

Enhanced Classic LAN in Cisco Nexus Dashboard Fabric Controller (NDFC) Release 12.1.3

Contents

Introduction	3
What is Cisco Nexus Dashboard Fabric Controller?	4
Use Cases of Classic LAN.....	4
Use Case 1: Classic LAN Networks (Brownfield and Greenfield)	5
Use Case 2: Coexistence of Brownfield Classic LAN and Greenfield VXLAN	5
Use Case 3: Journey from On-premises to Cloud	5
Cisco Nexus Dashboard Fabric Controller for Classic LAN	5
Topologies supported for Classic LAN	7
Topology 1: 3 Tier Hierarchical (Access, Aggregation, Core)	7
Topology 2: Collapsed Core (Access, Core)	7
External Connectivity from Enhanced Classic LAN	13
Guidelines and Limitations for Enhanced Classic LAN	13
Hardware and Software Recommendations	14
Using Enhanced Classic LAN	14
Prerequisites	14
Day 0 for Classic LAN	17
For the Access and Aggregation Layers	18
Step 1: Create the Fabric	19
Step 2: Discover the Switches in the Fabric	22
Step 3: Bootstrap (Power-on Auto-provisioning)	25
Step 4: Define the Roles	31
Step 5: Configure the vPC pairing	32
Step 6: Recalculate and Deploy	36
For the Core Layer	40
For a Group of Fabrics	41
Day 1 for Classic LAN	43
Layer 2 Network	46
Step 1: Create the Network	46
Step 2: Attach the Network	47
Step 3: Review Pending Configurations on Access and Aggregation	48
Step 4: Deploy the Configuration	49
Layer 3 Network in the Default VRF Instance	49
Step 1: Create the Network	50
Step 2: Attach the Network and Choose the FHRP Master per Network	51
Step 3: Review Pending Configurations on the Access and Aggregation Layer.....	51
Step 4: Deploy the Configuration	53

Layer 3 Network with a Custom VRF Instance	53
Step 1: Create the Network	54
Step 2: Create a VRF Instance to Link a Custom VRF Instance to This Layer 3 Network	55
Step 3: Attach the Network	57
VRF-Lite Extension Between the Aggregation and Core/Edge Layers	59
Step 1: Fabric Settings	59
Step 2: VRF Attachments	61
Step 3: Deploy on Aggregations	63
Step 4: Deploy on Core	64
VRF-Lite extension between Collapsed Core and WAN	65
Day 2 for Classic LAN	65
Integration of Classic LAN with Services	66
VRF-Lite Using Subinterfaces	67
VRF-Lite Using SVIs	73
VRF-Lite Using Routed Interfaces or Port Channels	76
Migration from Cisco Nexus 2000/5000/7000 Classic LAN networks to Cisco Nexus 7000/9000- based Classic LAN Networks	78
NDFC with Legacy Nexus Platforms	78
Host Port Resynchronizing	81
NDFC with Newer Cisco Nexus Platforms (Cisco Nexus 9000) Considering FEX	84
Layer 2/Layer 3 Demarcation at the Core Layer	86
Migration from Classic LAN and VXLAN Networks	88
Hybrid Cloud Connectivity with NDFC	92
Conclusion.....	92

Introduction

Given the broad deployment of classic hierarchical networks and the CLI-driven methodology to push changes, data centers need a simplified, automated, SDN driven approach for day 0, day 1, and day 2 aspects of such Ethernet-based networks. Such networks typically consist of Access, Aggregation and Core layers.

The Access layer is where the server chassis are connected. The Aggregation layer is the Layer 3 and Layer 2 boundary for the data center infrastructure. In common designs, the Aggregation layer is also the connection point for data center firewalls and other services. The Core layer provides the interconnection of multiple data center Aggregation modules. While Layer 2/Layer 3 demarcation is typically at the Aggregation layer, in some cases the demarcation can be at the Access layer (routed access networks, for example) or Core layer.

This document is a deep dive into Classic Ethernet networks and how the Enhanced Classic LAN fabric type introduced in Cisco Nexus Dashboard Fabric Controller (NDFC) release 12.1.3 can manage, maintain,

and monitor them. This whitepaper covers the end-to-end deployment of switches in such networks, the prerequisites to start using NDFC, and the Nexus hardware recommendations at each layer.

What is Cisco Nexus Dashboard Fabric Controller?

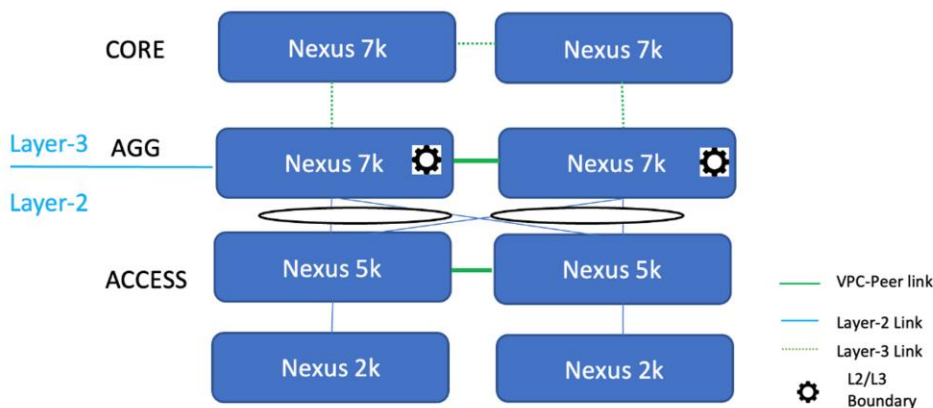
Cisco Nexus Dashboard Fabric Controller (NDFC), formerly known as Data Center Network Manager (DCNM), is the new, power-packed, feature-rich service exclusively available on the Cisco Nexus Dashboard compute platform. NDFC embraces a Kubernetes-based microservices architecture. With the introduction of NDFC release 12, you get a consistent experience across NDFC and other services hosted on Nexus Dashboard, such as Insights and Orchestrator. NDFC provides comprehensive lifecycle management, configuration, and automation for Cisco NX-OS, IOS-XE, IOS-XR, and non-Cisco devices for many deployments.

To begin with NDFC, you first need a Nexus Dashboard cluster. Nexus Dashboard 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 Nexus Dashboard cluster hosting NDFC depend on the scale of managed (or monitored) switches and whether NDFC will be used for LAN, SAN or Media Fabrics. It is also possible to co-host NDFC with services like Insights in the same Nexus Dashboard cluster and use NDFC to manage a variety of architectures concurrently such as Classic Ethernet and VXLAN.

You can use the [NDFC Capacity planning tool](#) to determine the number of Nexus Dashboard nodes required for your scale. After the Nexus Dashboard cluster is formed and healthy, you can install NDFC from the Cisco App Store, directly linked to the Nexus Dashboard. Upon enabling the NDFC service, the cluster intelligently determines the resources required, depending on the scale and features enabled.

Use Cases of Classic LAN

Cisco-based data center deployments consist of all ranges of Nexus platforms in the Cisco Nexus 2000 (aka Fabric Extender or FEX), 5000, 6000, 7000, 7700, and 9000 series. FEXes, were introduced to provide scalability and ease of management. These were seen connected to the Cisco Nexus 5000 and 6000 platforms acting as the Access switches, with the Cisco Nexus 7000/7700 switches as the Aggregation and the Core switch as shown in the following figure:



As time passed, and as newer, more powerful platforms came into play, the FEXes were replaced with the Cisco Nexus 9300 Cloud Scale series - a comprehensive line of switches with options including 1/10/25/50/100/400G multispeed ports and lower per-port cost. These 9ks were easy to manage even when no FEXes were in play. The Access and Aggregation layers started to see more penetration of the Cisco Nexus 9000 switches. Cisco Nexus 7000 and 7700 switches continues to be used as powerful Core

boxes. The Access, Aggregation, and Core layers with different configurations and connections are the major focus and use case of Classic LAN.

Most such Classic LAN deployments are yet to adapt to an SDN approach. Also, technologies like VXLAN are a future evolution, currently on the roadmap to be adopted and deployed by these customers. This is because, for such customers, there has not been a pressing need for overlays by their applications.

The use cases fall into three main categories (keeping in mind NDFC acting as the controller for these use cases)

Use Case 1: Classic LAN Networks (Brownfield and Greenfield)

This use case entails the management of single or multiple 2 or 3 Tier Hierarchical networks comprised of Nexus platforms. Existing networks managed by CLI or other mechanisms can be imported into NDFC with full support for Brownfield (i.e all intent will be learned by NDFC, and configurations on switches will be preserved), making this a non-disruptive operation. These networks can then be incrementally managed and maintained by NDFC. For any new deployments (Greenfield), Cisco best practice embedded templates in NDFC can be leveraged to provide end-to-end network connectivity.

Use Case 2: Coexistence of Brownfield Classic LAN and Greenfield VXLAN

This use case entails the coexistence of Classic LAN and VXLAN EVPN networks – A hybrid of vPC/Spanning Tree, 3-tier architecture, and an overlay-based Leaf-Spine VXLAN architecture, all within the same NDFC cluster. This option is for customers who plan to migrate workloads to an evolved VXLAN network but are currently on Classic Ethernet. NDFC can be leveraged for brownfield import of these existing Classic networks and manage them incrementally. NDFC templates can be used to build VXLAN BGP EVPN underlay and overlay from scratch (Greenfield). Once both architectures are up and running, NDFC can be used to migrate workloads from Classic to VXLAN networks. Classic LAN can thereafter be deprecated once all the migration is complete and customers are comfortable with VXLAN as a technology.

Use Case 3: Journey from On-premises to Cloud

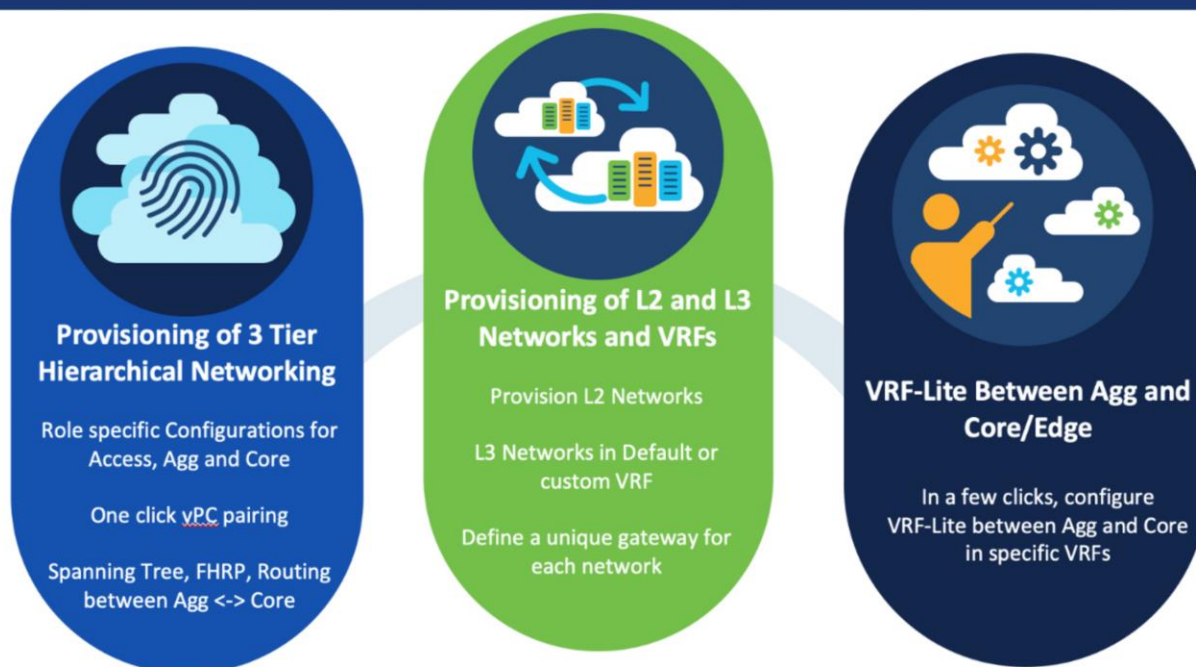
For customers planning to or already have workloads in the Cloud (AWS, Azure), NDFC has Hybrid Cloud support to provide seamless connectivity between on-premises and Cloud networks. This works in conjunction with Cisco Nexus Dashboard Orchestrator (NDO), which acts as a central policy orchestrator, and Cisco Cloud Network Controller (CNC), which is spun up as a SaaS instance to provide connectivity and policy translation for workloads in the Cloud. The hybrid cloud support is between VXLAN fabrics on-premises and Cloud networks. Hence, having VXLAN fabric(s) that will replace Classic Ethernet networks will be an evolution in this journey. NDFC helps provision the required secure underlay and overlay for this connectivity to Cloud either using IPsec over public networks or without IPsec in the presence of Direct Connect (AWS) or Express Route (Azure).

This whitepaper delves into each of the use cases.

Cisco Nexus Dashboard Fabric Controller for Classic LAN

The NDFC 12.1.3 release introduces a new fabric template called "Enhanced Classic LAN". This template is introduced to completely automate the Layer 2 and Layer 3 aspects of Access-Aggregation-Core, as per Cisco best practices. This minimizes the learning curve and makes it easy to move to an SDN-driven approach, all while preparing for the future by improving scalability, creating the opportunity to build overlays with VXLAN, and offering a wide variety of maintenance and operational features (discussed in the [Day 2 for Classic LAN](#) section).

Introducing Enhanced Classic LAN in NDFC 12.1.3



The following list contains examples of protocols that Enhanced Classic LAN provisions (as of the NDFC 12.1.3 release), keeping in mind Cisco recommended best practices and configurations for 3-Tier Access-Aggregation-Core and Layer 2/Layer 3 demarcation at Aggregation or Collapsed Core:

1. Routing protocols between Aggregation and Core/Edge layers: eBGP, OSPF, None

You can use the **None** option when:

- No routing is present (Access and Aggregation are both Layer 2)
- Using static routing/IS-IS between the Aggregation and Core/Edge layers (you can configure this using NDFC policies)

2. Spanning tree (with the Aggregation layer as bridge and root): RVPST+, MST, Unmanaged
3. FHRP (at the Aggregation layer): HSRP, VRRP, VRRPv3, None

You can use the **None** option when both the Access and Aggregation layers are Layer 2.

4. One-click vPC pairing: Between Aggregation pairs, between Access pairs, B2B vPC between Access and Aggregation (you can use the auto-pairing option)

NDFC's Enhanced Classic LAN template needs a few inputs from the user to learn their intent. At the time of fabric creation, the user must select the protocols of choice or stick with the default. Customizations for each protocol are possible under Fabric Settings. Once the fabric has been created and the respective switches have been discovered within this fabric, NDFC learns the topology and how the switches are connected. The user must thereafter specify roles for each switch. After the role definition, NDFC pushes respective configurations on 'Recalculate and Deploy'. After this, all configurations can be managed from a

Single Pane of Glass with the ability to do rollback at a granular level with the help of the "Change Control and Rollback" feature in NDFC.

This document covers all the items mentioned above in the [Day 0 for Classic LAN](#) and [Day 1 for Classic LAN](#) sections.

Topologies supported for Classic LAN

Two broad categories of topologies are considered for classic Ethernet in this whitepaper. The provisioning of these topologies is discussed in later sections.

