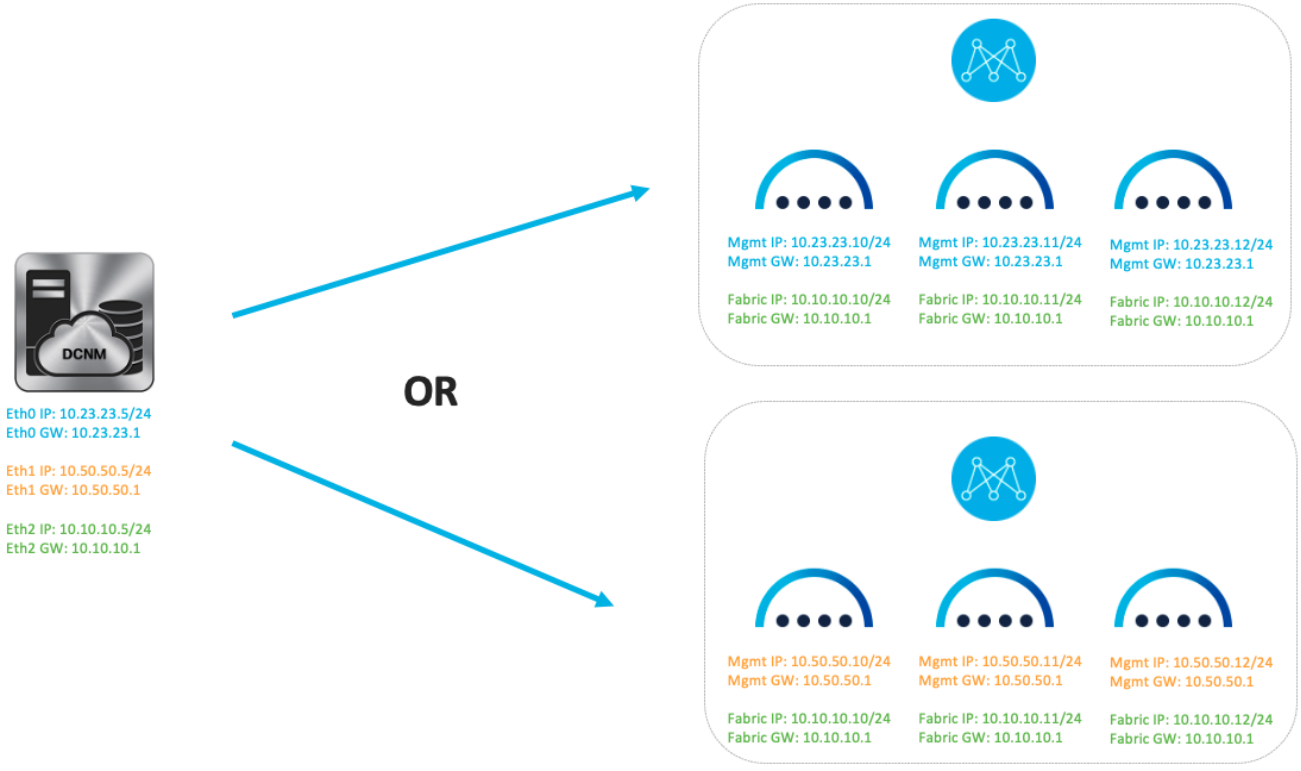

Figure 33. DCNM LAN Standalone to NDFC Migration Case 3

Case 4 - eth0/eth1 are in same subnet and eth2 is used for In-band.

