



Overview

- [Understanding NDFC for Legacy/Classic Networks, on page 1](#)
- [Support for Brownfield Deployments, on page 7](#)
- [Guidelines and Limitations, on page 8](#)
- [Recommended Platform and Software Versions, on page 10](#)
- [Prerequisites, on page 10](#)

Understanding NDFC for Legacy/Classic Networks

This document describes how Cisco Nexus Dashboard Fabric Controller (NDFC) can be used to manage, maintain and monitor legacy/classic networks. In order to better understand how NDFC can be used to manage legacy/classic networks, it's useful to first understand all the pieces that are used in the process:

- The following sections describe the software components that are used in this process:
 - [Understanding ND, on page 1](#)
 - [Understanding NDFC, on page 2](#)
- [Understanding Legacy/Classic Deployments, on page 2](#) describes information about standard legacy/classic deployments.
- [Supported Legacy/Classic Network Topologies, on page 3](#) describes the two topology types that are supported when managing a legacy/classic network through NDFC.
- [Understanding Access-Aggregation Device Attachments, on page 6](#) provides the different options for attaching the Access devices to the Aggregation devices.
- [Understanding How NDFC Fabric Templates Are Used to Manage Legacy/Classic Networks, on page 6](#) describes how NDFC fabric templates are used to manage these types of standard legacy/classic deployments.

Understanding ND

Cisco Nexus Dashboard (ND) is a central management console for multiple data center sites and a common platform for hosting Cisco data center operation applications, such as Nexus Dashboard Insights and Nexus Dashboard Orchestrator. These applications are universally available for all the data center sites and provide real time analytics, visibility, and assurance for network policies and operations. Cisco Nexus Dashboard Orchestrator can also run on Nexus Dashboard as a hosted app.

Nexus Dashboard provides a common platform and modern technology stack for the above-mentioned micro services-based applications, simplifying the life cycle management of the different modern applications and reducing the operational overhead to run and maintain these applications. It also provides a central integration point for external third-party applications with the locally hosted applications.

Each Nexus Dashboard cluster typically consists of one or three `master` nodes. For three-node clusters, you can also provision a number of `worker` nodes to enable horizontal scaling and `standby` nodes for easy cluster recovery in case of a master node failure. For maximum number of `worker` and `standby` nodes supported in this release, see the "Verified Scalability Limits" sections of the [Cisco Nexus Dashboard Release Notes](#).

Understanding NDFC

Cisco Nexus Dashboard Fabric Controller, or NDFC (formerly known as Data Center Network Manager, or DCNM) is a service available exclusively on the Cisco Nexus Dashboard (ND) that uses a Kubernetes-based microservices architecture. NDFC provides comprehensive lifecycle management, configuration, and automation for a myriad of deployments using Cisco devices, such as NX-OS, IOS-XE, and IOS-XR devices, as well as non-Cisco devices.

In order to begin using NDFC, you must first have an ND cluster, where ND is deployed as a cluster of master and worker nodes in a virtual or physical form factor. The type and number of nodes required in a given cluster hosting NDFC depends on the scale of the managed switches, and whether NDFC will be used for LAN, SAN or Media Fabrics. It is possible to co-host NDFC with services like Insights in the same cluster, and to use NDFC for a variety of architectures at the same time, such as classic Ethernet and VXLAN.


You can use the [NDFC Capacity planning tool](#) to determine the number of ND nodes required for your scale. When the ND cluster is formed and healthy, NDFC can be installed from the Cisco App store, which is directly linked to the Nexus Dashboard. On enabling the service, the cluster will intelligently determine the resources required, depending on the scale and features enabled.