Legends

- vPC Peer-Link
- Layer-2 Link
- - - Layer-3 Link
- L3

Topology 1: 3 Tier Hierarchical (Access, Aggregation, Core)

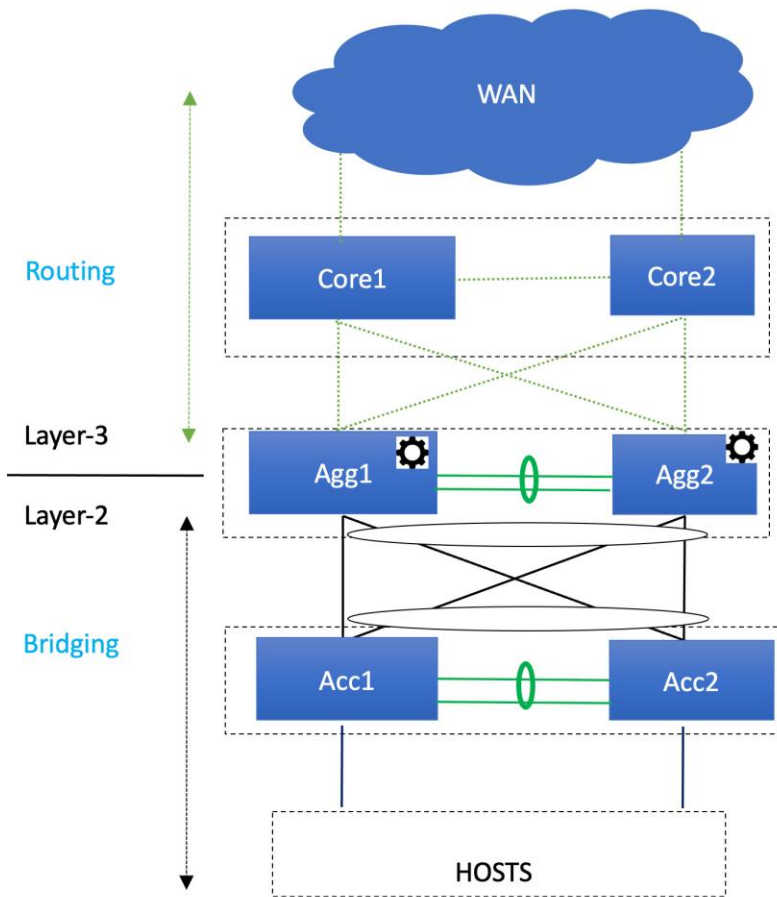



Figure 1. Layer 2/Layer 3 Boundary at Aggregation

Topology 2: Collapsed Core (Access, Core)

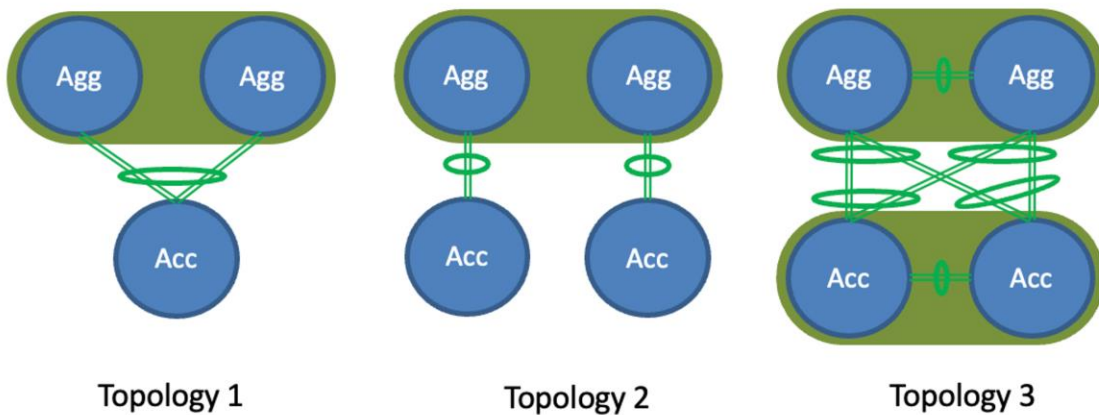
Layer 2/Layer 3 boundary at the collapsed Aggregation/Core or Edge layer.

This is the same as topology 1, except that Core and Aggregation layers are combined into a single collapsed layer.



Aggregation/Collapsed Core typically present the Layer 2/Layer 3 boundary so that you can enable the appropriate SVIs with a first hop redundancy protocol (FHRP) of choice at this layer. All routed or intra-subnet traffic is forwarded through the Aggregation layer.

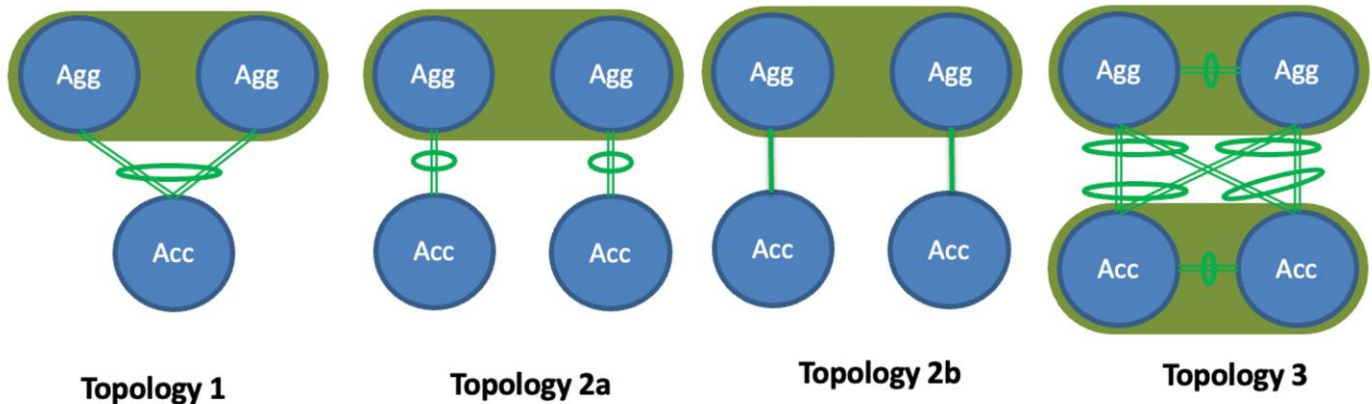
The following illustration shows the supported topologies for connectivity between the Access and Aggregation layers for a greenfield deployment:



- Topology 1: vPC domain at the Aggregation layer with the same Access switch connected with a port channel to both Aggregation switches.


- Topology 2: vPC domain at the Aggregation layer with Access switches connected in “straight-through” mode to a single Aggregation switch.
- Topology-3: vPC domain at the Aggregation with back to back vPC connectivity to a pair of access switches.

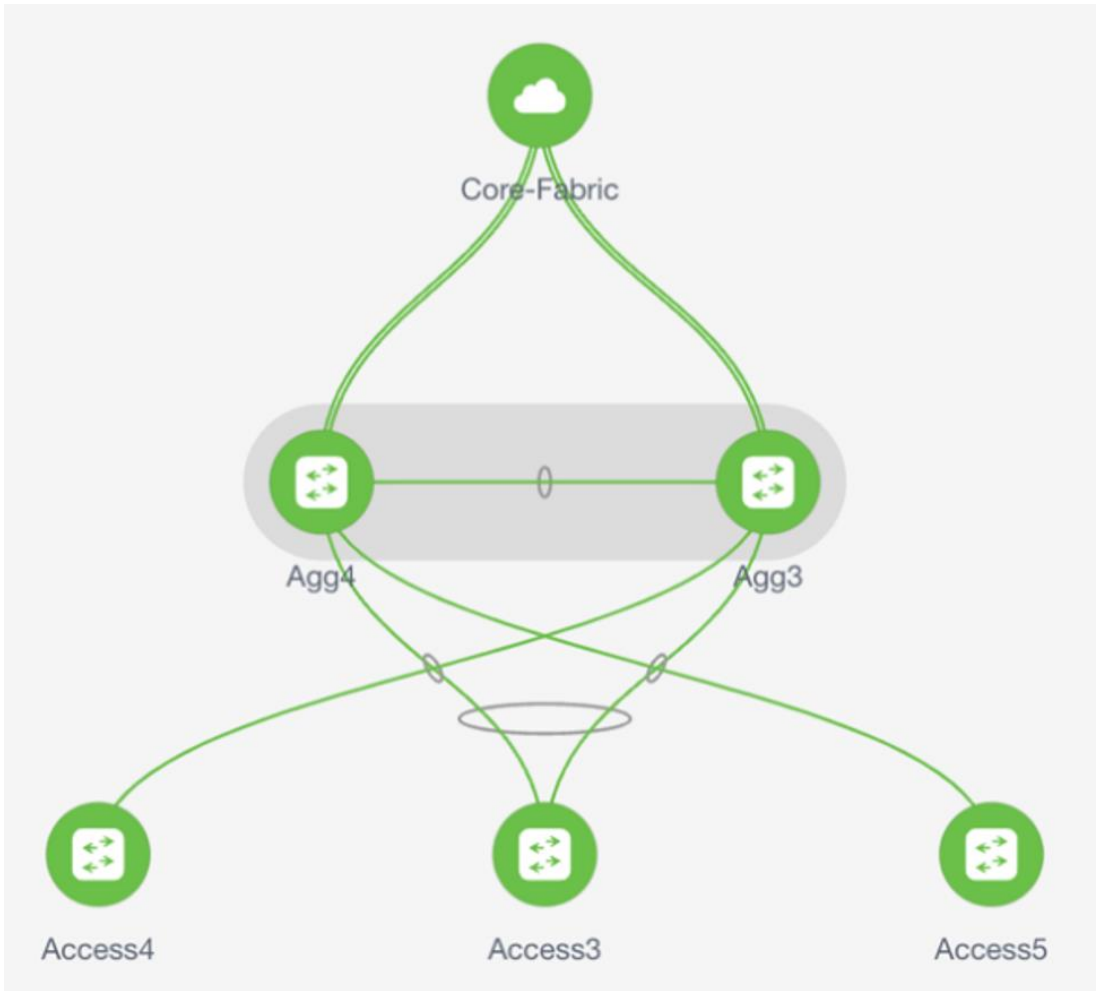
The following illustration shows the supported topologies for connectivity between the Access and Aggregation layers for a brownfield deployment:

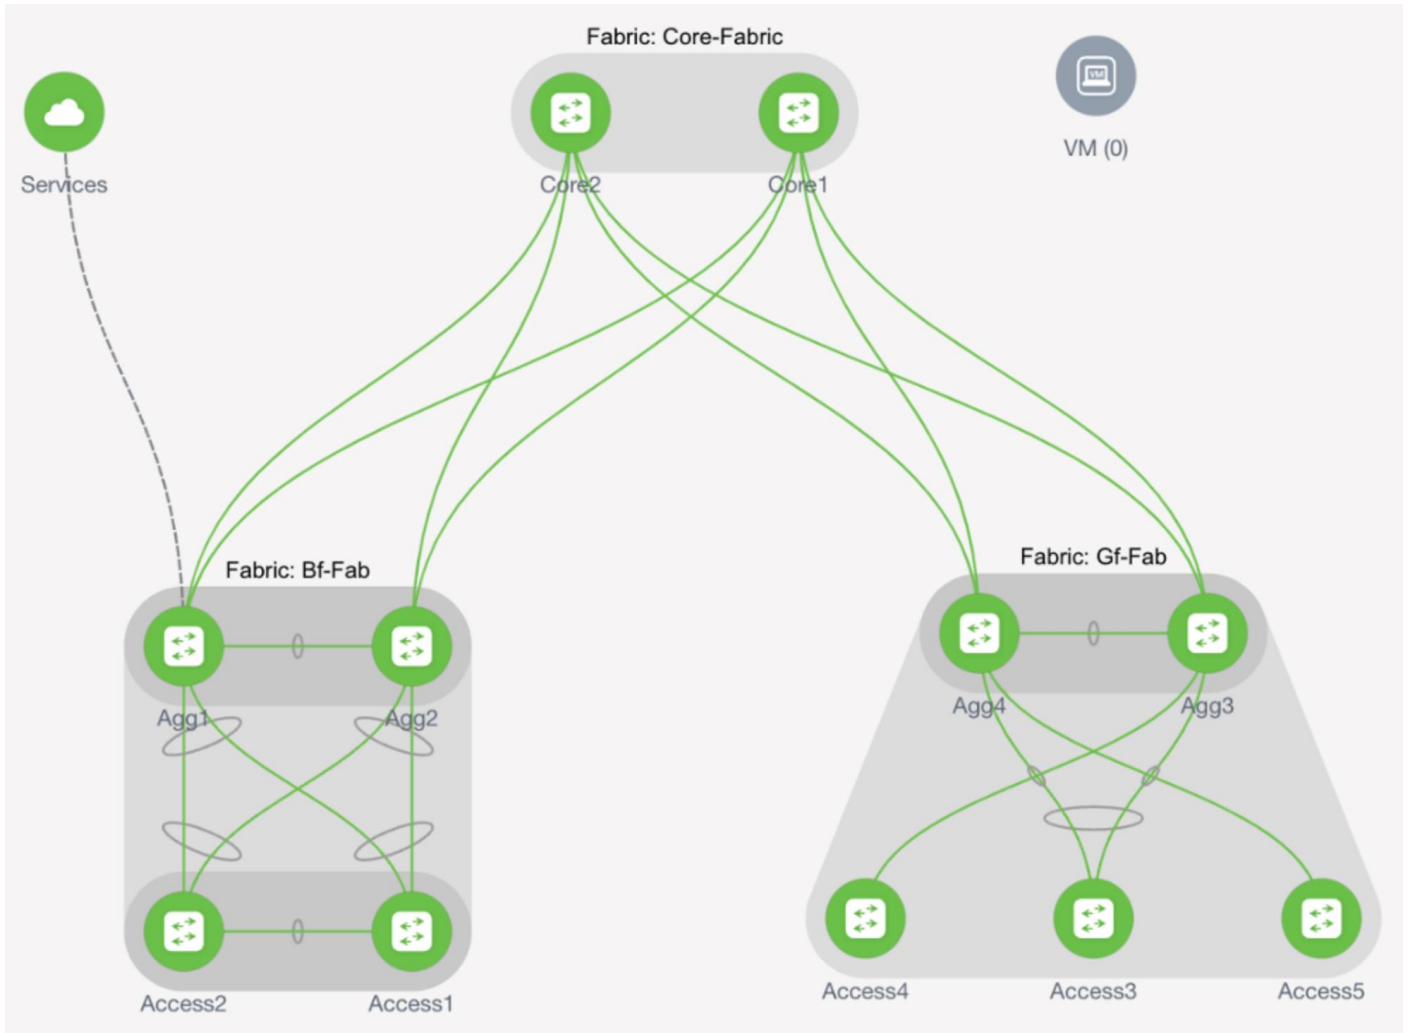


A brownfield deployment has an additional topology compared to greenfield deployments--labelled 2b-- where Access switches may be connected to the Aggregation peer through an Ethernet trunk port and not through a Layer 2 port channel.

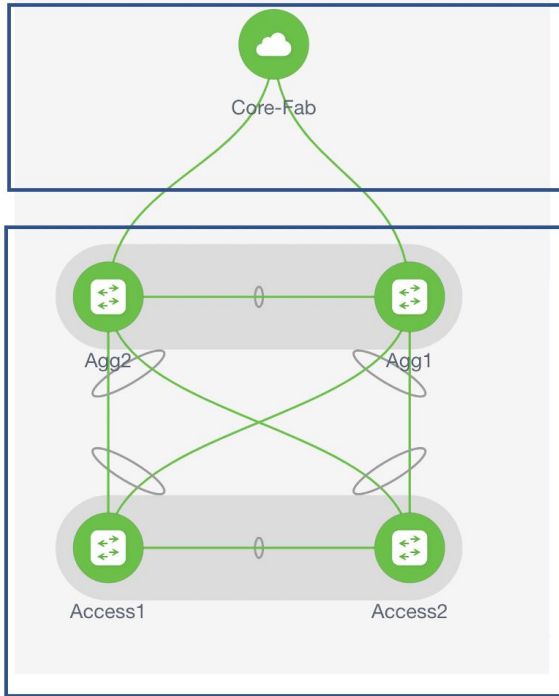
There can be several variations to the topologies based on how the switches are connected. The following figures show some of the variations:







In these topologies, Access and Aggregation switches use the "Enhanced Classic LAN" fabric type in NDFC. For Core/Edge, you must use the external connectivity network fabric type.