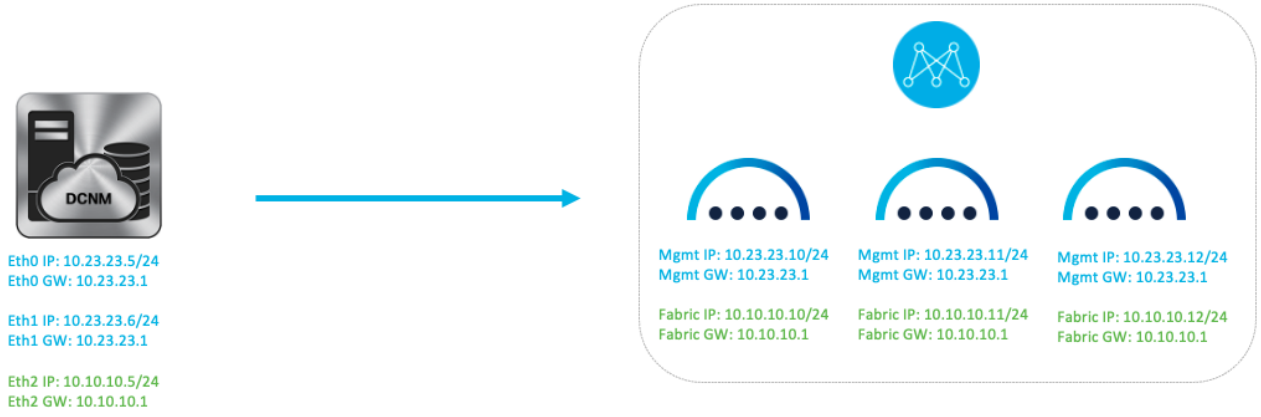

Figure 34. DCNM LAN Standalone to NDFC Migration Case 4

DCNM LAN Native HA to NDFC Migration

Case 1 - eth0/eth1 are in the different subnets and eth2 is not used

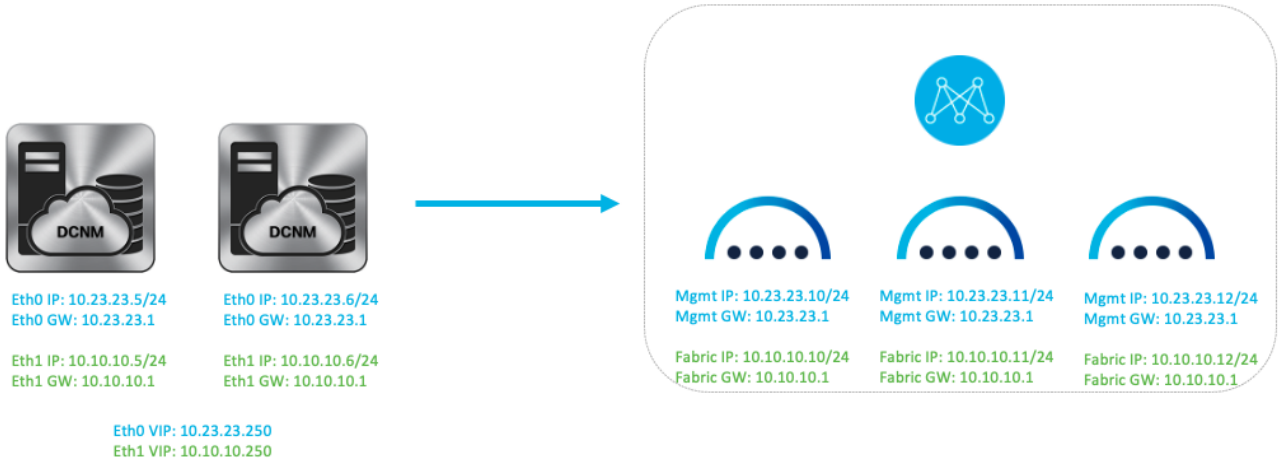

Figure 35. DCNM LAN Native HA to NDFC Migration Case 1

Case 2 - eth0/eth1 are in the same subnet and eth2 is not used