Understanding Legacy/Classic Deployments

Typically, a legacy/classic data center deployment consists of three tiers or layers, as described in [Data Center Multi-Tier Model Design](#):

- **Access layer:** The access layer provides the physical level attachment to the server resources. The access layer is the first oversubscription point in the data center because it aggregates the server traffic onto Gigabit EtherChannel or 10 GigE/10 Gigabit EtherChannel uplinks to the aggregation layer. Spanning tree routing protocols are extended from the aggregation layer into the access layer, depending on which access layer model is used.
- **Aggregation layer:** The aggregation layer is the Layer 3 and Layer 2 boundary for the data center infrastructure. Usually the aggregation layer is also the connection point for data center firewalls and other services. The aggregation layer, with many access layer uplinks connected to it, has the primary responsibility of aggregating the thousands of sessions leaving and entering the data center. The aggregation switches must be capable of supporting many 10 GigE and GigE interconnects while providing a high-speed switching fabric with a high forwarding rate. The aggregation layer also provides value-added services, such as server load balancing, firewalling, and SSL offloading to the servers across the access layer switches.
- **Core layer:** The core layer provides the interconnection of multiple data center aggregation devices, providing a fabric for high-speed packet switching between multiple aggregation modules. This layer serves as the gateway to the campus core where other modules connect, including the extranet, WAN, and Internet edge. All links connecting the data center core are terminated at Layer 3 and typically use 10 GigE interfaces for supporting a high level of throughput and performance, and to meet oversubscription levels.

The following figure shows an example of a data center multi-tier model topology.



Supported Legacy/Classic Network Topologies

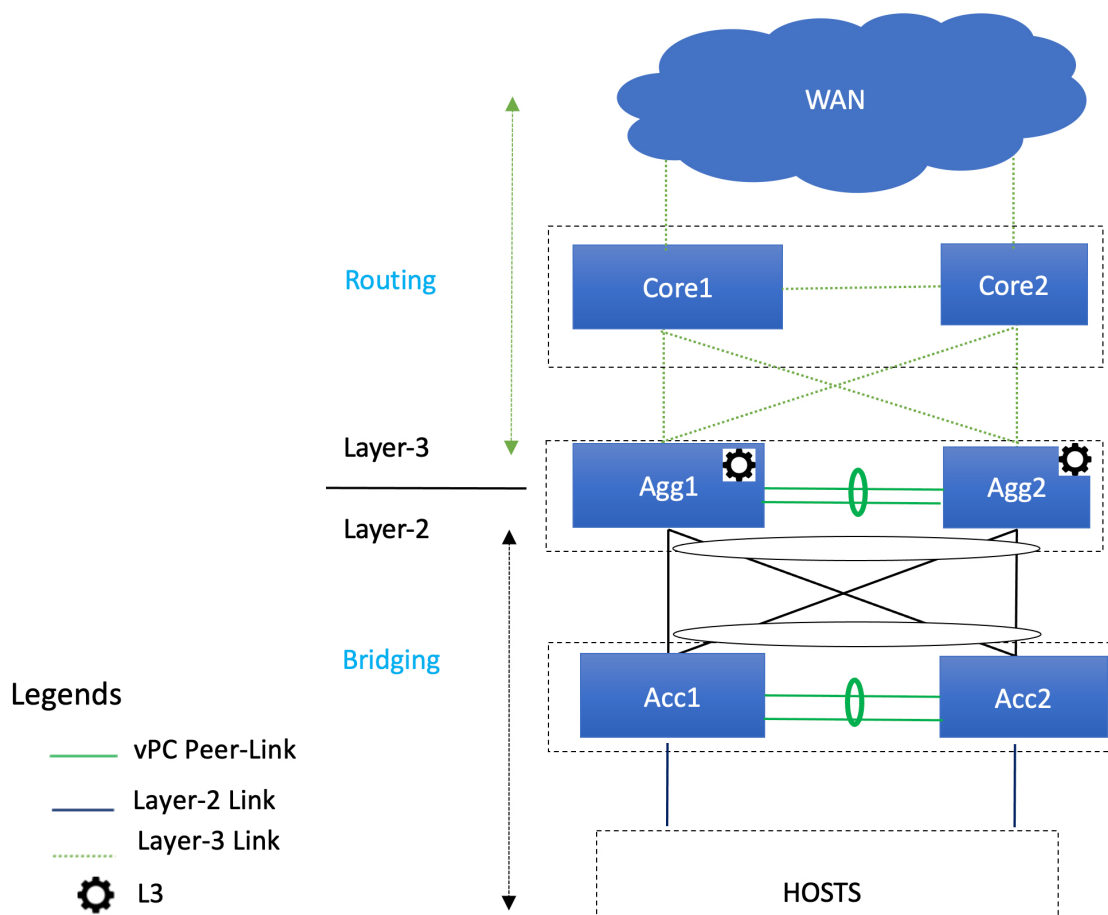
There are two types of legacy/classic network topologies that are supported, as described in the following sections.