'External Connectivity Network' fabric with Core switch

'Enhanced Classic LAN' fabric with Access-Aggregation switches

External Connectivity from Enhanced Classic LAN

You must have external connectivity from the data center. That is, you must have reachability for workloads that are part of the data center fabric that can communicate with external network resources over WAN/backbone services. Use the VRF-Lite connectivity option between Aggregation devices and the Core/Edge router for connecting the fabric to an external Layer 3 domain for north-south traffic communication and for interconnecting several fabrics. Regarding Core and Edge, NDFC supports both roles as external connections/exit points of Enhanced Classic LAN networks. These Core and Edge platforms are placed and supported in the external connectivity network fabric type. As long as these are Nexus devices, NDFC supports the auto Virtual Routing and Forwarding-Lite (VRF-Lite) option.

Note: VRF-Lite assumes the existence of multiple VRF instances, but in classic networks that may not be the case. NDFC supports the existence of a single VRF instance as well.

In Enhanced Classic LAN, the Core or Edge roles are identical and they have no functionality differences. You can use these roles interchangeably for external connectivity.

Guidelines and Limitations for Enhanced Classic LAN

The following guidelines and limitations apply:

- Support for greenfield and brownfield.
- Supports for IPv4 and IPv6 for switch discovery as well as for VRF-Lite.
- Support for FEX, Nexus 7000, 7700, and 9000 platforms.
- You must create a vPC at the Aggregation Layer.
- Support for multiple Aggregation pairs with no overlapping VLANs as of NDFC 12.1.3..

- Support for interface groups, multi-attach/detach, quick attach/detach, and shared policies for switches.
- Support for change control and rollback.
- If Core/Edge are Nexus platforms, NDFC auto-generates VRF-Lite and routing configurations between Aggregation and Core/Edge. Manual VRF-Lite is also supported. If these are non-Nexus platforms (Cisco IOS-XE, IOS-XR, or ASR), you must manually apply templates for VRF-Lite and routing as documented in the [VRF Lite, Release 12.1.3](#) article.
- Brownfield with multiple Aggregation vPC pairs in the same Enhanced Classic LAN fabric is not supported as of release 12.1.3.
- No support for in-band management and in-band power-on auto-provisioning (POAP) for devices of the Enhanced Classic LAN fabric type.

For configuration information of the listed features, see the [Enhanced Classic LAN, Release 12.1.3](#) article.

Hardware and Software Recommendations

To use Enhanced Classic LAN, we recommend that you use Cisco Nexus Dashboard release 3.0.1 and NDFC release 12.1.3. Enhanced Classic LAN supports the Cisco Nexus 2000, 7000, 7700, and 9000 platforms.

At the Access layer, we recommend Cisco Nexus 9000 platforms. We recommend the Cisco Nexus 9000, 7000, and 7700 platforms at the Aggregation and Core layer. We do not require any specific platforms per layer.

For the models of Cisco Nexus platforms and Cisco NX-OS software supported in NDFC release 12.1.3, see the [Cisco Nexus Dashboard Fabric Controller Release Notes, Release 12.1.3](#).

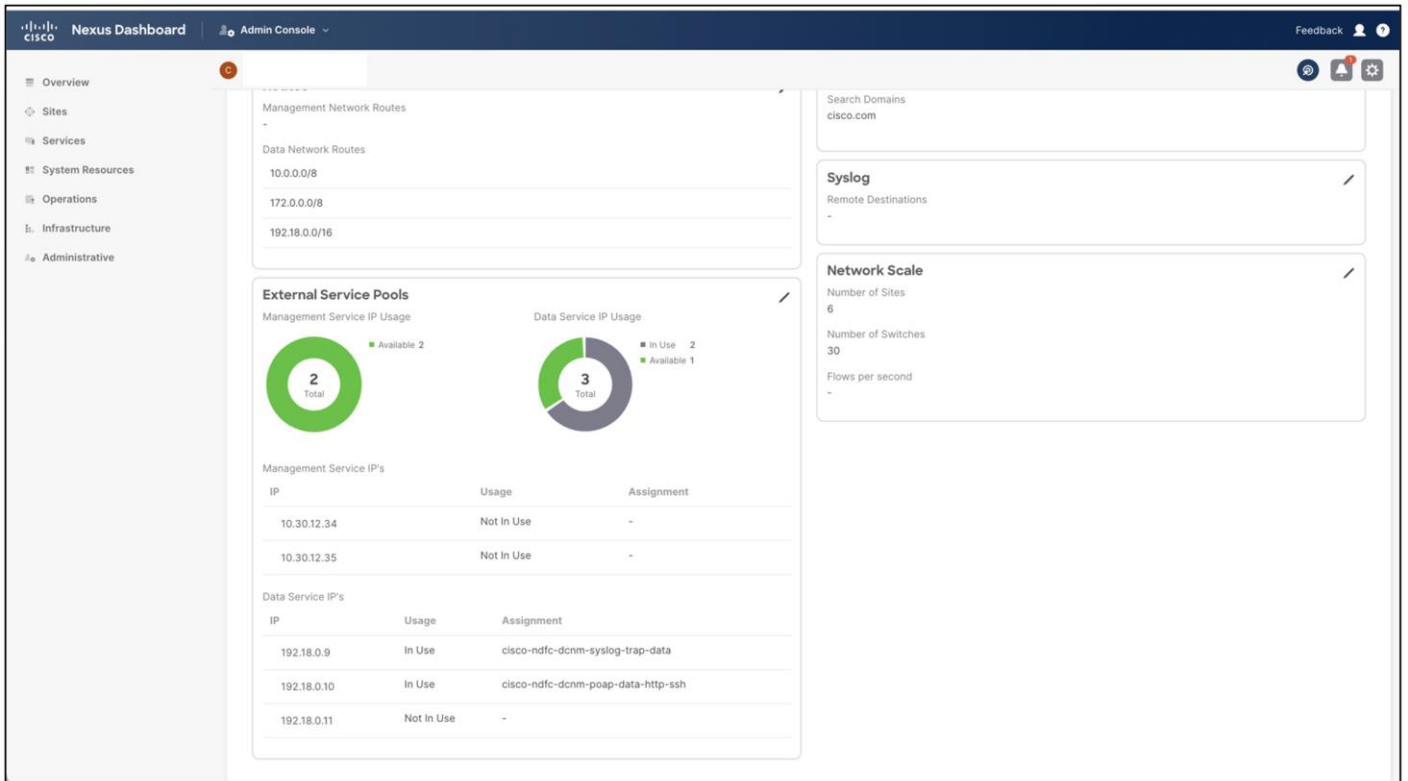
Using Enhanced Classic LAN

Prerequisites

You must meet the following requirements to start provisioning Classic Ethernet networks:

- You must have at least one Cisco Nexus Dashboard cluster and a healthy NDFC service before you can perform the other operations/setup steps.
- Cisco Nexus Dashboard (virtual or physical) nodes to form a cluster.
 - For the sizing guide for the number of nodes per form factor and supported scale, see the [Nexus Dashboard Capacity Planning](#).
 - The cluster nodes can be Layer 2 or Layer 3 adjacent on the data interface.
 - We recommend that you have a standby node for HA purposes.
 - For more information, see the [Cisco Nexus Dashboard Deployment Guide, Release 2.3.x](#).
- Cisco Nexus Dashboard Fabric Controller (NDFC).
 - You must install the service on your Cisco Nexus Dashboard cluster. For more information, see the [Cisco Nexus Dashboard Fabric Controller Installation and Upgrade Guide, Release 12.1.3](#).
 - For deployment information, see the [Cisco Nexus Dashboard Fabric Controller Deployment Guide](#).

- Reachability between the NDFC service and the switches to be managed.
 - Classic LAN only supports out-of-band (OOB) management of switches. You must decide if you want to manage your switches OOB using the Nexus Dashboard management or data interface.
 - Define appropriate routes for the reachability of the switches from the Nexus Dashboard cluster in Nexus Dashboard under **Infrastructure -> Cluster Configuration**. Define external service pools for SNMP and POAP over management or data subnet.

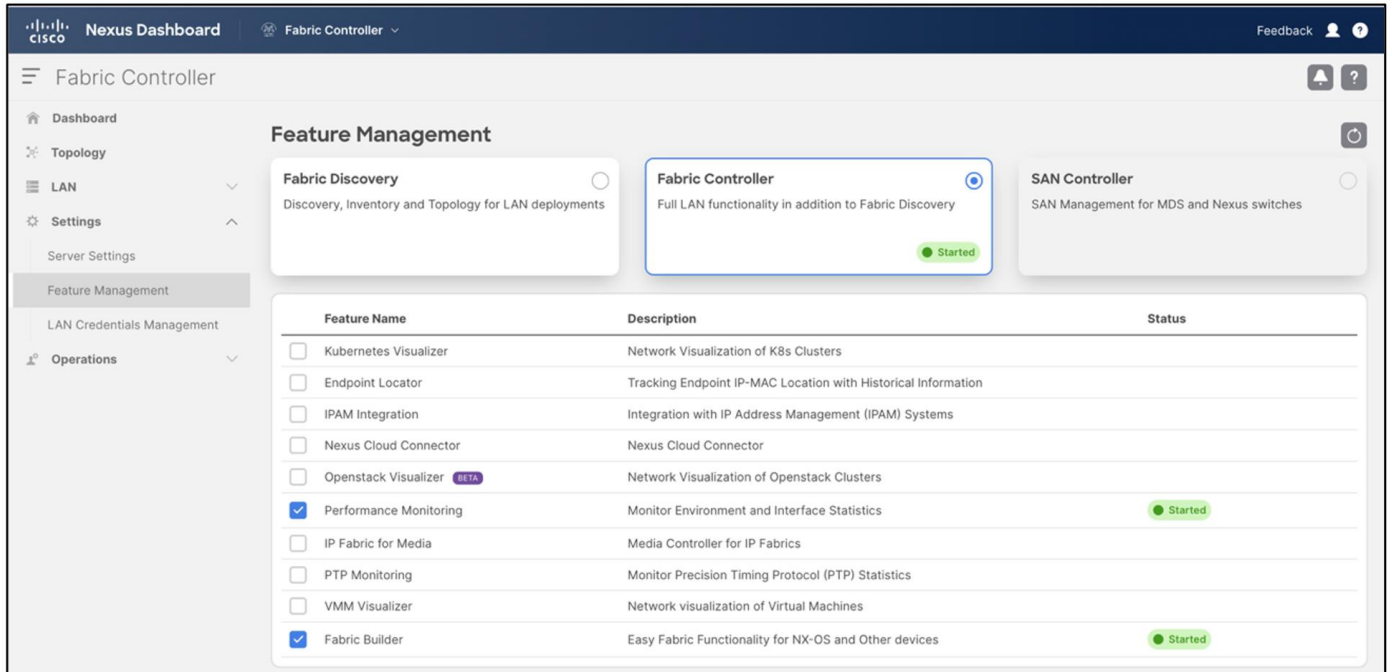


- NDFC "LAN Device Management Connectivity" in the Server Settings is set to "Management" by default. If you use the default setting, you must spawn external IP addresses associated with the Nexus Dashboard management subnet and ensure that Nexus Dashboard uses the management interface to communicate with the mgmt0 interfaces of the switches. You might need to add static routes that are associated with the Nexus Dashboard management interface. You do not need to add these routes only if the Nexus Dashboard management interface is deployed in the same subnet of the mgmt0 interfaces of the switches.
- If you change the setting to "Data" (as seen below), you must spawn external IP addresses associated to the Nexus Dashboard data subnet, and you can establish that connectivity between the Nexus Dashboard data interface and the mgmt0 interfaces of the switches. You do not need any static routes for this, as NDFC uses the default route associated with the data interface by default if that the Nexus Dashboard management interface is not in the same subnet as the mgmt0 interfaces of the switches, as discussed in the previous paragraph. You can change this setting to "Data" as well.
- To change the default option, navigate to the "LAN Device Management Connectivity" option in NDFC, which is found under **Settings -> Server Settings -> Admin**. These settings are discussed in the [Cisco Nexus Dashboard Fabric Controller Deployment Guide](#).

The screenshot displays the Cisco Nexus Dashboard Fabric Controller interface. The top navigation bar shows 'Nexus Dashboard' and 'Fabric Controller'. The left sidebar contains a menu with 'Settings' expanded to 'Server Settings'. The main panel, titled 'Server Settings', has tabs for 'Alarms', 'Events', 'Reports', 'LAN-Fabric', 'Discovery', 'SSH', 'PM', 'VMM', 'SNMP', 'Admin', 'SMTP', and 'Debug'. The 'Admin' tab is active. The configuration fields are as follows:

- LAN Device Management Connectivity*: Data
- Specify connection pool, max active connection*: 20
- Specify connection validation*:
- Specify validation query for database*: select 1
- Database performance test interval*: 20
- Database history tables maintenance interval (in days)*: 90

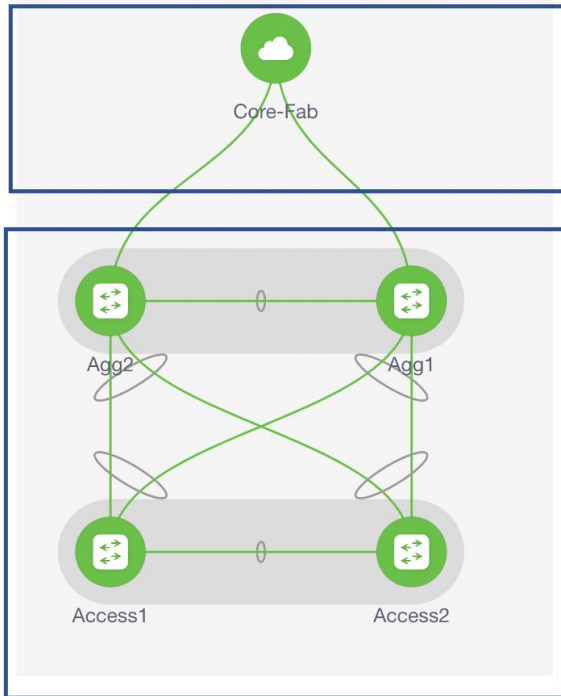
- Feature Management within NDFC
 - Fabric Builder (at the very minimum).
 - Performance Monitoring (optional) for SNMP-based performance monitoring (CPU, memory, traffic, temperature, interface, and links) at a switch level.



- Fabric Type to be used
 - Enhanced Classic LAN
- Switches and roles that can be managed in Enhanced Classic LAN Fabric
 - Cisco Nexus 2000, 7000, 7700, and 9000.
 - Enhanced Classic LAN supports only the Access and Aggregation roles. For Core/Edge, you must use the external connectivity network fabric type.

Day 0 for Classic LAN

This section discusses the day 0 for classic LAN topology. For Collapsed Core topologies, see the [Step 4: Define the Roles](#) section.



‘External Connectivity Network’ fabric with Core switch

‘Enhanced Classic LAN’ fabric with Access-Aggregation switches

This section first discusses the steps for configuring a fabric consisting of the [Access and Aggregation layers](#), then discusses the steps for creating a [fabric for the Core layer](#).

Note: All features highlighted for the Core layer are also supported for the Edge role.

There are two separate fabrics because typical deployments comprise a shared Core layer. The Core routers reside in a separate fabric shared by fabrics with the Access-Aggregation switches.