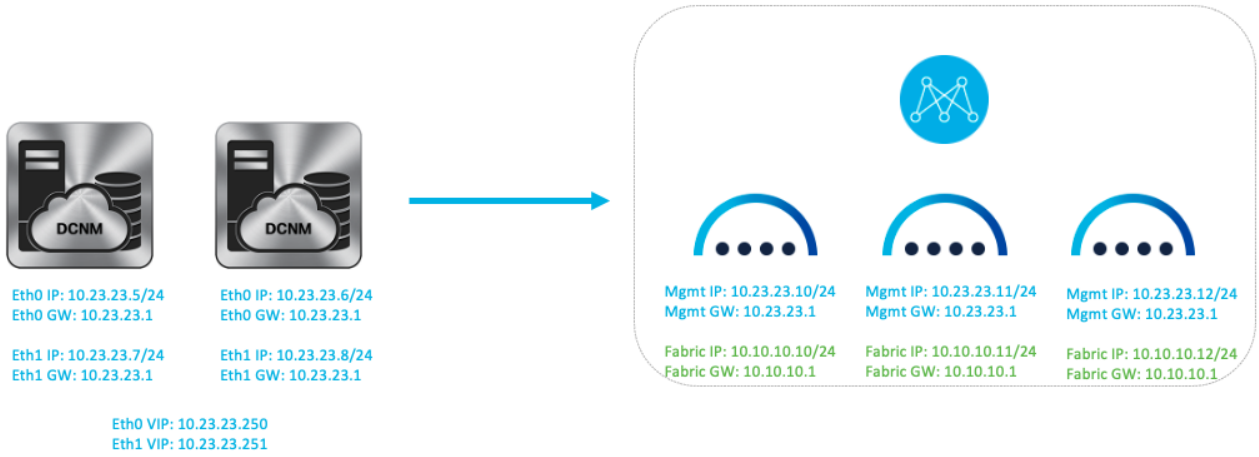

Figure 36. DCNM LAN Native HA to NDFC Migration Case 2

Case 3 - eth0/eth1 are in different subnets and eth2 is used for In-band

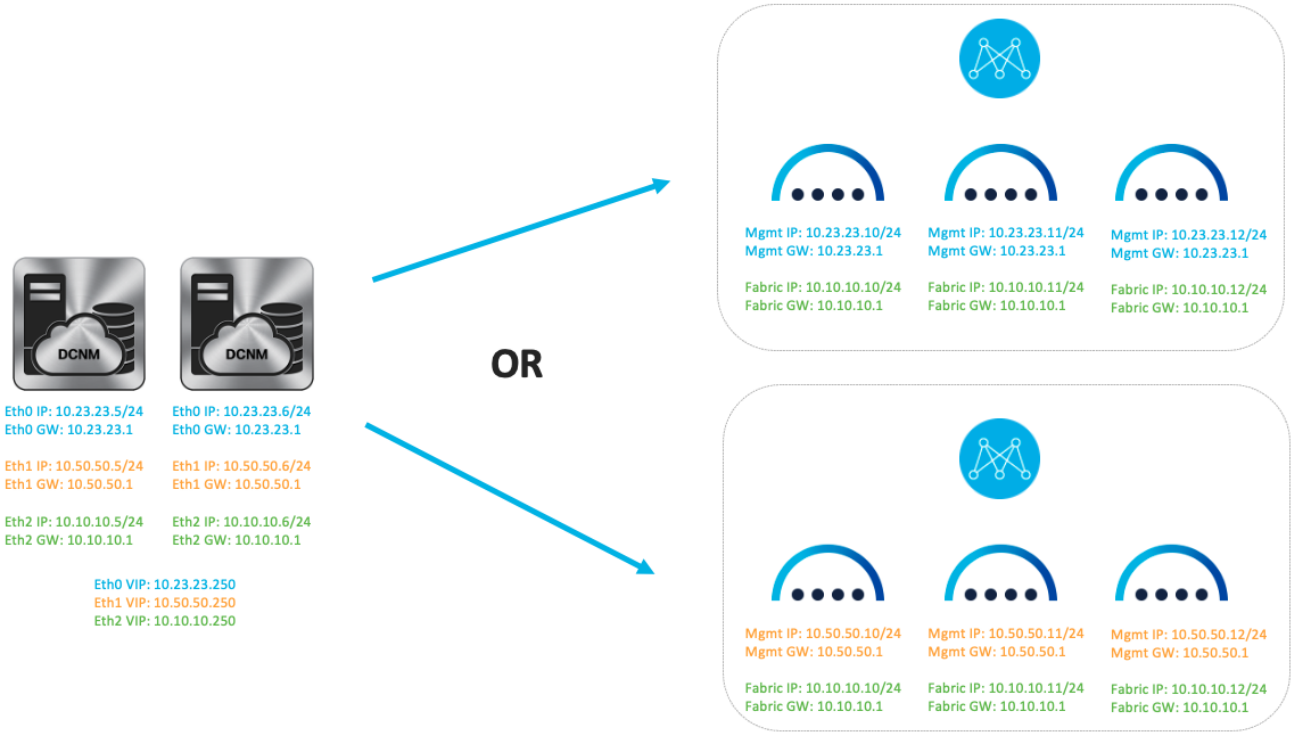

Figure 37. DCNM LAN Native HA to NDFC Migration Case 3

Case 4 - eth0/eth1 are in same subnet and eth2 is used for In-band