Three-Tier Hierarchical Topology

In this type of legacy/classic network, the topology is split into three tiers:

- Access
- Aggregation
- Core

The following graphic shows an example of this three-tier topology.



For this topology:

- In this topology, the Layer 2/Layer 3 boundary occurs at the Aggregation level.
- You will use the Enhanced Classic LAN fabric template to configure the fabric that will be used for the Access and Aggregation tiers, where you will be setting roles for the switches in this fabric to one of the following roles:
 - **Access role:** Access devices are usually attached to the end hosts or servers and have a Layer 2 configuration. You can configure Access devices as a vPC pair for better redundancy.
 - **Aggregation role:** Aggregation devices are always deployed in an Enhanced Classic LAN fabric as a vPC pair. There can be multiple Aggregation vPC pairs in a single Enhanced Classic LAN fabric. Aggregation or distribution devices typically present the Layer-2/Layer-3 boundary, so you can enable the appropriate SVIs with your desired First Hop Redundancy Protocol (FHRP) at this layer. All routed (intra-subnet) traffic is forwarded through the Aggregation layer.
- You will use the the External Network Connectivity fabric template to configure the fabric that will be used for the Core tier, where you will be setting the role for the switches in this fabric to the **Core Router** or the **Edge Router** role.

See [Understanding How NDFC Fabric Templates Are Used to Manage Legacy/Classic Networks](#), on page 6 for more information on the two fabric template types.

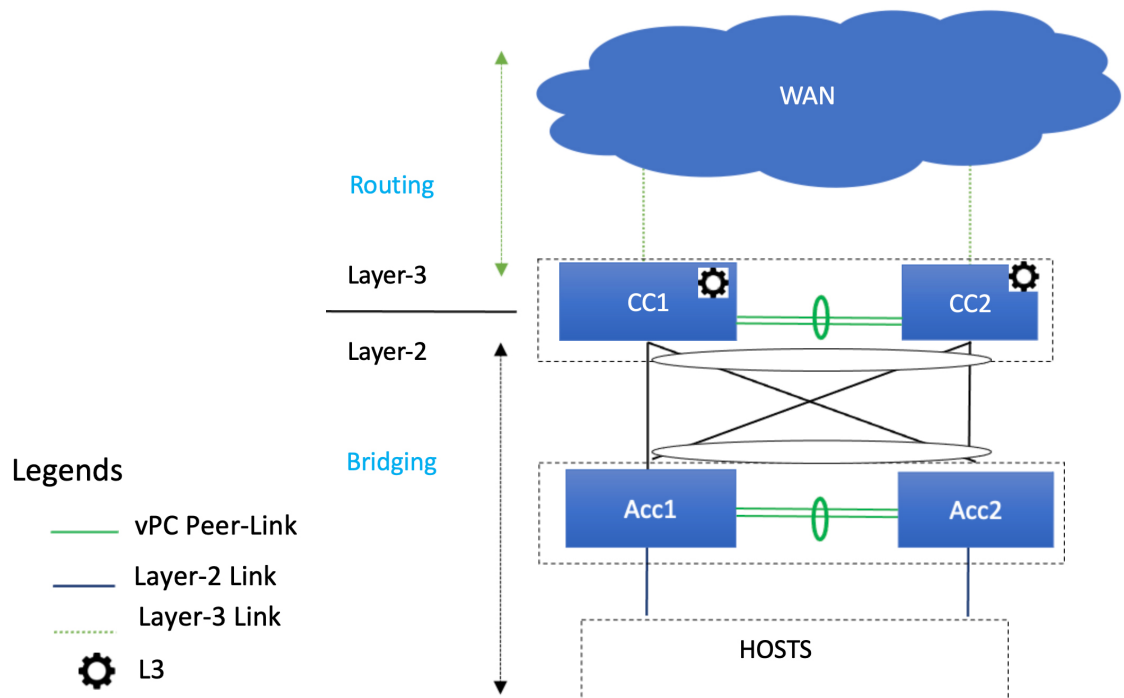


Note There is only one Aggregation pair shown in the preceding figure; however, for greenfield deployments, you could have multiple Aggregation pairs when you have this sort of topology managed by Cisco Nexus Dashboard Fabric Controller.