For the Access and Aggregation Layers

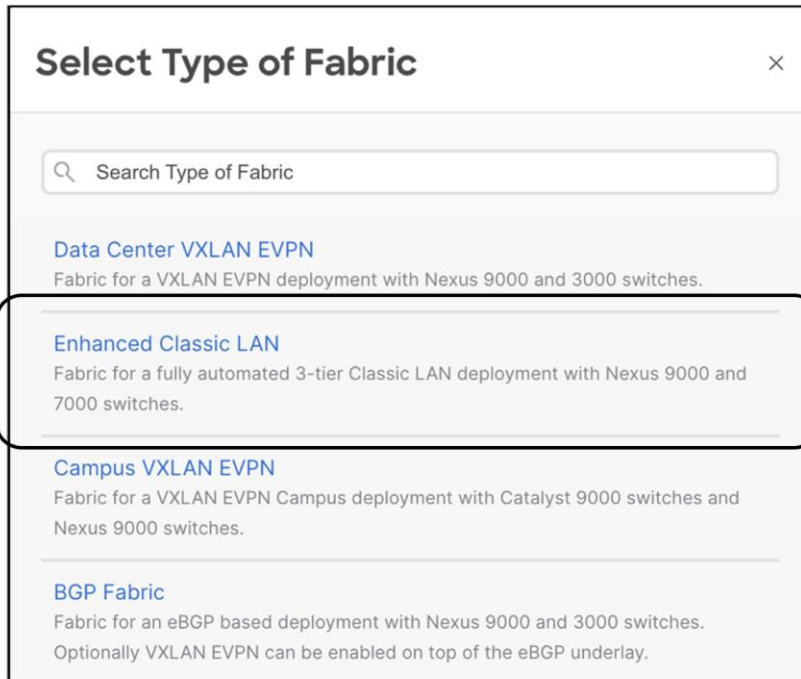
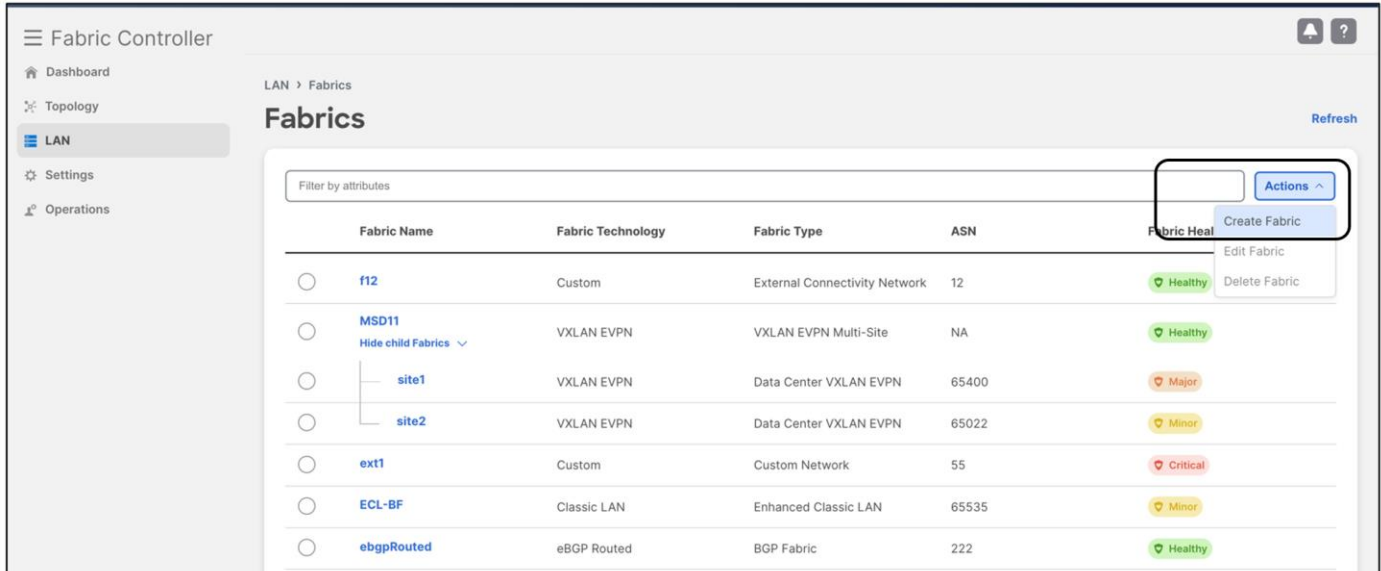
The overall process for configuring the Access and Aggregation layers is as follows:

1. [Create the fabric.](#)
2. [Discover the switches in the fabric.](#)
 - [Greenfield Import](#)
 - [Brownfield Import](#)
3. [Bootstrap the switches \(Power-On Auto-Provisioning\).](#)
4. [Define the roles.](#)
5. [Configure the vPC pairing.](#)
6. [Recalculate and deploy.](#)

The following sections discuss each of the steps in detail. The screenshots of pending configurations are examples of a single Access, Aggregation, and Core layer for explanatory purposes and do not include all the pairs that are in the topology.

Step 1: Create the Fabric

The first step is to create a fabric using the "Enhanced Classic LAN" template. This template is for Access and Aggregation switches. The fabric-level template helps define parameters that apply to the respective network layers and the fabric as a whole.



After you choose the fabric template, you can optionally customize the fabric-level settings or leave them at their default values. The default values are per Cisco best practice recommendations, and thus you should not change the values unless required.

The only mandatory fields are as follows:

- **Fabric Name** and **ASN** when you use eBGP as a peering protocol between the Aggregation and Core layers

- **Fabric Name, OSPF Tag, and Area ID** when you use OSPF as a peering protocol between the Aggregation and Core layers

The default options for **General Parameters** are as follows:

- FHRP – HSRP
 - Supported options: HSRP, VRRP, VRRP3, and None.
- Peering Protocol between Agg and Core/Edge – eBGP
 - Supported options: eBGP, OSPF, and None.
- Enable Performance Monitoring – No
 - This option collects SNMP-based stats, such as CPU, memory, and temperature.

The default options for spanning tree parameters are as follows:

- Spanning Tree Protocol – RPVST+
 - Supported options: RPVST+, MST, and Unmanaged.

All values and timers are per Cisco best practice and can be customized.

Classic-Demo

Pick Fabric

[Enhanced Classic LAN >](#)

General Parameters **Spanning Tree** VPC Protocols Advanced Resources Manageability Bootstrap Configuration Backup Flow Monitor

Spanning Tree Root Bridge Protocol

rpvst+ Protocol to be used for configuring Root Bridge: rpvst+: Rapid Per-VLAN Spanning Tree, mst: Multiple Spanning Tree, unmanaged (default): STP Root not managed by NDFC. Note: Spanning Tree Settings and Bridge Configs are applicable at Aggregation layer only.

rpvst+ mst unmanaged

Vlan Range Vlan range, Example: 1,3-5,7,9-11, Default is 1-3967 (Applicable only for Aggregation devices)

MST Instance Range MST instance range, Example: 0-3,5,7-9, Default is 0 (Applicable only for Aggregation devices)

Spanning Tree Bridge Priority Bridge priority for the spanning tree in increments of 4096 (Applicable only for Aggregation devices)

Spanning Tree Hello Interval Set the number of seconds between generation of config bpdu, default is 2 (Applicable only for Aggregation devices)

Spanning Tree Forward Delay Set the number of seconds for the forward delay timer, default is 15 (Applicable only for Aggregation devices)

Spanning Tree Max Age Interval Set the maximum number of seconds the information in a bpdu is valid, default is 20 (Applicable only for Aggregation devices)

Spanning Tree Pathcost Method long: Use 32 bit based values, short (default): Use 16 bit based values for default port path costs (Applicable only for

The vPC defaults in the following screenshot are per Cisco's best practices and you can customized the values:

Fabric Name
Classic-Demo

Pick Fabric
Enhanced Classic LAN >

General Parameters Spanning Tree **VPC** Protocols Advanced Resources Manageability Bootstrap Configuration Backup Flow Monitor

vPC Auto Recovery Time (In Seconds)
360 (Min:240, Max:3600)

vPC Delay Restore Time (In Seconds)
150 (Min:1, Max:3600)

vPC Peer Link Port Channel ID
500 (Min:1, Max:4096)

vPC IPv6 ND Synchronize
 Enable IPv6 ND synchronization between vPC peers

vPC Domain Id Range
1-1000 vPC Domain id range to use for new pairings

vPC Layer-3 Peer-Router Option
 Enable Layer-3 Peer-Router on all Aggregation Devices

You can customize the following additional settings:

- **Protocols** for OSPF and BGP authentication
- **Advanced** for AAA, NXAPI, templates to be used for sub-operations, CoPP profile, as well as group freeform configuration for Access and Aggregation switches
- **Resources** for the default IP address and subnet ranges
- **Manageability** for DNS, NTP, and syslog server settings
- **Bootstrap** for POAP and DHCP server settings; use this option to enable POAP at a fabric level
- **Configuration Backup** to define the cadence of automatic fabric level backups
- **Flow Monitor** to enable Netflow

The Recalculate & Deploy (R&D) process auto-generates the appropriate best practice configurations for both the Access and Aggregation layers.

Note: After you create a fabric, you cannot edit the values for **Fabric Name**, **First Hop Redundancy Protocol**, and **VRF Lite Agg-Core or Collapsed Core-WAN Peering Protocol Options**.

Step 2: Discover the Switches in the Fabric

After you create the fabric, you can import the switches using the **Seed IP** and **Credentials**. Make sure the reachability exists between NDFC and these switches. The seed IP address must be the management IP address of one of the switches. This fabric type supports only Out-of-Band management of the switches as of NDFC release 12.1.3.

Greenfield Import

If you do not put a check in the **Preserve Config** box, NDFC performs a greenfield import. NDFC erases all existing configurations except the management IP address, default gateway, and boot variables, and pushes a fresh configuration. You can then manage all switches from scratch.

Switch Addition Mechanism*

Discover

Seed Switch Details

Seed IP*
192.18.0.13
Ex: "2.2.2.20" or "10.10.10.40-60" or "2.2.2.20, 2.2.2.21"

Authentication Protocol*
MD5

Username*
admin

Password*
.....

Max Hops*
2

Preserve Config ← Greenfield
Unchecking this will clean up the configuration on switch(es)

In the case of a greenfield addition of Cisco Nexus 9000 switches, by default NDFC learns the basic intent from the switch and performs a write erase and reload followed by restoring only the basic intent on that switch. For Cisco Nexus 7000 switches, given that they are VDC-based, the greenfield addition follows a different path. Specifically, because NDFC does not support VDC POAP, NDFC performs the clean-up on the Nexus 7000 device without a reload. This works similarly to the greenfield clean-up on Nexus 9000 switches without the reload option. You can disable the reload for Nexus 9000 switches in the **Fabric** settings under the **Advanced** tab.

Greenfield Cleanup Option

Enable Enable to clean switch configuration without reload when PreserveConfig=no

Brownfield Import

If you put a check in the **Preserve Config** box, NDFC performs a brownfield import, which preserves all existing configurations.

Switch Addition Mechanism*

Discover

Seed Switch Details

Seed IP*

192.18.0.13

Ex: "2.2.2.20" or "10.10.10.40-60" or "2.2.2.20, 2.2.2.21"

Authentication Protocol*

MD5

Username*

admin

Password*

Max Hops*

2

Preserve Config ←Brownfield

Unchecking this will clean up the configuration on switch(es)

In Enhanced Classic LAN with brownfield import, NDFC learns preserves all configurations in the switches. Thereafter, NDFC incrementally manages the switches. The prerequisite is that the fabric of imported switches must be fully functional with configurations as per Cisco best practices. If not, the brownfield import fails and NDFC generates relevant error messages to guide the user.

When you perform a brownfield import on NDFC, adhere to Cisco's best practices and recommendations before the import. In addition, you must meet the following prerequisites:

- No support for brownfield of fabric path-based infrastructures
- Enhanced Classic LAN can support 3-tier and 2-tier with the Collapsed Core layer
- You must have vPCs at the Aggregation layer
- A Layer 2/Layer 3 boundary is supported only with the Aggregation/Collapsed Core layers

Note: For brownfield deployments, you must create an Enhanced Classic LAN Fabric and configure the fabric settings in accordance with their existing legacy 3-tier or 2-tier deployment. For example, if you use eBGP as a VRF Lite protocol between the Aggregation and Core layers, then you must choose eBGP in the settings and provide the appropriate ASN. In addition, you must set the appropriate spanning tree-related parameters in the fabric settings. You must disable NX-API if that is not required because these options are enabled in the fabric settings by default. You must set the role of the Aggregation switches, as by default all roles are set to Access (the role definition is discussed in this section). Switches are placed in "Migration Mode" if you perform a brownfield addition. After this is done, you must run a Recalculate and Deploy (R&D) for the fabric as described below.

After NDFC discovers the switches, NDFC shows a list and the user can select the appropriate switches and add them to the fabric. Depending on whether this is a greenfield or brownfield import, NDFC performs specific actions as described above.

Switch Addition Mechanism*
Discover

Seed Switch Details

Fabric	Switch	Authentication Protocol	Username
Bf-Fab	192.18.0.15-20	MD5	admin
Password	Max Hops	Preserve config	
● Set	1	● Enabled	

← Back

Discovery Results

Status == manageable X Edit Clear All

Switch Name	Serial Number	IP Address	Model	Version	Status	Progress
Agg1	9EEM1T58D89	192.18.0.16	N9K-C9300v	9.3(9)	● Manageable	
Access1	9HMPQ8NTXEV	192.18.0.13	N9K-C9300v	9.3(9)	● Manageable	
Access2	9MZSTD2N4ML	192.18.0.14	N9K-C9300v	9.3(9)	● Manageable	
Agg2	9UEXH9V6Z9L	192.18.0.17	N9K-C9300v	9.3(9)	● Manageable	

Close Add Switches

Step 3. Bootstrap (Power-on Auto-provisioning)

To bring up a new switch with the management IP address, default route, and start-up configurations, you can use power-on auto-provisioning (POAP) from NDFC. NDFC supports only Out-Of-Band POAP for switches in Enhanced Classic LAN fabric type and supports IPv4 and IPv6-related POAP options. NDFC can be the local DHCP server providing a management IP address and a default route for reachability when the switch is bootstrapped. You can also push the desired startup configurations and optionally an image with which to boot the switch. Alternatively, you can use an external DHCP server.

POAP Process

After you power up a switch, attach the cables, and add the switch to a POAP loop, the switch sends out DHCP requests on all the interfaces that are UP. Any DHCP server can respond to this request. The server providing the DHCP offer will be printed in the POAP logs in the switch. You must ensure that multiple DHCP servers are not deployed on the same segment, otherwise the POAP process may be impacted.

In this case, let's consider NDFC to be the local DHCP server that is reachable from the switch. The POAP script on the switch tries to download the startup configuration from NDFC, which is provided after the switch is bootstrapped from NDFC. The switch tries to download the configurations from NDFC, fails, and repeats the process until the switch is provisioned. In the meantime, NDFC hands out temporary management IP addresses and a default gateway to the switch (as defined in Fabric Settings). After the switch is bootstrapped, the management IP address that you provided for each switch replaces the temporary IP address.

POAP in NDFC

In this example, we are trying to bootstrap two Access switches into an existing Enhanced Classic LAN fabric with Access and Aggregation switches that were discovered using their seed IP addresses. Alternatively, you can use a fresh fabric with no switches; both options are supported.

The first step is to enable bootstrapping in the Fabric Settings and optionally enable a local DHCP server (use NDFC as a DHCP server). You must also define the subnet scope and default gateway that NDFC will use temporarily while the switch is in its POAP loop after the switch has been powered up.