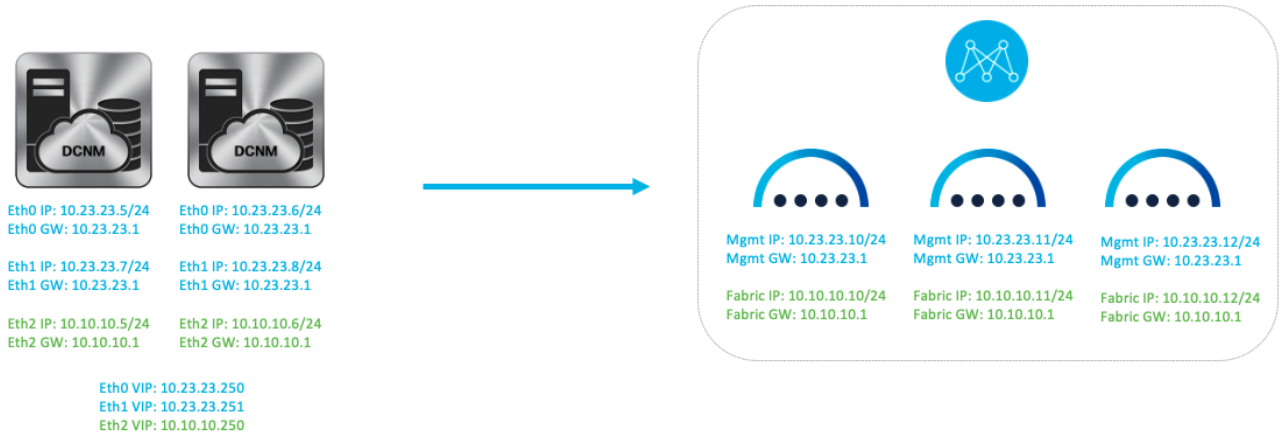

Figure 38. DCNM LAN Native HA to NDFC Migration Case 4

DCNM LAN Cluster to NDFC Migration

DCNM 11.x Cluster mode supports 350 managed and 750 monitor switches. NDFC 12.0.1a release supports 80 switches with 3 node deployment. NDFC 12.0.2 supports up to 350 switches (release time subjected to change) with 5 node vND or 3 node pND.