Two-Tier Collapsed Core Topology

In this type of legacy/classic network, the topology is split into two tiers:

- Access
- Collapsed Core



For this topology:

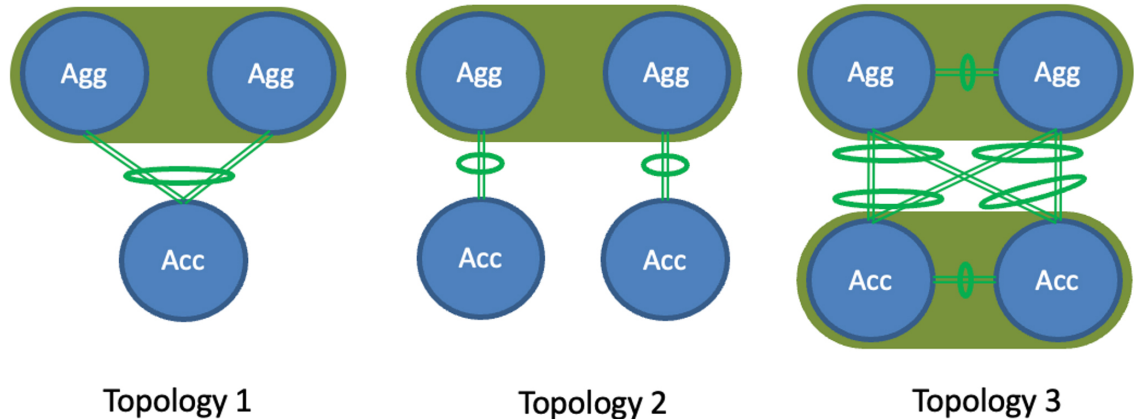
- In this topology, the Core and Aggregation tiers are collapsed into a single, combined tier, called the "Collapsed Core" tier. The Layer 2/Layer 3 boundary occurs at the Collapsed Core tier.
- You will use the Enhanced Classic LAN fabric template to configure the fabric that will be used for the Collapsed Core tiers. See [Understanding How NDFC Fabric Templates Are Used to Manage Legacy/Classic Networks, on page 6](#) for more information.

Switches Used At Each Tier

In typical legacy/classic networks using Cisco equipment, Cisco Nexus 7000 Series switches are deployed at the aggregation and core layers, and Cisco Nexus 5000 and 6000 Series switches are deployed at the access layer. You might also have Cisco Nexus 2000 Series Fabric Extender switches attached to either the aggregation or the access layers. Cisco Nexus 9000 and 3000 Series switches might be used at the core, aggregation or access layers; however, when using the new Enhanced Classic LAN fabric template in NDFC, only Cisco Nexus 7000 and 9000 Series switches are supported.

Understanding Access-Aggregation Device Attachments

For the three-tier hierarchical topology described in [Supported Legacy/Classic Network Topologies, on page 3](#), the Access devices can be attached to the Aggregation devices using one of the following options:



- **Topology 1:** vPC Aggregation pairing with the same Access (one Access device connected to both Aggregation devices), such as using a Fabric Extender in active-active (FEX-AA) mode for the Access device.
- **Topology 2:** vPC Aggregation pairing with a separate Access (each Access device connected to one of the Aggregation devices), such as using a Fabric Extender in straight through (FEX-ST) mode for the Access device.
- **Topology 3:** vPC Aggregation pairing with vPC Access through back-to-back vPC pairings.

For any of these Access-Aggregation device connections, when you select **Recalculate and Deploy** after you have finished entering the necessary configuration information in the Enhanced Classic LAN fabric template, NDFC will automatically detect the connectivity between the Access and Aggregation layer devices and will generate the appropriate configurations based on the supported topologies above that were detected.



Note If you have an Access-Aggregation device connection that does not fall into one of the supported topologies above, NDFC will return an error when you select **Recalculate and Deploy** after you have finished entering the necessary configuration information in the Enhanced Classic LAN fabric template.