Pick Fabric
Enhanced Classic LAN >

General Parameters Spanning Tree VPC Protocols Advanced Resources Manageability **Bootstrap** Configuration Backup Flow Monitor

Enable Bootstrap Automatic IP Assignment For POAP

Enable Local DHCP Server Automatic IP Assignment For POAP From Local DHCP Server

DHCP Version
DHCPv4

DHCP Scope Start Address*
192.168.92.100 Start Address For Switch POAP

DHCP Scope End Address*
192.168.92.110 End Address For Switch POAP

Switch Mgmt Default Gateway*
192.168.92.1 Default Gateway For Management VRF On The Switch

Switch Mgmt IP Subnet Prefix*
24 (Min:8, Max:30)

DHCPv4 Multi Subnet Scope ▲
 lines with # prefix are ignored here

Enable AAA Config Include AAA configs from Manageability tab during device bootup

Bootstrap Freeform Config
 Additional CLIs required during device bootup/login e.g. AAA/Radius

Under the specific fabric, **Add Switches -> Bootstrap** tab, the switches in the POAP loop are listed. At this point, NDFC hands out only temporary management IP addresses to the switches.

Switch Addition Mechanism*

Discover Bootstrap(POAP) Pre-provision

Switch Credentials

Admin password*

For discovery, use*

Admin user and supplied password Specify a new user

Switches to Bootstrap

Filter by attributes Refresh

<input type="checkbox"/>	Serial Number	Model	Version	IP Address	Hostname	Gateway	Role	Action
<input type="checkbox"/>	98Z8C5LLPTL	N9K-C9300v	10.2(4)			192.168.92.1/24		Edit
<input type="checkbox"/>	9REYJCVCSI	N9K-C9500v	10.2(4)			192.168.92.1/24		Edit
<input type="checkbox"/>	9UXZ1T3RXUW	N9K-C9300v	10.3(1)			192.168.92.1/24		Edit

To bootstrap the switches and send the startup configuration down from NDFC to the switch, you must enter an Admin password.

Note: The primary use case for **Specify a new user** is AAA.

The AAA configurations must be part of the **Manageability** tab or Access/Aggregation freeform under Fabric Settings. In the **Bootstrap** tab thereafter, you must put a check in the **Enable AAA Config** box. That way, all configurations provided are used during bootstrap.

During the bootstrap process, the specified discovery user must be a valid AAA user. NDFC uses this user for switch discovery.

You must edit the properties in the **Edit bootstrap switch** dialog for each switch. That is, you must edit the mgmt0 IP address (this will be the permanent management IP address), hostname, switch role (in this case, Access or Aggregation), and, optionally, an image policy with which to boot the switch. The image policy as well as the image must be present in NDFC prior to choosing the image policy. For more information, see "Image Management" in the "Operations" chapter of the [Cisco NDFC-Fabric Controller Configuration Guide](#). The **Data** field is automatically populated.

Edit bootstrap switch

IP Address*

192.168.92.201

Hostname*

Access3

Image Policy

Select...

Switch Role

Access

Data*

```
{
  "modulesModel": [
    "N9K-X9364v",
    "N9K-vSUP"
  ],
  "gateway": "192.168.92.1/24"
}
```

After you enter all the details and you click **Import Selected Switches**, the switches receive the respective startup configuration from NDFC and NDFC replaces the temporary mgmt0 IP address with the address that you entered in this step.

Add Switches - Fabric: Demo1

Switch Addition Mechanism*

Discover Bootstrap(POAP) Pre-provision

Switch Credentials

Admin password*

For discovery, use*

Admin user and supplied password Specify a new user

Switches to Bootstrap

Filter by attributes Refresh

<input checked="" type="checkbox"/>	Serial Number	Model	Version	IP Address	Hostname	Gateway	Role	Action
<input checked="" type="checkbox"/>	98Z8C5LLPTL	N9K-C9300v	10.2(4)	192.168.92.201	Access3	192.168.92.1/24	access	Edit
<input checked="" type="checkbox"/>	9REYJCVCSI	N9K-C9500v	10.2(4)	192.168.92.202	Access4	192.168.92.1/24	access	Edit

Close Import Selected Switches

Now, you see the two switches bootstrapped under the **Switches** tab with **Config Status** as "NA," but the roles are defined.

Overview Switches Links Interfaces Policies Networks VRFs Event Analytics History Resources									
Filter by attributes									Actions
<input type="checkbox"/>	Switch	IP Address	Role	Serial Number	Config Status	Oper Status	Discovery Status	Model	VPC Role
<input type="checkbox"/>	Access1	10.30.12.16	Access	9L73CUD6EB9	In-Sync	Healthy	Ok	N9K-C9300v	Primary
<input type="checkbox"/>	Access2	10.30.12.17	Access	9NUXUAWS5IO	In-Sync	Healthy	Ok	N9K-C9300v	Secondary
<input checked="" type="checkbox"/>	Access3	192.168.92.201	Access	98Z8C5LLPTL	NA	Healthy	Unreachable	N9K-C9300v	
<input checked="" type="checkbox"/>	Access4	192.168.92.202	Access	9REYJCVCVSI	NA	Healthy	Unreachable	N9K-C9500v	
<input type="checkbox"/>	Agg1	10.30.12.15	Aggregation	92PV30FWOHR	Out-Of-Sync	Healthy	Ok	N9K-C9300v	Primary
<input type="checkbox"/>	Agg2	10.30.12.18	Aggregation	922L6WNL45G	Out-Of-Sync	Healthy	Ok	N9K-C9300v	Secondary

The startup configuration that NDFC pushed to the switches is as follows (for Access3):

```

ipv6 switch-packets Ila
power redundancy-mode ps-redundant
no password strength-check
hostname Access3
username admin password xyz role network-admin
vrf context management
 ip route 0.0.0.0/0 a.b.c.d
interface mgmt0
 vrf member management
interface mgmt0
 no shutdown
 no cdp enable
 ip address a.b.c.e/24

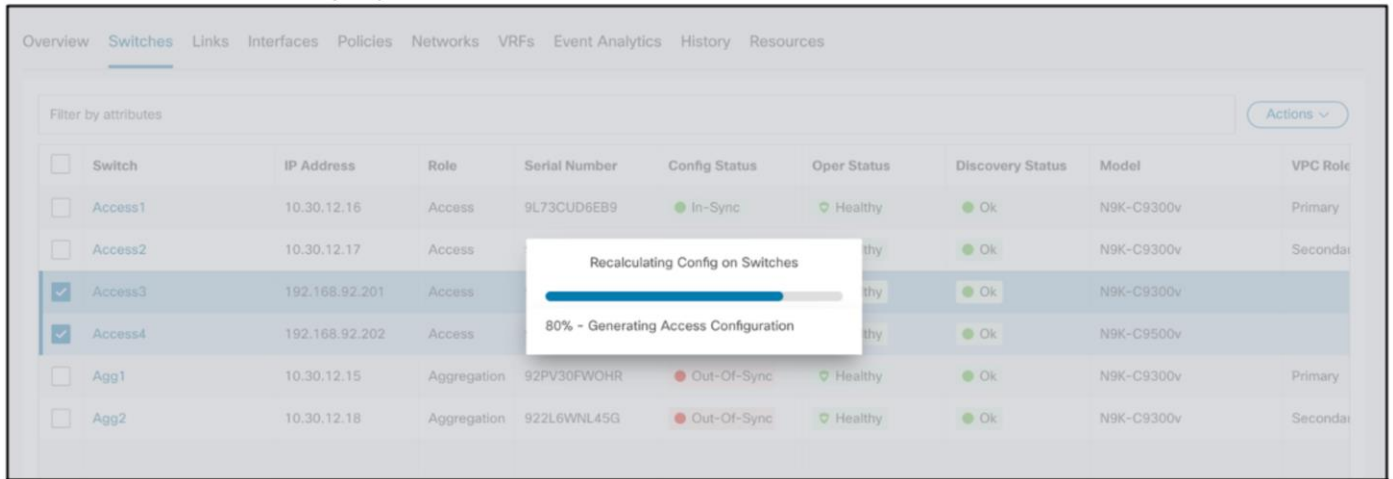

```

To inherit the fabric settings with respect to routing, spanning tree, FHRP, and so on, perform a "Recalculate and Deploy" :

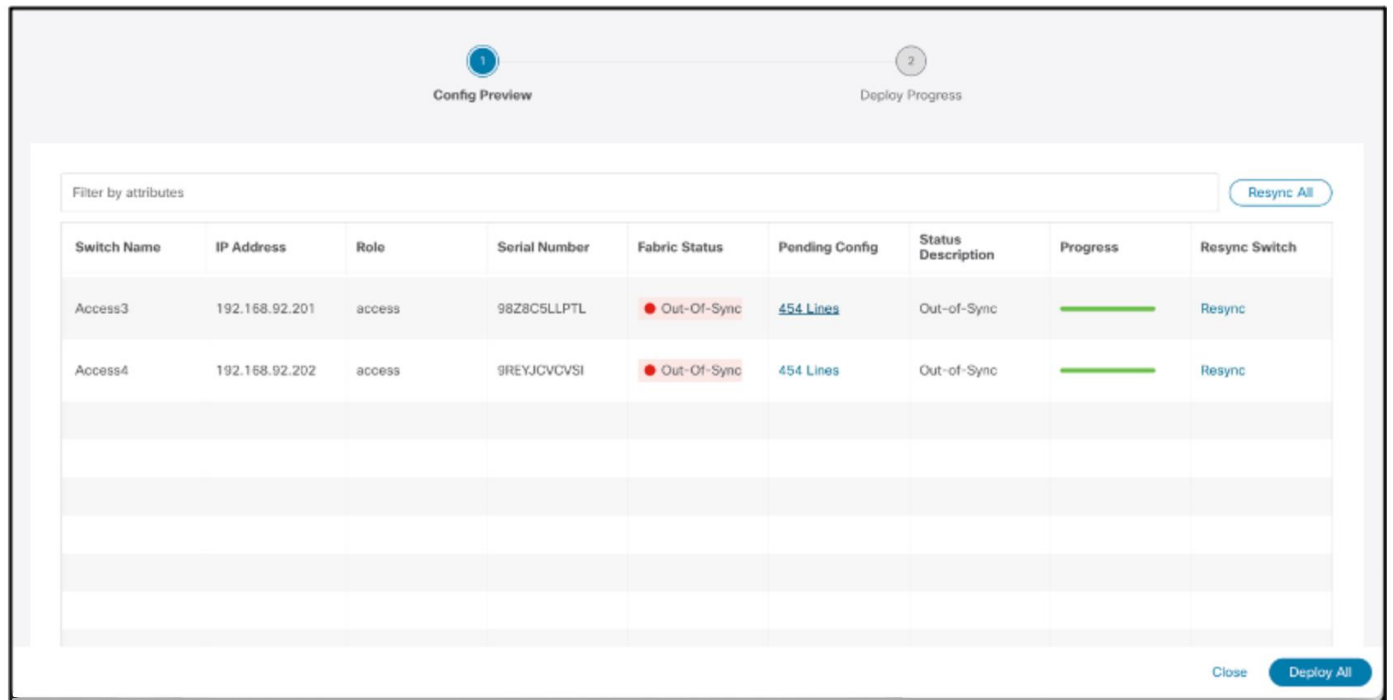
1. Select the switches.

Overview Switches Links Interfaces Policies Networks VRFs Event Analytics History Resources									
Filter by attributes									Actions
<input type="checkbox"/>	Switch	IP Address	Role	Serial Number	Config Status	Oper Status	Discovery Status	Model	VPC Role
<input type="checkbox"/>	Access1	10.30.12.16	Access	9L73CUD6EB9	In-Sync	Healthy	Ok	N9K-C9300v	Primary
<input type="checkbox"/>	Access2	10.30.12.17	Access	9NUXUAWS5IO	In-Sync	Healthy	Ok	N9K-C9300v	Secondary
<input checked="" type="checkbox"/>	Access3	192.168.92.201	Access	98Z8C5LLPTL	NA	Healthy	Ok	N9K-C9300v	
<input checked="" type="checkbox"/>	Access4	192.168.92.202	Access	9REYJCVCVSI	NA	Healthy	Ok	N9K-C9500v	
<input type="checkbox"/>	Agg1	10.30.12.15	Aggregation	92PV30FWOHR	Out-Of-Sync	Healthy	Ok	N9K-C9300v	Primary
<input type="checkbox"/>	Agg2	10.30.12.18	Aggregation	922L6WNL45G	Out-Of-Sync	Healthy	Ok	N9K-C9300v	Secondary

2. From the Fabric Overview page, choose **Actions > Recalculate and Deploy**.



3. Review the configuration preview.



4. If everything looks accurate, click **Deploy All** to deploy the configuration for both switches.

Switch Name	IP Address	Status	Status Description	Progress
Access3	192.168.92.201	STARTED	Deployment in progress.	Executed 309 / 454
Access4	192.168.92.202	STARTED	Deployment in progress.	Executed 168 / 454

After deploying the configuration, the following screenshot shows all of the switches back "In-Sync" :

Switch	IP Address	Role	Serial Number	Config Status	Oper Status	Discovery Status	Model	VPC Role
Access1	10.30.12.16	Access	9L73CUD6EB9	In-Sync	Healthy	Ok	N9K-C9300v	Primary
Access2	10.30.12.17	Access	9NUXUAWS5IO	In-Sync	Healthy	Ok	N9K-C9300v	Secondary
Access3	192.168.92.201	Access	98Z8C5LLPTL	In-Sync	Healthy	Ok	N9K-C9300v	
Access4	192.168.92.202	Access	9REYJCVCSI	In-Sync	Healthy	Ok	N9K-C9500v	

Step 4: Define the Roles

This section onward discusses switches imported using seed IP addresses and not POAP.

After you import the switches, you can begin defining your intent, as in what do we want this switch to be: an Access or an Aggregation? Based on this role, appropriate configuration is generated and pushed to the switches by NDFC.

Enhanced Classic LAN has two roles: Access and Aggregation. If the user has a Core layer, you must create a separate external connectivity network fabric to place this Core switch, which is discussed in the next section.

Let's revisit the two topologies to define the roles appropriately:

1. 3-tier hierarchical network with a Layer 2/ Layer 3 boundary at the Aggregation and Core layers connecting to the WAN
2. Collapsed Core where the Core and Aggregation layers are collocated on the same switch

For #1, you can use the Access layer to connect to servers, the Aggregation layer as Layer 2/Layer 3 demarcation, and the Core layer in a separate shared fabric. The aggregation will also act as the spanning tree bridge and, optionally, a gateway with the relevant FHRP configurations.

For #2, because the Core and Aggregation layers are unified, you can use the Aggregation role as a collapsed Core layer. While serving as an Layer 2/Layer 3 demarcation, a bridge, and a gateway, this switch will also connect to the WAN (optionally using VRF-Lite, which is fully supported in the Aggregation layer). Day 1 aspects remain identical for a Collapsed Core topology, as discussed in the following

sections. However, VRF-Lite will be between the Aggregation and WAN device instead of the Aggregation and Core layers. For more information, see the [Day1 for Classic LAN](#) section.