Case 1 - eth0/eth1 are in different subnets

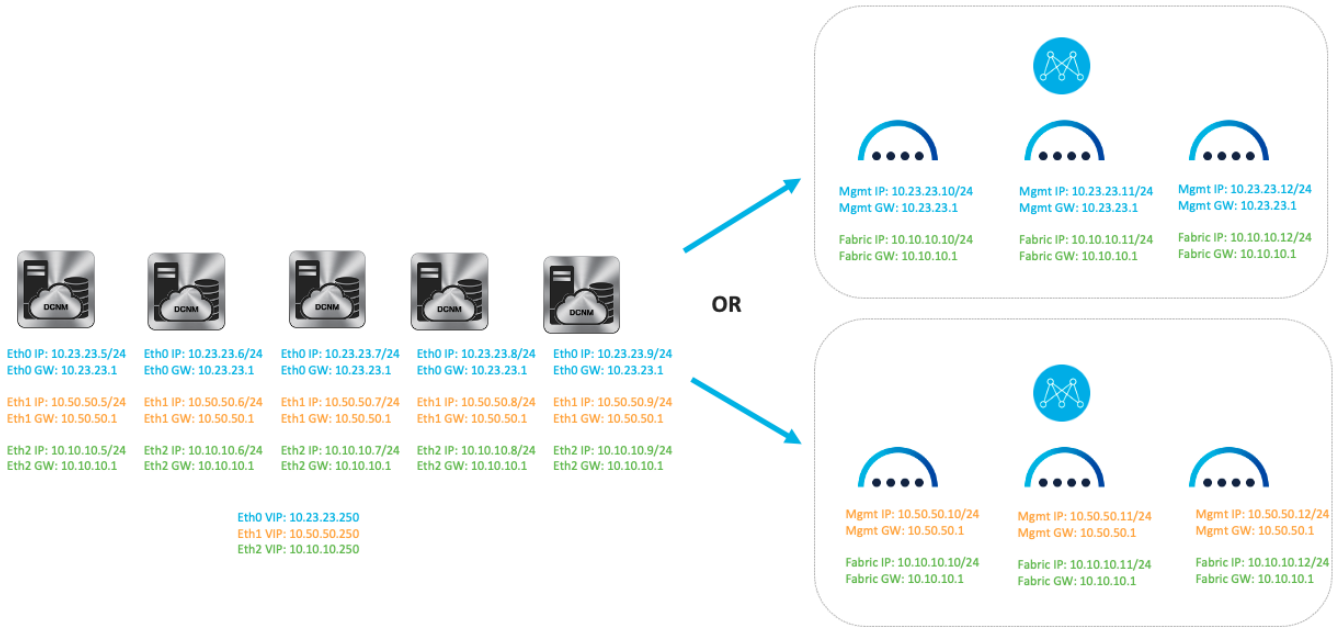

Figure 39. DCNM LAN Cluster to NDFC Migration Case 1

Case 2 – eth0/eth1 are in same subnet

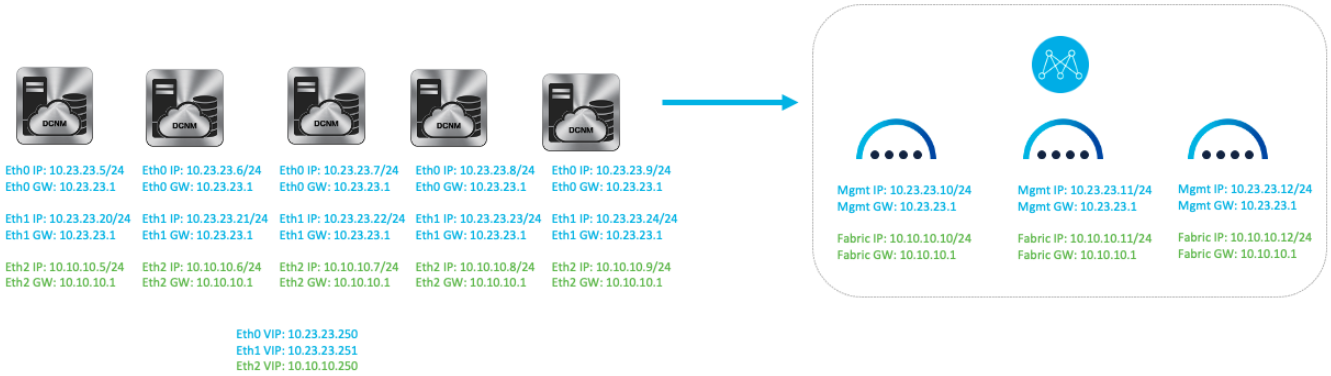

Figure 40. DCNM LAN Cluster to NDFC Migration Case 2

IPFM Fabrics

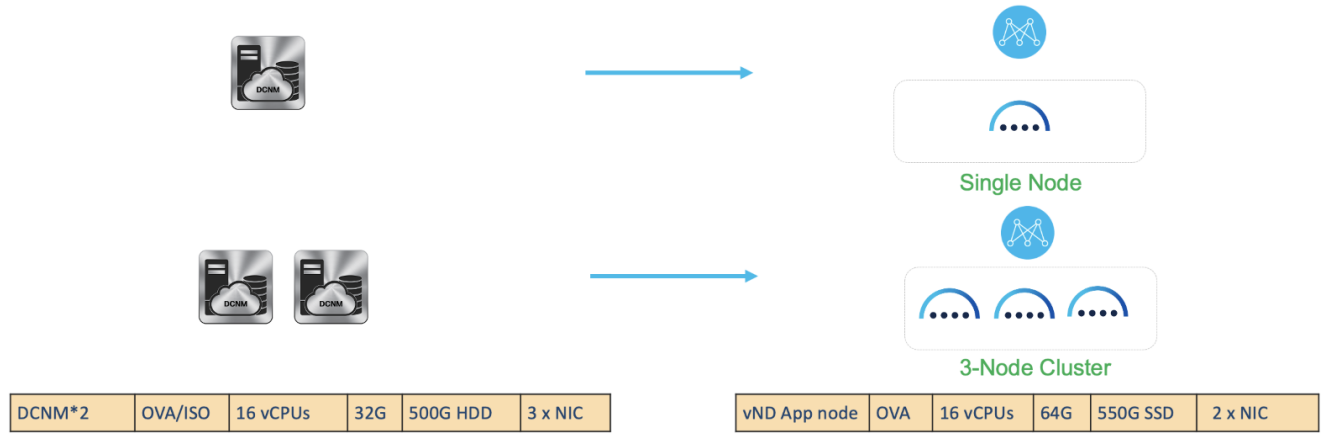

DCNM*2	OVA/ISO	16 vCPUs	32G	500G HDD	3 x NIC

vND App node	OVA	16 vCPUs	64G	550G SSD	2 x NIC

DCNM IPFM Standalone to NDFC Migration

This document does not cover the scale requirements for IPFM. See Additional Information section for links that provide access to the Cisco websites specific to IPFM scalability guide.