Understanding How NDFC Fabric Templates Are Used to Manage Legacy/Classic Networks

As part of the process for managing legacy/classic networks through NDFC, you will use different NDFC fabric templates to create the fabrics that will be used for the different tiers:


- **Enhanced Classic LAN fabric template:** The Enhanced Classic LAN fabric template is used to configure the fabric for one of the following types of tiers, as described in [Supported Legacy/Classic Network Topologies, on page 3](#):
 - For a three-tier topology, you will use the Enhanced Classic LAN fabric template to configure the fabric that will be used for the Access and Aggregation tiers.

- For a two-tier (Collapsed Core) topology, you will use the Enhanced Classic LAN fabric template to configure the fabric that will be used for both tiers in that topology.
- **External Network Connectivity fabric template:** The External Network Connectivity fabric template is used to configure the fabric specifically for the Core tier in a three-tier topology, as described in [Supported Legacy/Classic Network Topologies, on page 3](#).

Support for Brownfield Deployments

You can import existing brownfield Enhanced Classic networks into NDFC, where all intent will be learned by NDFC and the configurations on the switches will be preserved, resulting in a non-disruptive operation. These imported networks can then be incrementally managed and maintained by NDFC.

For brownfield deployments, the following topologies are supported:



- **Topology 1:** vPC Aggregation pairing with the same Access (one Access device connected to both Aggregation devices), such as using a Fabric Extender in active-active (FEX-AA) mode for the Access device.
- **Topology 2a:** vPC Aggregation pairing with a separate Access (each Access device connected to one of the Aggregation devices), such as using a Fabric Extender in straight through (FEX-ST) mode for the Access device.
- **Topology 2b:** vPC Aggregation pairing with a separate Access (each Access device connected to one of the Aggregation devices), where Access devices are connected to the Aggregation peer through an ethernet trunk port.
- **Topology 3:** vPC Aggregation pairing with vPC Access through back-to-back vPC pairings.

For brownfield deployments, you must create an Enhanced Classic LAN fabric and set the fabric settings in accordance with your existing legacy 3-tier deployment. For example, if eBGP is used as a VRF Lite protocol between the Aggregation and Core layer, then you should make that selection when configuring the Enhanced Classic LAN fabric and provide the appropriate ASN.

Make the following additional configurations for brownfield deployments:

- You must set the appropriate spanning tree related parameters in the fabric settings.
- Disable NX-API if it's not required, since by default these options are enabled in the fabric settings.
- Import the switches into this Enhanced Classic LAN fabric using the **Preserve Config= Yes** option.

- Set the role of the Aggregation devices because, by default, all roles will be set to Access.

After you have made the necessary configurations in the Enhanced Classic LAN fabric template for the brownfield deployment, select **Recalculate and Deploy**. NDFC will then perform various pre-checks on the switches:

- Aggregation devices must be configured as a vPC pair; otherwise, NDFC will return an error.
- vPC consistency checks should indicate CONSISTENT on the vPC pairs. vPC pairs are mandatory at the Aggregation layer but are optional at the Access layer. If configured on the Access layer, the vPC pair should be consistent.
- Various topology checks will be performed to ensure that the current deployment being imported into the Enhanced Classic LAN fabric uses one of the supported connectivity options described earlier in this section. If any other topology is discovered, NDFC will return an error.
- The First Hop Redundancy Protocol (FHRP) configurations that you entered in fabric settings must match what is configured on the Aggregation layer switches.

Once all the pre-checks pass, as part of the brownfield import process, NDFC will collect the `show vlan` and `show vrf` related outputs from the Aggregation and Access layers:

- A VLAN is considered as a top-down VLAN if it is either Layer 2 only or has a Layer 3 SVI configured on the Aggregation layer with the appropriate FHRP as set in the fabric settings.
- All VRFs are considered as top-down VRFs.


All vPC pairing related information (such as the vPC domain, the vPC peer KPA, and the vPC peer link) will be learned for the Aggregation devices and, if applicable, the Access layer devices. In addition, all interface-related configurations will be learned during the brownfield import, such as configurations related to access, trunk, routed, subinterface, port-channels, vPCs, and so on. The port-channels and vPCs connected between the Aggregation and Access layers will be appropriately mapped to the `uplink_access` policies, along with the mapping of Access devices to Aggregation devices. In addition, for the network/VRF attachments, VRF Lite-related configurations will also be learned on the Enhanced Classic LAN fabric as part of the brownfield import.