The following screenshots show how you can select the roles, with the default role being Access:

Switch	IP Address	Role	Serial Number	Mode	Config Status	Oper Status	Discovery Status	Model	VPC Role
<input type="checkbox"/> Access1	192.18.0.13	Access	9HMPQ6NTXEV	Normal	NA	Healthy	Ok	N9K-C9300v	Secondary
<input type="checkbox"/> Access2	192.18.0.14	Access	9MZSTD2N4ML	Normal	NA	Healthy	Ok	N9K-C9300v	Primary
<input checked="" type="checkbox"/> Agg1	192.18.0.16	Access	9EEM1T58D89	Normal	NA	Healthy	Ok	N9K-C9300v	Primary
<input checked="" type="checkbox"/> Agg2	192.18.0.17	Access	9UEXH9V6Z9L	Normal	NA	Healthy	Ok	N9K-C9300v	Secondary

Select Role

- Access (current)**

- Aggregation

Step 5: Configure the vPC pairing

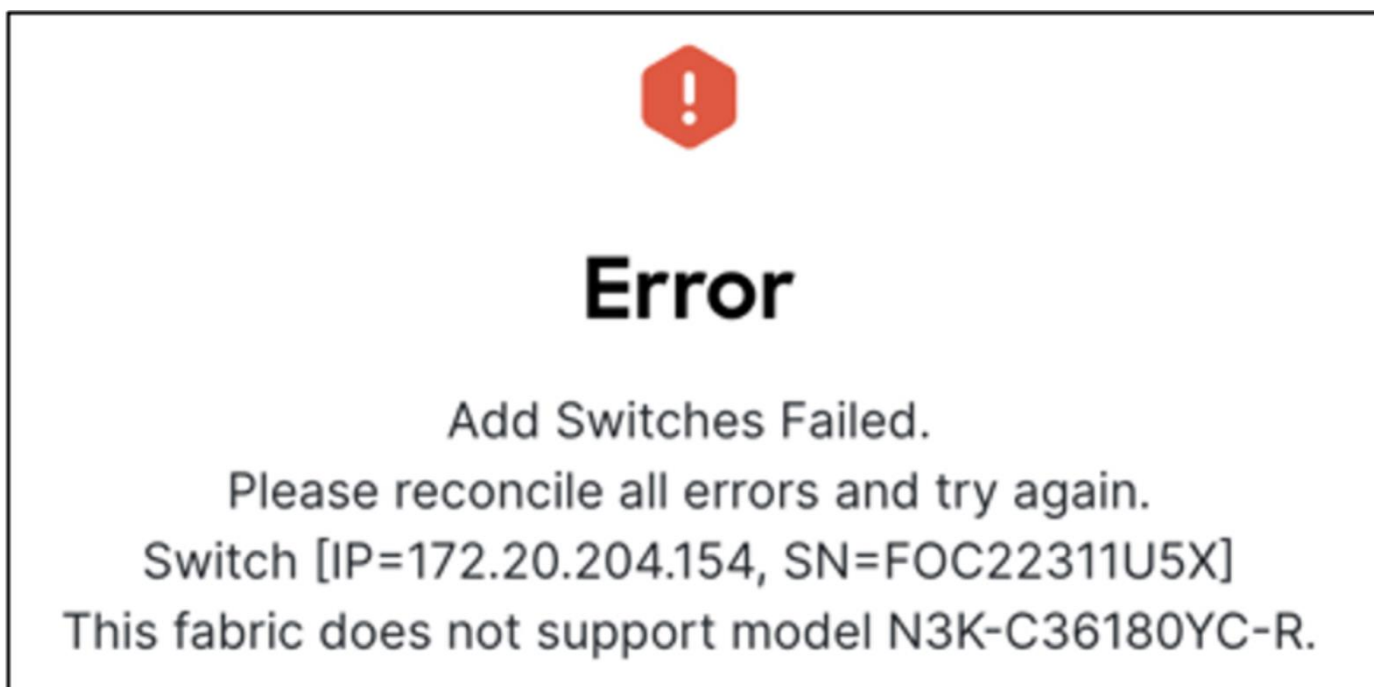
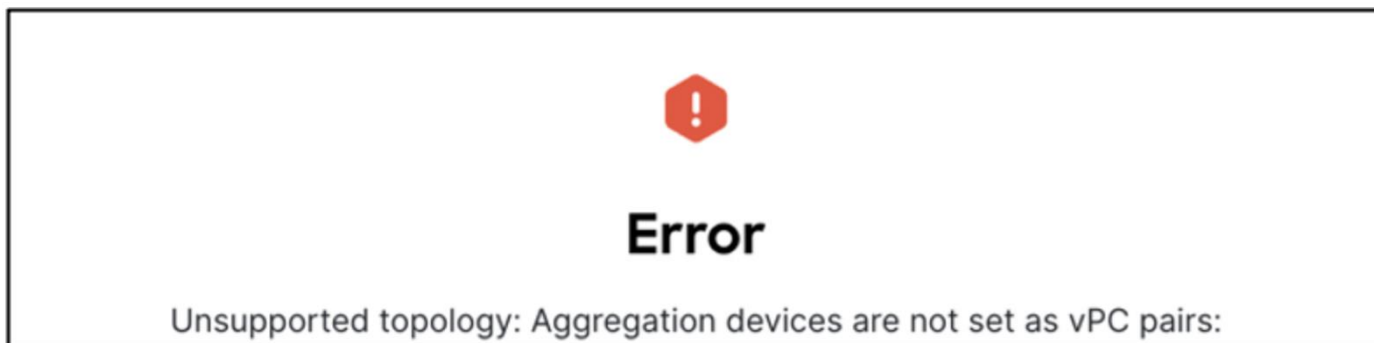
After you defined the roles, you can pair the vPCs at the Access/Aggregation layer. vPCs at the Aggregation layer is mandatory in Enhanced Classic LAN, we it recommend as per Cisco best practices. A knob is placed in the **Advanced** tab of the fabric settings to automatically detect and pair Access-Aggregation for optimal traffic Engineering. By default, the auto Aggregation-Access pairing option is enabled. This means that on a Recalculate & Deploy (R&D), NDFC automatically detects the connectivity between the Access and Aggregation switches and generate appropriate configurations based on the detected supported topologies. The configurations include vPC domains that NDFC automatically pushes to both Access and Aggregation pairs. The links between these tiers are bundled into a common vPC logical construct.

Enable Agg/Access Auto Pairing



Automatically pair aggregation and access devices based on topology

In cases where you have wired the Access/Aggregation layer such that it does not fit within the supported topologies or platforms, NDFC returns an appropriate error. The following screenshots show a few examples of the errors:



For vPC pairing at the Aggregation and the Access layer, the default option is to use the mgmt0 interface of the switches as the vPC Peer Keep Alive (PKA) link. However, if you configure a dedicated Layer 3 link for the vPC PKA, that will be honored by NDFC. You must configure this before you perform the Recalculate & Deploy step.

Access vPC Pairing

After you choose the Access switch for vPC pairing, NDFC shows the recommended devices. There is a back-to-back vPC between Access and Aggregation, which is auto detected, as shown in the following screenshot:

Overview **Switches** Links Interfaces Interface Groups Policies Networks VRFs Event Analytics History Resources Virtual Infrastructure Metrics

Filter by attributes Actions ^

Switch	IP Address	Role	Serial Number	Mode	Config Status	Oper Status	Discovery Status	Model	VPC Role
<input checked="" type="checkbox"/> Access1	192.18.0.13	Access	9HMPQ6NTXEV	Migration	NA	Healthy	Ok	N9K-C9300v	Secondary
<input type="checkbox"/> Access2	192.18.0.14	Access	9MZSTD2N4ML	Migration	NA	Healthy	Ok	N9K-C9300v	Primary
<input type="checkbox"/> Agg1	192.18.0.16	Aggregation	9EEM1T58D89	Migration	NA	Healthy	Ok	N9K-C9300v	Primary
<input type="checkbox"/> Agg2	192.18.0.17	Aggregation	9UEXH9V6Z9L	Migration	NA	Healthy	Ok	N9K-C9300v	Secondary

- Add Switches
- Preview
- Deploy
- Discovery >
- Set Role
- vPC Pairing**
- ToR/Access Pairing
- vPC Overview
- More >

vPC Pairing ? - X

Select vPC Peer for Access1

Filter by attributes

Device	Recommended	Reason	Serial Number	IP Address
<input checked="" type="radio"/> Access2	True	Switches are connected and have same role	9MZSTD2N4ML	192.18.0.14
<input type="radio"/> Agg2	False	Switches have different roles	9UEXH9V6Z9L	192.18.0.17
<input type="radio"/> Agg1	False	Switches have different roles	9EEM1T58D89	192.18.0.16

10 Rows Page 1 of 1 << 1-3 of 3 >>

Cancel Save

You do not explicitly need to vPC pair Access and Aggregation switches if you kept the Auto Pairing flag at its default value. Instead, NDFC automatically pairs the switches after this step.

Aggregation vPC Pairing

You must vPC pair Aggregation switches.

Overview **Switches** Links Interfaces Interface Groups Policies Networks VRFs Event Analytics History Resources Virtual Infrastructure Metrics

Filter by attributes Actions ^

Switch	IP Address	Role	Serial Number	Mode	Config Status	Oper Status	Discovery Status	Model	VPC Role
<input type="checkbox"/> Access1	192.18.0.13	Access	9HMPQ6NTXEY	Migration	NA	Healthy	Ok	N9K-C9300v	Secondary
<input type="checkbox"/> Access2	192.18.0.14	Access	9MZSTD2N4ML	Migration	NA	Healthy	Ok	N9K-C9300v	Primary
<input checked="" type="checkbox"/> Agg1	192.18.0.16	Aggregation	9EEM1T58D89	Migration	NA	Healthy	Ok	N9K-C9300v	Primary
<input type="checkbox"/> Agg2	192.18.0.17	Aggregation	9UEXH9V6Z9L	Migration	NA	Healthy	Ok	N9K-C9300v	Secondary

- Add Switches
- Preview
- Deploy
- Discovery >
- Set Role
- vPC Pairing**
- ToR/Access Pairing
- vPC Overview
- More >

vPC Pairing ? - X

Select vPC Peer for Agg1

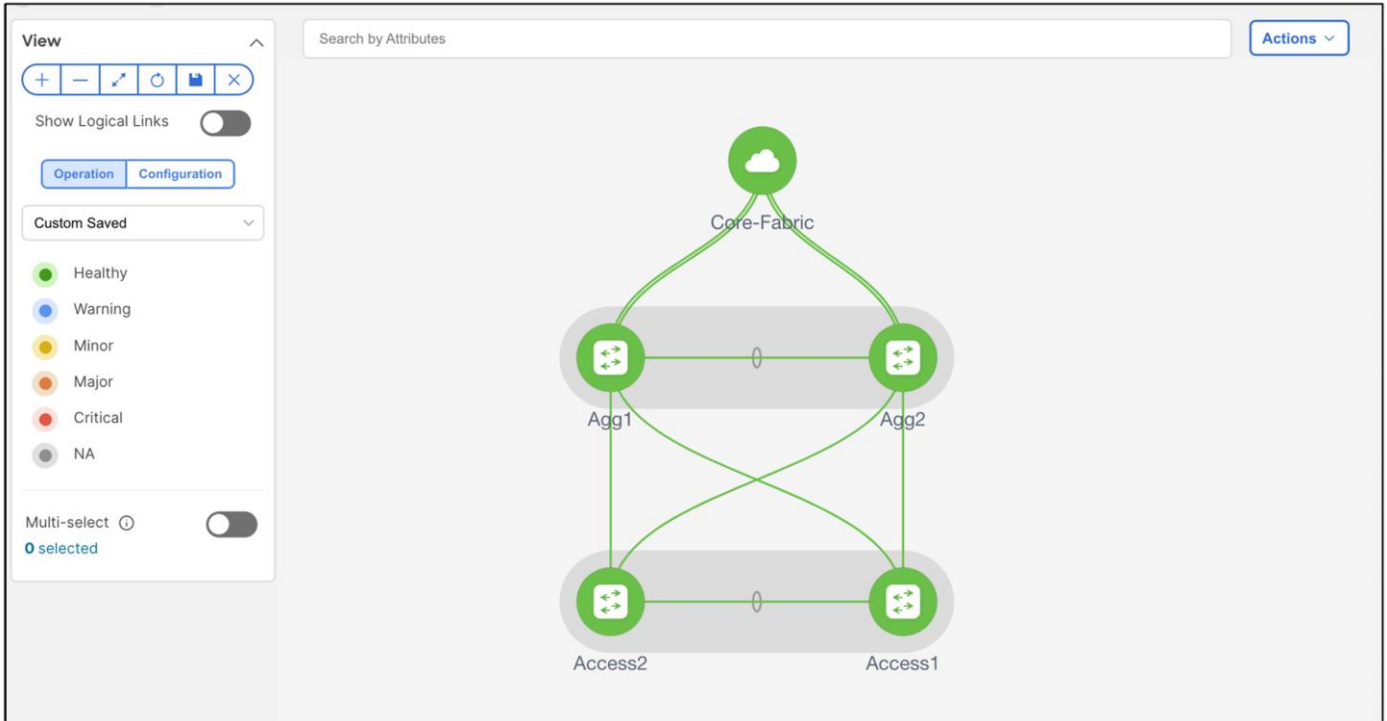
Filter by attributes

Device	Recommended	Reason	Serial Number	IP Address
<input checked="" type="radio"/> Agg2	True	Switches are connected and have same role	9UEXH9V6Z9L	192.18.0.17
<input type="radio"/> Access1	False	Already paired with 9MZSTD2N4ML (Access2)	9HMPQ6NTXEY	192.18.0.13
<input type="radio"/> Access2	T False	Already paired with 9HMPQ6NTXEY (Access1)	9MZSTD2N4ML	192.18.0.14

10 Rows Page 1 of 1 << 1-3 >>

Cancel Save

To visualize the pairing, you can navigate to the NDFC **Topology** page. As shown in the following screenshot, Access1 is paired with Access2, and Aggregation1 is paired with Aggregation2 based on user intent:



There is a back-to-back vPC present between Access and Aggregation, for which you do not need to take any explicit action to pair.

You will generate and push the vPC domain and back-to-back vPC configurations in the next step (On Recalculate & Deploy).

Step 6: Recalculate and Deploy

After you define the intent concerning the fabric, roles, and vPC, NDFC needs you to perform a "Recalculate and Deploy" (R&D). This means NDFC starts calculating the configurations required for each switch in the fabric. When doing so, it considers fabric as well as switch intent and shows you a preview of the configuration, which, after you approve, can be deployed.