Case 1 – eth0/eth1 are in different subnets

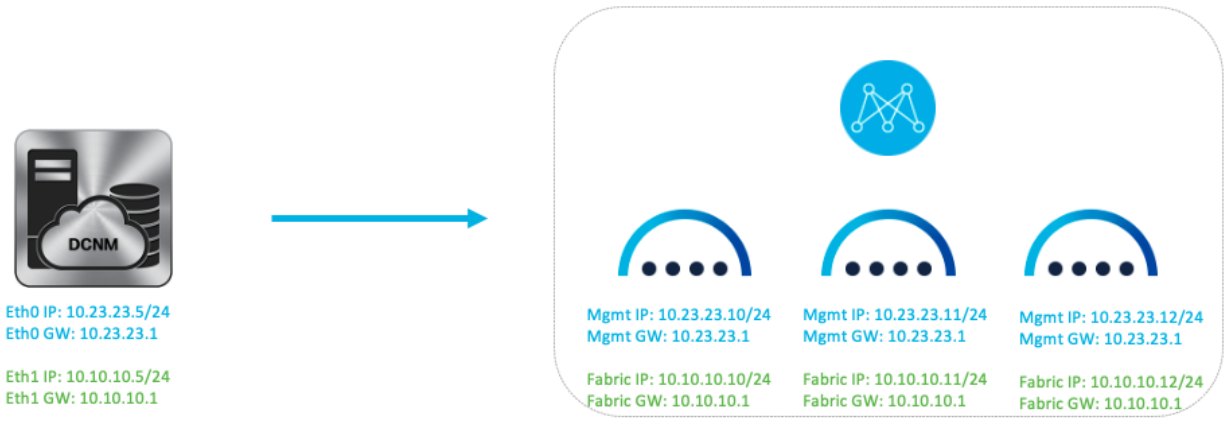

Figure 41. DCNM IPFM Standalone to NDFC Migration Case 1

Case 2 - eth0/eth1 are in same subnet

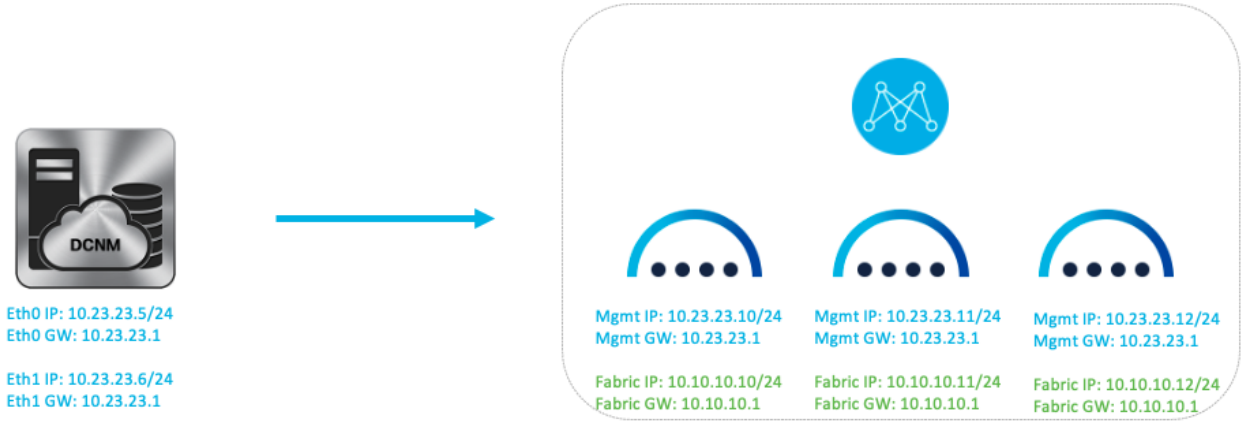

Figure 42. DCNM IPFM Standalone to NDFC Migration Case 2

DCNM IPFM Native HA to NDFC Migration

Case 1 - eth0/eth1 are in different subnets

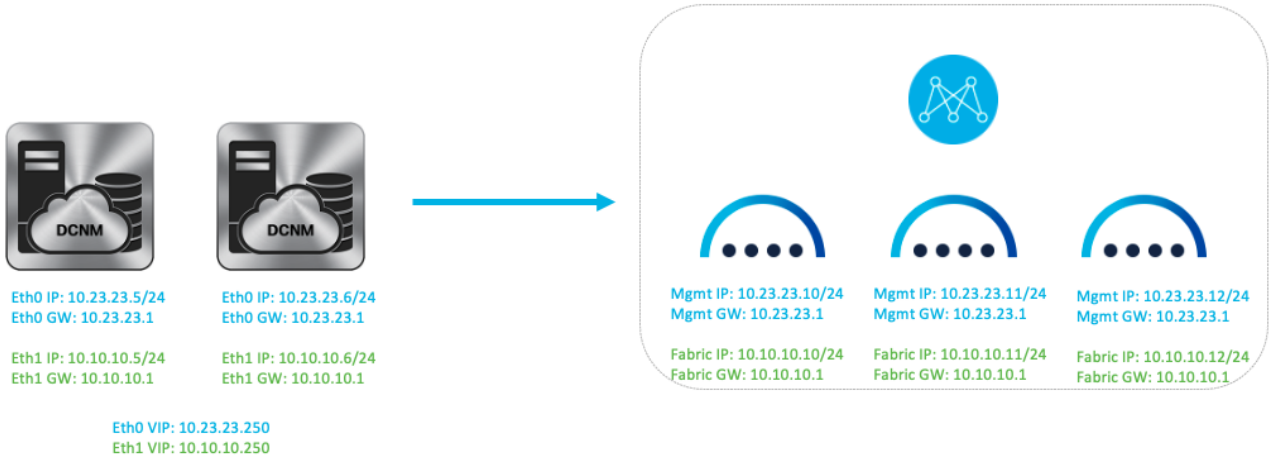

Figure 43. DCNM IPFM Native HA to NDFC Migration Case 1

Case 2 - eth0/eth1 are in same subnet

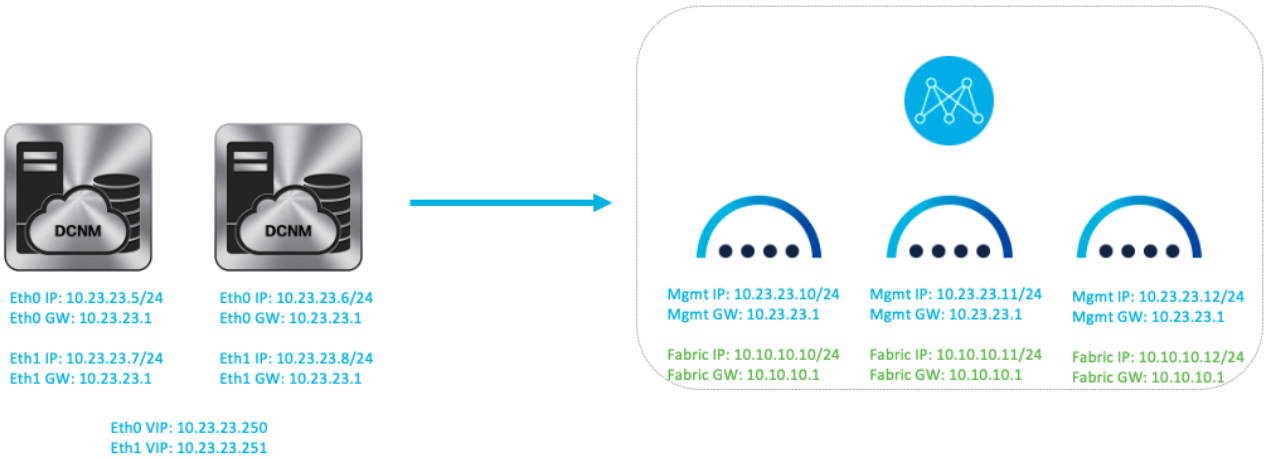