Guidelines and Limitations

Following are the guidelines and limitations when configuring NDFC for legacy/classic networks:

- The Enhanced Classic LAN fabric does not support admin VDC on Cisco Nexus 7000 Series switches. The fabric only supports default and user VDC.
- Support is available for configuring greenfield or brownfield legacy/classic networks with NDFC.
 - For brownfield legacy/classic networks, it is assumed that you have a fully-functional, running fabric before bringing that brownfield fabric into NDFC to manage.
 - A legacy/classic network configured with Cisco Fabric Path is not supported as a brownfield import into NDFC.
- Inband management and inband POAP is not supported for devices in the Enhanced Classic LAN fabric type.

- You must have vPC configured at the following tiers:
 - At the Aggregation tier for the 3-tier topology
 - At the Collapsed Core tier for the 2-tier topology
- For greenfield deployments, the following guidelines and limitations apply:
 - You can have multiple Aggregation pairs configured at the Aggregation tier for the 3-tier topology; however, you must configure vPC for each Aggregation pair.
 - The following topologies are supported for a brownfield deployment, as described in [Support for Brownfield Deployments, on page 7](#):



However, for a greenfield deployment, Topology 2b is not supported.

If you are configuring a **Topology 2a** type of deployment for a greenfield deployment, follow these steps to distinguish this configuration from the **Topology 2b** configuration that is not supported on greenfield deployments:

1. In the **Fabric Overview** page for the Enhanced Classic LAN fabric, click the **Switches** tab, then click on one of the Aggregation switches, which would be part of a vPC pair, and click **Actions > ToR/Access Pairing**.
2. In the **ToR/Access Pairing** window, select the Access switches and click **Save**.



Note Do not select the **Complete ToR/Access Pairing as VPC Pair** checkbox.

3. In the **Fabric Overview** page, click **Actions** at the top of the page and select **Recalculate and Deploy**.

Wait for the deployment to complete and the configuration status for the switches to turn to **In-Sync**, then continue with your configurations steps as you normally would.

- The following switches are supported with the Enhanced Classic LAN template:
 - Cisco Nexus 7000/9000 switches
 - Cisco Nexus 2000 Series Fabric Extender switches
- The following topologies are supported:

- 3-tier topology (Access/Aggregation/Core tiers)
- 2-tier Collapsed Core topology, where:
 - The first tier (the Access tier) is the same as the Access tier in the 3-tier topology
 - The second tier (the Collapsed Core tier) combines the Aggregation and Core tiers from the 3-tier topology into a single unified layer

Recommended Platform and Software Versions

Following are the recommended versions for the platforms and software:

- Recommended software versions:
 - **Nexus Dashboard:** Release 3.0.1 and later
 - **NDFC:** Release 12.1.3 and later
- Recommended platforms:
 - **Access and aggregation layers:**
 - Cisco Nexus 9000 Series switches (recommended release version: NX-OS release 9.3.11 and later)
 - Cisco Nexus 7000 Series switches
 - Cisco Nexus 2000 Series Fabric Extender switches
 - **Core layer:**
 - Cisco Network Convergence System (NCS) 5500 Series
 - Cisco Nexus 7000 or 9000 Series switches
 - Cisco Catalyst 9000 Series switches
 - Cisco ASR 9000 Series Aggregation Services Routers

Prerequisites

Before configuring NDFC with legacy/classic networks, complete the following prerequisite tasks to configure the Nexus Dashboard cluster and verify that the NDFC service is healthy:

1. Configure the virtual or physical Nexus Dashboard (ND) nodes to form a cluster.
 - Use the [Nexus Dashboard Capacity Planning tool](#) to determine the number of nodes per form factor and the supported scale for your deployment.
 - The cluster nodes can be Layer 2 or Layer 3 adjacent on the data interface.
 - We recommend having a standby node for high availability purposes.

Deploy the Nexus Dashboard using the procedures provided in the [Cisco Nexus Dashboard Deployment Guide](#), release 2.3.x or later.