In the case of a brownfield import, when you perform R&D, as part of the process NDFC performs various pre-checks on the switches comprising of the following things:

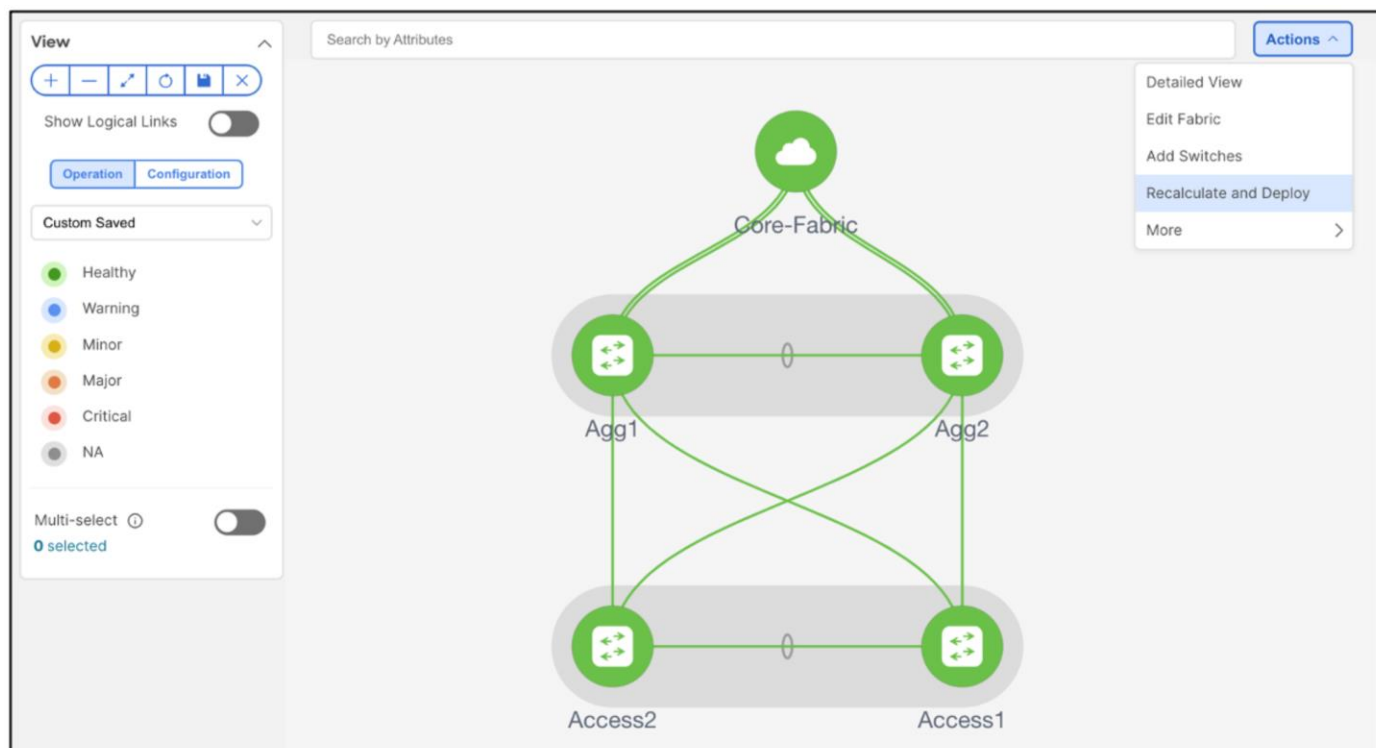
1. You must configure the Aggregation switches as a vPC pair, otherwise NDFC returns an error.
2. vPC consistency checks should indicate CONSISTENT on the vPC pairs. vPC pairs are mandatory at the Aggregation layer, but optional at the Access layer. If configured on the access layer, the vPC pair should be consistent.
3. NDFC performs various topology checks to ensure that the current deployment being imported into the ECL fabric has the right connectivity in terms of fitting it into the supported topologies. If NDFC discovers any other topology, NDFC displays an appropriate error.
4. Appropriate FHRP protocol configured in fabric settings must match what is configured on the Aggregation switches.

Note: On a successful brownfield import, NDFC learns the existing state and configurations (NDFC can now incrementally manage these things).

All vPC pairing-related information including the vPC domain, the vPC peer keepalive (KPA), and the vPC peer link are learned for the Aggregation and Access layer switches (if applicable). All interface-related configurations are learned, including access, trunk, routed, subinterface, port channels, and vPCs. The port channels or vPCs connected between the Aggregation and Access layers will be appropriately mapped to the "uplink_access" policy, along with the mapping of which Access switches map to which Aggregation switches. In addition to the network/VRF attachments, VRF-Lite related configurations are also learned. The NDFC Resource Manager will have appropriate accounting of various resources used on the switches, including but not limited to: VLANs, port channel IDs, vPC IDs, and loopback IDs.

The following procedure performs R&D:

1. Choose **Recalculate and Deploy**.



2. Preview the configuration. The configuration to be provisioned will be more substantial for a greenfield import compared to a brownfield import.

1
Config Preview
2
Deploy Progress

Filter by attributes
Resync All

Switch Name	IP Address	Role	Serial Number	Fabric Status	Pending Config	Status Description	Progress	Resync Switch
Access1	10.30.12.16	access	9L73CUD6EB9	● Out-Of-Sync	463 Lines	Out-of-Sync	<div style="width: 100%; height: 5px; background-color: green;"></div>	Resync
Access2	10.30.12.17	access	9NUXUAWSSIO	● Out-Of-Sync	463 Lines	Out-of-Sync	<div style="width: 100%; height: 5px; background-color: green;"></div>	Resync
Agg1	10.30.12.15	access	92PV30FWOHR	● Out-Of-Sync	463 Lines	Out-of-Sync	<div style="width: 100%; height: 5px; background-color: green;"></div>	Resync
Agg2	10.30.12.18	access	922L6WNL45G	● Out-Of-Sync	463 Lines	Out-of-Sync	<div style="width: 100%; height: 5px; background-color: green;"></div>	Resync

Close Deploy All

Example of Access Configuration:

Pending Config
Side-by-Side Comparison

```

feature lACP
feature lldp
feature vpc
snmp-server host 192.93.0.174 traps version 2c public udp-port 2162
spanning-tree mode mst
vpc domain 1
  peer-keepalive destination 10.30.12.17 source 10.30.12.16 hold-timeout 3
  peer-switch
  auto-recovery reload-delay 360
interface port-channel500
  switchport
  switchport mode trunk
  description "vpc-peer-link Access1--Access2"
  no shutdown
  spanning-tree port type network
  switchport trunk allowed vlan 1-4094
  vpc peer-link
interface ethernet1/5
  description "P0 500 (vpc-peer-link) member Access1-Ethernet1/5 to Access2-Ethernet1/5"
  channel-group 500 force mode active
  no shutdown
interface ethernet1/1
  switchport
  switchport mode trunk
  mtu 9216
  spanning-tree port type edge trunk

```

Example of Aggregation Configuration:

```

Pending Config Side-by-Side Comparison
feature lldp
feature vpc
snmp-server host 192.93.0.174 traps version 2c public udp-port 2162
spanning-tree mode mst
vpc domain 2
  peer-keepalive destination 10.30.12.15 source 10.30.12.18 hold-timeout 3
  peer-switch
  auto-recovery reload-delay 360
interface port-channel500
  switchport
  switchport mode trunk
  description "vpc-peer-link Agg2--Agg1"
  no shutdown
  spanning-tree port type network
  switchport trunk allowed vlan 1-4094
  vpc peer-link
interface ethernet1/4
  description "P0 500 (vpc-peer-link) member Agg2-Ethernet1/4 to Agg1-Ethernet1/4"
  channel-group 500 force mode active
  no shutdown
interface ethernet1/1
  switchport
  switchport mode trunk
  mtu 9216
  spanning-tree port type edge trunk
  no shutdown

```

3. Deploy and make sure the **Config Status** is "In-Sync".

Overview **Switches** Links Interfaces Interface Groups Policies Networks VRFs Event Analytics History Resources Virtual Infrastructure Metrics

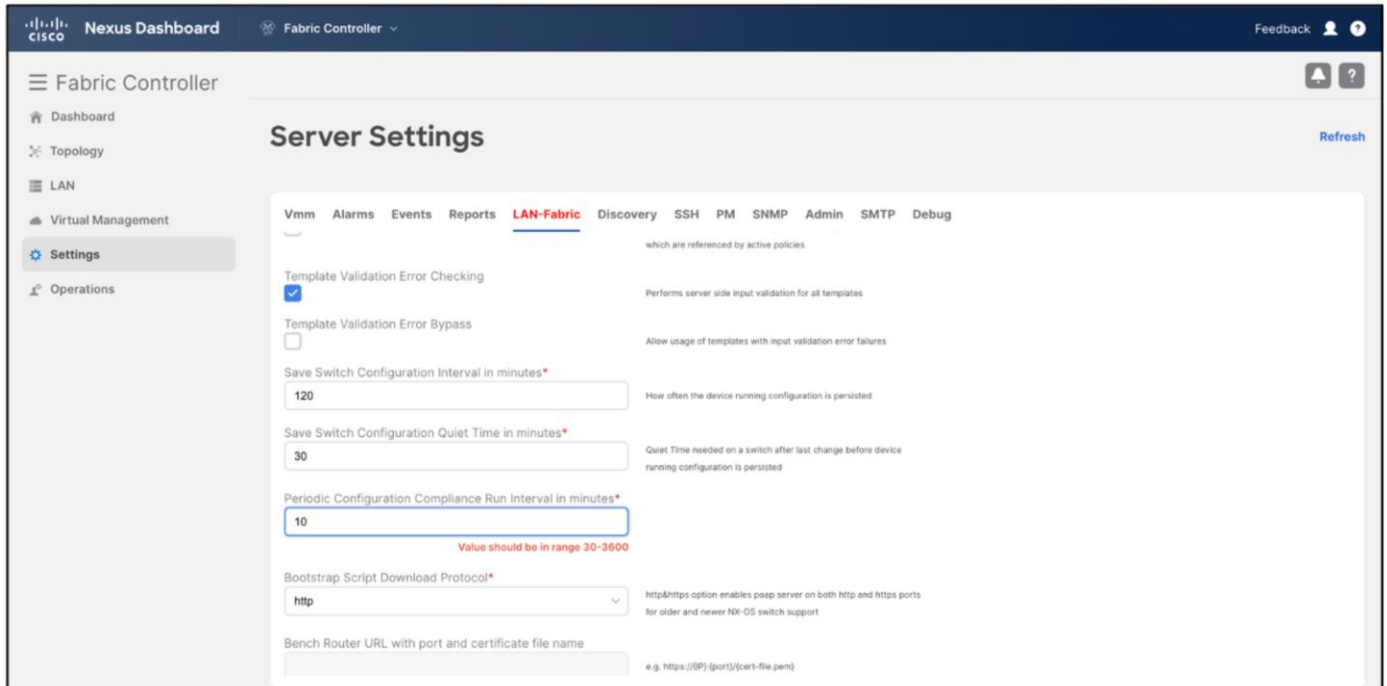
Filter by attributes Actions

<input type="checkbox"/>	Switch	IP Address	Role	Serial Number	Mode	Config Status	Oper Status	Discovery Status	Model	VPC Role	VPC Peer
<input type="checkbox"/>	Access1	192.18.0.13	Access	9HMPQ6NTXEV	Normal	In-Sync	Healthy	Ok	N9K-C9300v	Secondary	Access2
<input type="checkbox"/>	Access2	192.18.0.14	Access	9MZSTD2N4ML	Normal	In-Sync	Healthy	Ok	N9K-C9300v	Primary	Access1
<input type="checkbox"/>	Agg1	192.18.0.16	Aggregation	9EEM1T58DB9	Normal	In-Sync	Healthy	Ok	N9K-C9300v	Primary	Agg2
<input type="checkbox"/>	Agg2	192.18.0.17	Aggregation	9UEXH9V6Z9L	Normal	In-Sync	Healthy	Ok	N9K-C9300v	Secondary	Agg1

Hereafter, Configuration Compliance (CC) kicks in. Any deviation from what is intended by NDFC is flagged and the switch is marked as being "Out-of-Sync." This can be either a pending change that has not been pushed from NDFC, or an out-of-band change made using the CLI (for example). To bring it back to "In-Sync," you must deploy any pending changes at a the switch level.

For an overview of the configuration compliance, see the [NDFC Configuration Compliance](#) video on Cisco's YouTube channel.

The default setting is for CC to run once every day, but you can customize CC to be run every 30 minutes up to 3,600 minutes.



For the Core Layer

Because the Core routers are not part of the Enhanced Classic LAN fabric, you import them into a fabric type called "external connectivity network." For more information, see the [Cisco NDFC-Fabric Controller Configuration Guide, Release 12.1.1e](#).

The process is very similar to the above, except that you use a different fabric type and a different role:

- Create fabric using external connectivity network fabric template
- Discover the switch(es)
- Define role as Core
- Recalculate and Deploy

The following screenshot shows an example of creating a fabric using the external connectivity network fabric type:

Fabric Name
Core-Fabric

Pick Fabric
[External Connectivity Network >](#)

General Parameters | Advanced | Resources | Configuration Backup | Bootstrap | Flow Monitor

BGP AS #*
65534 1-4294967295 | 1-65535[0-65535] It is a good practice to have a unique ASN for each Fabric.

Fabric Monitor Mode If enabled, fabric is only monitored. No configuration will be deployed

Enable Performance Monitoring (For NX-OS Switches Only)

The following screenshot shows a switch defined as Core, with a Recalculate and Deploy executed:

Fabric Overview - Core-Fabric Actions Refresh Help Close

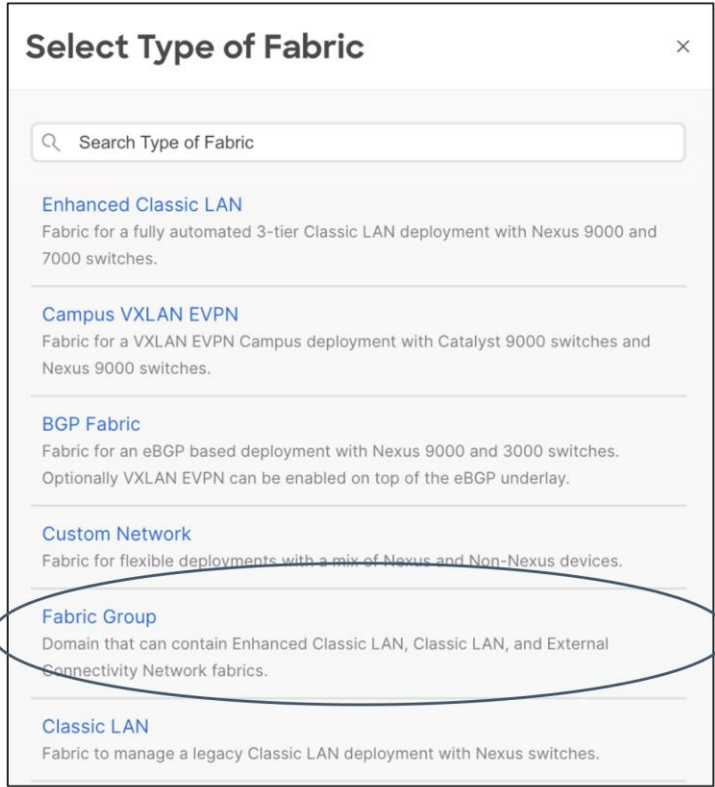
Overview | **Switches** | Links | Interfaces | Policies | Event Analytics | History | Resources | Virtual Infrastructure | Metrics

Filter by attributes Actions

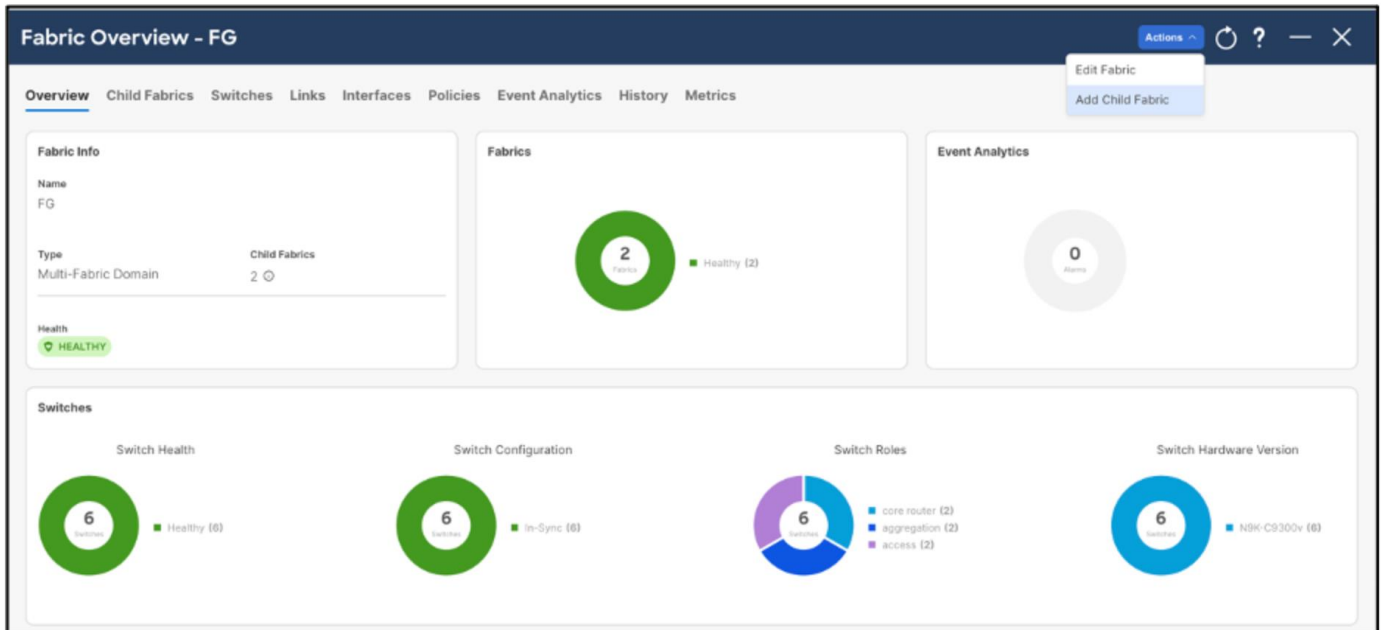
Switch	IP Address	Role	Serial Number	Mode	Config Status	Oper Status	Discovery Status	Model	VPC Role	VPC Peer
<input type="checkbox"/> Core1	192.18.0.22	Core Router	9FBR2OMC7SJ	Normal	In-Sync	Healthy	Ok	N9K-C9300v		
<input type="checkbox"/> Core2	192.18.0.23	Core Router	92NWQIKMBHH	Normal	In-Sync	Healthy	Ok	N9K-C9300v		