Figure 44. DCNM IPFM Native HA to NDFC Migration Case 2

SAN Fabrics

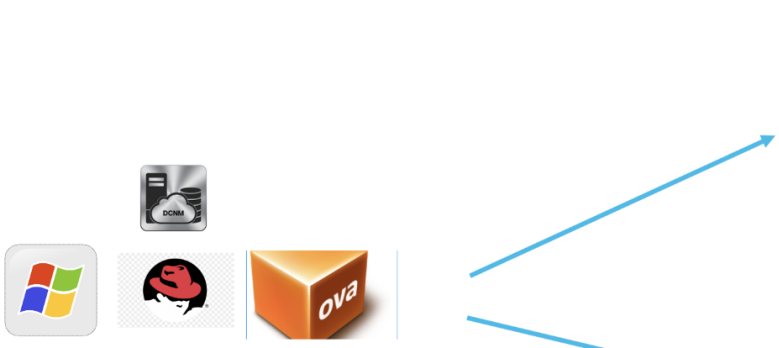

Single Node

OR

pND	SE Appliance

3-Node Cluster

DCNM-large	OVA/ISO	16 vCPUs	32G	500G HDD	2 x NIC
DCNM-huge	OVA/ISO	32 vCPUs	128G	2TB HDD	2 x NIC

vND App node	OVA	16 vCPUs	64G	550G SSD	2 x NIC
vND Data node	OVA	32 vCPUs	128G	3 TB SSD	2 x NIC

OR

pND*3	SE Appliance

This document does not cover the scale requirements for SAN. See Additional Information section for links that provide access to the Cisco websites specific to SAN scalability guide.