2. Configure the Nexus Dashboard Fabric Controller (NDFC).

Install the NDFC service on the ND cluster that you configured in the previous step using the procedures provided in the [Cisco Nexus Dashboard Fabric Controller Installation and Upgrade Guide](#), release 12.1.2 or later, and the [Cisco Nexus Dashboard Fabric Controller Deployment Guide](#).

3. Verify the reachability between the NDFC service and the switches that you will be managing through NDFC.

- In Nexus Dashboard, navigate to **Infrastructure** > **Cluster Configuration** to define the appropriate routes for the reachability of the switches from the ND cluster.

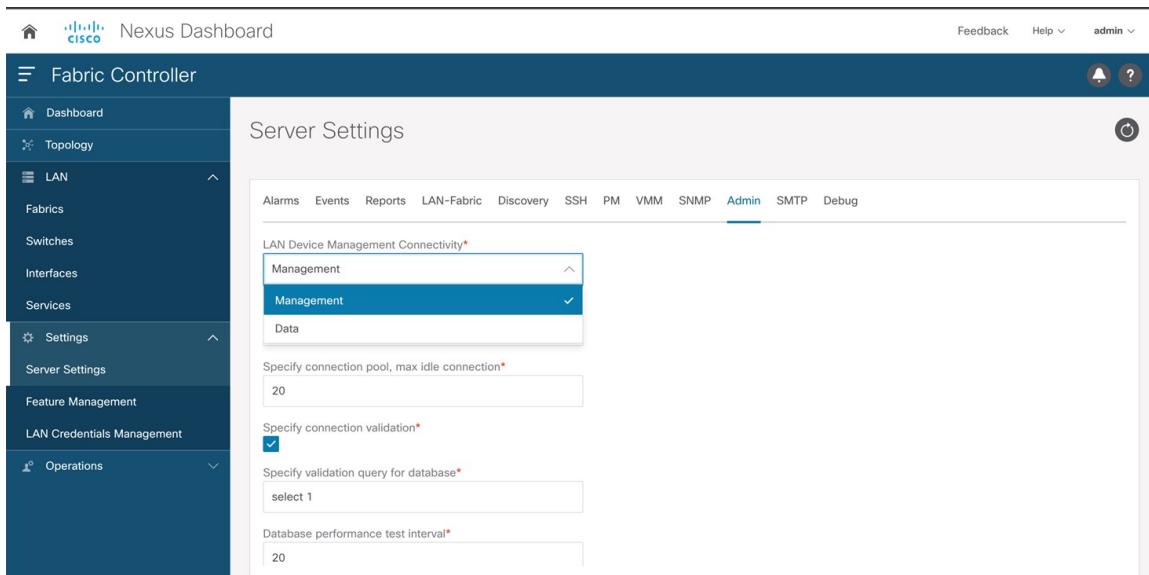
The screenshot shows the Cisco Nexus Dashboard Admin Console interface. The left sidebar contains navigation options: Overview, Sites, Services, System Resources, Operations, Infrastructure, Cluster Configuration, Resource Utilization, Intersight, App Infra Services, and Administrative. The main content area is titled 'Cluster Configuration' and shows 'Routes' for 'Management Network Routes' and 'Data Network Routes'. Below the routes, there are two circular gauges for 'External Service Pools': 'Management Service IP Usage' (0 Total) and 'Data Service IP Usage' (2 Total, with 2 in use). A table below shows 'Data Service IP's' with columns for IP, Usage, and Assignment.

IP	Usage	Assignment
192.93.0.174	In Use	cisco-ndfc-dcnm-syslog-trap-data
192.93.0.175	In Use	cisco-ndfc-dcnm-poap-data-http-ssh

- Define the External Service Pools for SNMP and POAP over the management or data subnet.

Because classic LAN only supports the out-of-band management of switches, configure the following setting in NDFC:

- In Nexus Dashboard Fabric Controller, navigate to **Settings** > **Server Settings**, then click the **Admin** tab.
- In the **LAN Device Management Connectivity** field, choose **Management**.



4. Configure the necessary features in NDFC.

In NDFC, click **Feature Management** and click the **Fabric Controller** box, then make the following selections in the **Feature Name** column:

- **Fabric Builder:** Mandatory, to allow for easy fabric functionality for NX-OS and other devices.
- **Performance Monitoring:** Optional, to monitor environment and interface statistics.