Case 1 – eth0 is used for Web and Device access, eth1 is no used

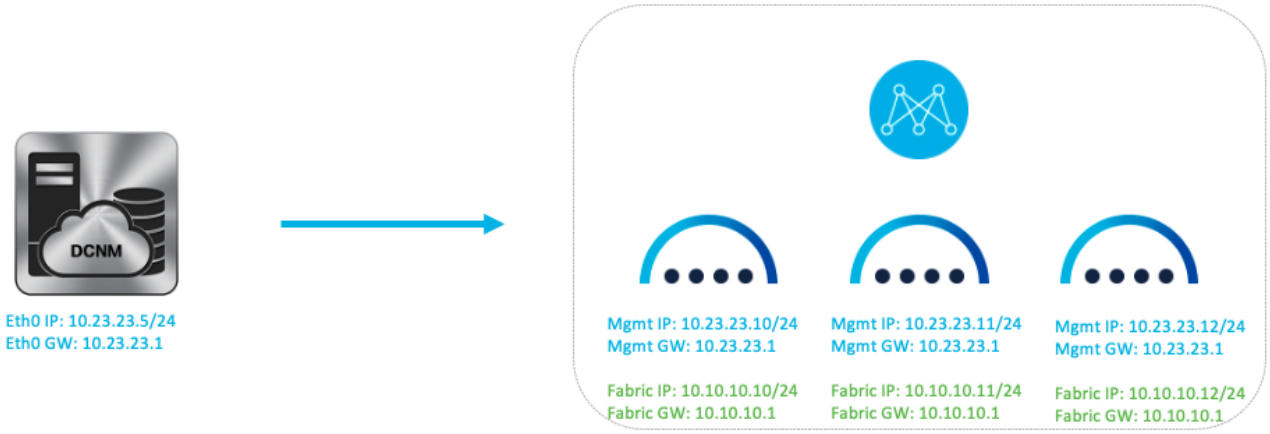

Figure 45. DCNM SAN to NDFC Migration Case 1

Case 2 – eth0 is used for Web access and eth1 is used for Device access

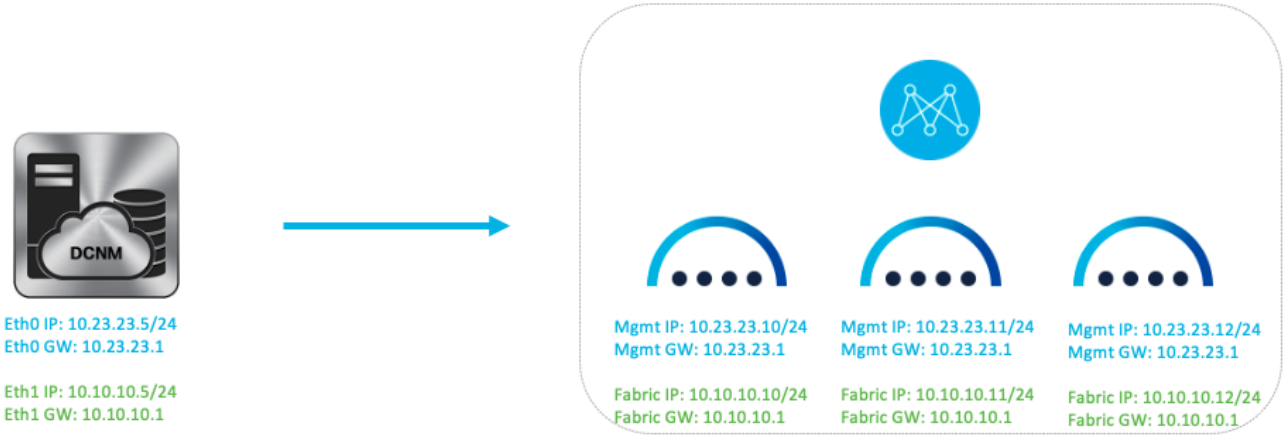

Figure 46. DCNM SAN to NDFC Migration Case 2

Additional Information

Additional documentation about Cisco Nexus Dashboard and Cisco Nexus Fabric Controller and related topics can be found at the sites listed here.

Nexus Dashboard

Deployment Guide: https://www.cisco.com/c/en/us/td/docs/dcn/nd/2x/deployment/cisco-nexus-dashboard-deployment-guide-211/nd-deploy-overview-21x.html

User Guide: https://www.cisco.com/c/dam/en/us/td/docs/dcn/nd/2x/user-guide/cisco-nexus-dashboard-user-guide-211.pdf

Release and Compatibility Matrix: https://www.cisco.com/c/en/us/support/data-center-analytics/nexus-dashboard/series.html#ReleaseandCompatibility

Nexus Dashboard Fabric Controller

Release Notes: https://www.cisco.com/c/en/us/td/docs/dcn/ndfc/1201/release-notes/cisco-ndfc-release-notes-1201.html

Compatibility Matrix: https://www-author3.cisco.com/content/en/us/td/docs/dcn/ndfc/1201/compatibility/cisco-ndfc-compatibility-matrix-guide-1201.html

Scalability Guide: https://www.cisco.com/c/en/us/td/docs/dcn/ndfc/1201/verified-scalability/cisco-ndfc-verified-scalability-1201.html

Configuration Guide: https://www.cisco.com/c/en/us/support/cloud-systems-management/prime-data-center-network-manager/products-installation-and-configuration-guides-list.html

Legal Information