For a Group of Fabrics

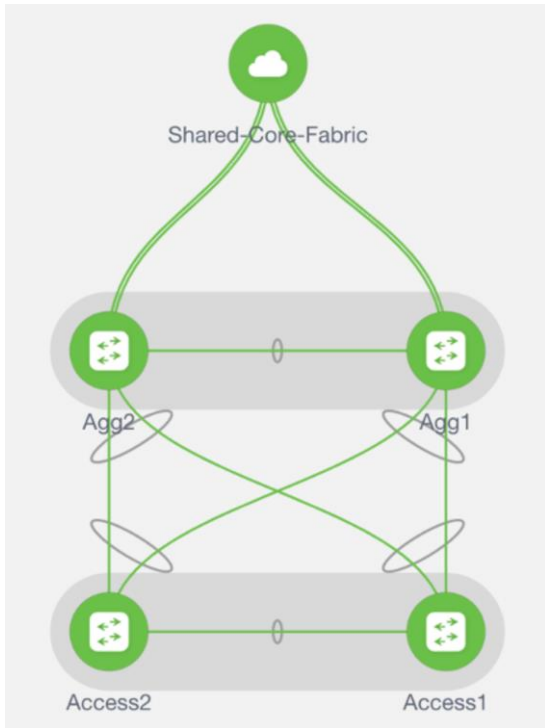
If you require group visualization for a topological view or an option to do a group deployment for switches in different fabrics (example Core and Aggregation), you can create the Fabric Group fabric type can with the Access-Aggregation and Core fabrics as child members of this group. Optionally, you can add other Enhanced Classic LAN fabrics to the group. These fabrics are considered child fabrics of the Fabric Group.



After you create the group, you must add the child fabrics:



The following screenshot shows how the Topology will appear:



You can choose right-click operations from the Topology page per switch or fabric. From the Fabric Group, you can use a deploy option for a switch or group of switches that are part of the child fabrics, which is useful for VRF-Lite level operations as discussed in the [Day 1 section](#).

Day 1 for Classic LAN

After you created the fabrics with the appropriate switches and the vPC-based topology is up and running, it is time to deploy the networks and VRF instances, and provision VRF-Lite. NDFC supports both IPv4 and IPv6 options. Classic LAN deployments often have single or very few VRF instances.

For a brownfield import scenario, the existing networks will be learned with a suffix "Auto" as shown in the following screenshot:

Network Name	VRF Name	IPv4 Gateway/Suffix	IPv6 Gateway/Prefix	Network Status	VLAN ID	Interface Group
<input type="checkbox"/> Auto_Net_VLAN301	default	192.168.31.1/24		DEPLOYED	301	
<input type="checkbox"/> Auto_Net_VLAN300	NA			DEPLOYED	300	
<input type="checkbox"/> Auto_Net_VLAN2303	NA			DEPLOYED	2303	
<input type="checkbox"/> Auto_Net_VLAN2302	myvrf_50002	192.168.1.1/24		DEPLOYED	2302	
<input type="checkbox"/> Auto_Net_VLAN2304	NA			DEPLOYED	2304	
<input type="checkbox"/> Auto_Net_VLAN32	vrf-prod	192.168.32.1/24		DEPLOYED	32	
<input type="checkbox"/> Auto_Net_VLAN34	vrf-prod		2001::192:168:34:1/120	DEPLOYED	34	
<input type="checkbox"/> Auto_Net_VLAN2301	NA			DEPLOYED	2301	
<input type="checkbox"/> Auto_Net_VLAN2300	nonprod	1.1.1/24	2001::1:1:1:1/112	DEPLOYED	2300	

The existing VRF instances are learned as well, as shown in the following screenshot:

Overview Switches Links Interfaces Interface Groups Policies Networks **VRFs** Event Analytics History Resources Virtual Infrastructure Metrics

Filter by attributes Actions

<input type="checkbox"/> VRF Name	VRF Status
<input type="checkbox"/> default	DEPLOYED
<input type="checkbox"/> vrf-prod	DEPLOYED
<input type="checkbox"/> myvrf_50002	DEPLOYED
<input type="checkbox"/> nonprod	DEPLOYED

You can always perform an edit operation for learned networks and VRF instances using NDFC. The existing configurations have been mapped to a predefined template that provides an intuitive workflow to create or edit as shown in the following screenshot:

Edit Network

Network Name*
Auto_Net_VLAN34

Layer 2 Only

VRF Name*
vrf-prod Create VRF

VLAN ID*
34 Propose VLAN

Network Template*
Network_Classic >

General Parameters **Advanced**

IPv4 Gateway/NetMask
 Example 192.0.2.1/24. Address for FHRP VIP

Interface IPv4 addr on active
 example 192.0.2.2. Interface IP address on the active/master device

Interface IPv4 addr on standby
 example 192.0.2.3. Interface IP address on the standby/backup device

IPv6 Gateway/NetMask
 IPv6 address for VIP. For VRRPv3, this is the VRRP secondary global IPv6 address.

Interface IPv6 addr on active*
 Interface IPv6 address on the active/master device.

Interface IPv6 addr on standby*

Edit VRF

VRF Name*
vrf-prod

VLAN ID*
2003 Propose VLAN

VRF Template*
[VRF_Classic >](#)

General Parameters **Advanced**

Routing Protocol
eBGP VRF Lite Agg-Core or Collapsed Core-WAN Peering Protocol (from Fabric Settings)

OSPF Process Tag
 OSPF Routing Process Tag (from Fabric Settings)

OSPF Area Id
 OSPF Area Id in IP address format

OSPFv3 Process Tag
 OSPFv3 Routing Process Tag (from Fabric Settings)

OSPFv3 Area Id
 OSPFv3 Area Id in IP address format

VRF VLAN Name
 If > 32 chars enable:system vlan long-name

VRF Interface Description

VRF Description

Now, let's look at the creation of new networks and VRF instances. This is applicable for both brownfield and greenfield networks.

Day 1 workflows fall in the following categories:

1. [Layer 2 Network](#)
2. [Layer 3 Network in Default VRF](#)
3. [Layer 3 Network with Custom VRF](#)
4. [VRF-Lite extension between the Aggregation and Core layers](#)
5. [VRF-Lite extension between Collapsed Core and WAN](#)

Two new templates, Network_Classic and VRF_Classic, have been introduced to incorporate use cases for classic Ethernet.

Note: Network Names and **VRF Names** are auto-populated on creation. NDFC also has the "Propose VLAN" option for networks and VRF instances. You can customize all of these fields. The NDFC Resource Manager also tracks all these parameters, which keeps a database of used resources to avoid conflicts.

Layer 2 Network

A Layer 2 network is easy to create. The gateway for a Layer 2 network resides outside of the fabric; hence the IP addresses are left empty. You can input an associated VLAN or let NDFC 'propose a VLAN based on the available resources (the range is customizable in the fabric settings).

After you create the networks, you can attached the networks to host-facing ports on the Access switch, which will thereby allow the VLAN on these Trunk or Access ports, and also on the vPC, port channel, and standalone ports between the Access and Aggregation layers.

You only need to specify intent to attach networks on host-facing ports. All the other interfaces between the Access and Aggregation layers as well as the Aggregation layer will automatically inherit the respective VLANs to allow end-to-end communication without you having to define this explicitly, making the operation trivial.

Step 1: Create the Network

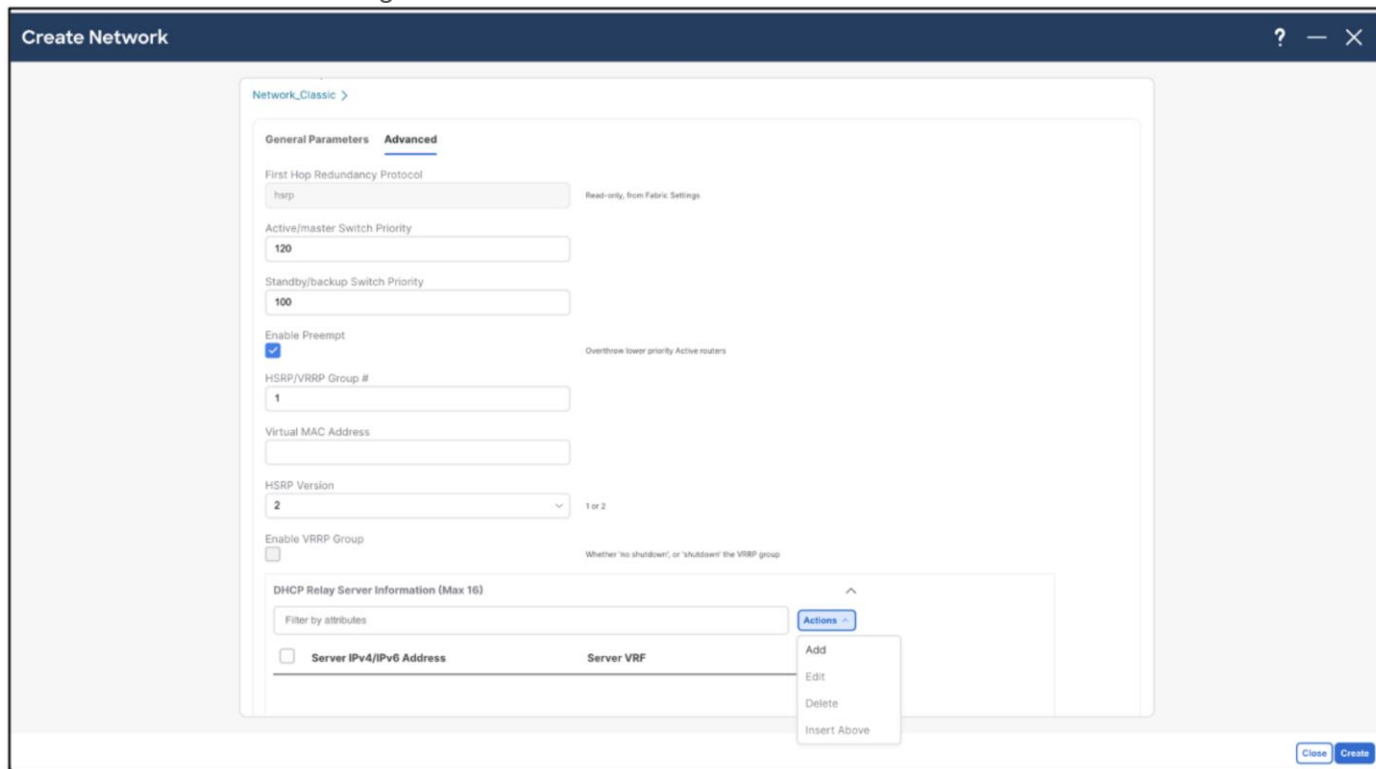
The screenshot shows the 'Create Network' configuration window. At the top, there's a title bar with a question mark, a minus sign, and a close button. The main content area is a form with the following sections:

- Network Name***: A text input field containing 'MyNetwork_30009'.
- Layer 2 Only**: A checkbox that is checked.
- VRF Name***: A dropdown menu showing 'NA' and a 'Create VRF' button.
- VLAN ID***: A text input field containing '2305' and a 'Propose VLAN' button.
- Network Template***: A dropdown menu showing 'Network_Classic' with a right-pointing arrow.
- General Parameters**: A section with two tabs, 'General Parameters' (selected) and 'Advanced'. It contains several input fields with example text:
 - IPv4 Gateway/NetMask**: An empty text field. Example: 192.0.2.1/24. Address for FHRP VIP.
 - Interface IPv4 addr on active**: An empty text field. Example: 192.0.2.2. Interface IP address on the active/master device.
 - Interface IPv4 addr on standby**: An empty text field. Example: 192.0.2.3. Interface IP address on the standby/backup device.
 - IPv6 Gateway/NetMask**: An empty text field. Example: IPv6 address for VIP. For VRRPv3, this is the VRRP secondary global IPv6 address.

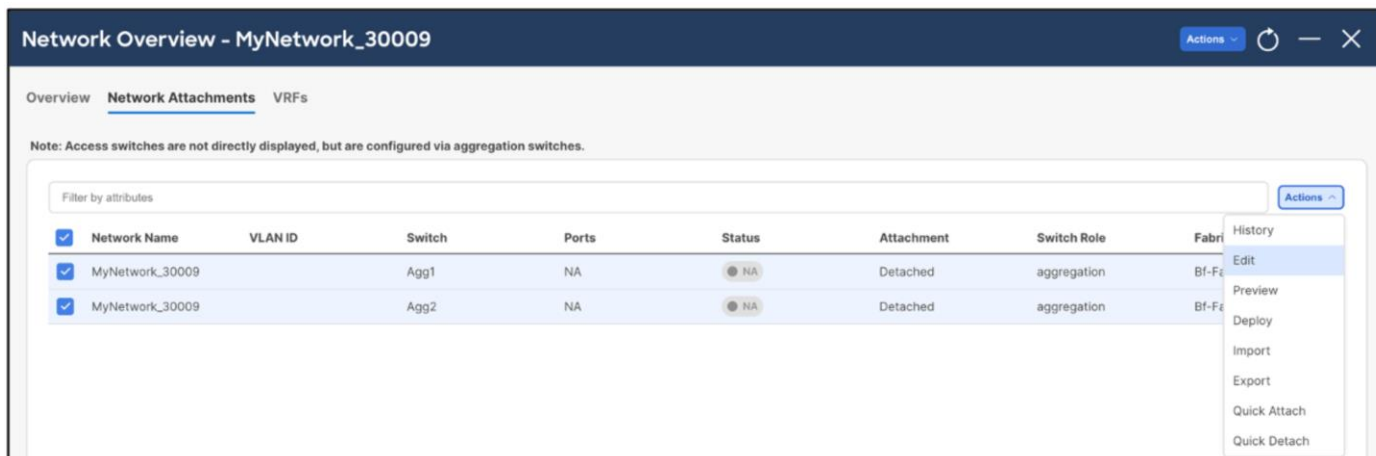
At the bottom right of the form, there are 'Close' and 'Create' buttons.

Advanced settings include adding DHCP relay server information and editing the default HSRP/VRRP settings.

Note: You cannot change the FHRP protocol from HSRP to VRRP or back from this screen. These are inherited from the fabric settings.

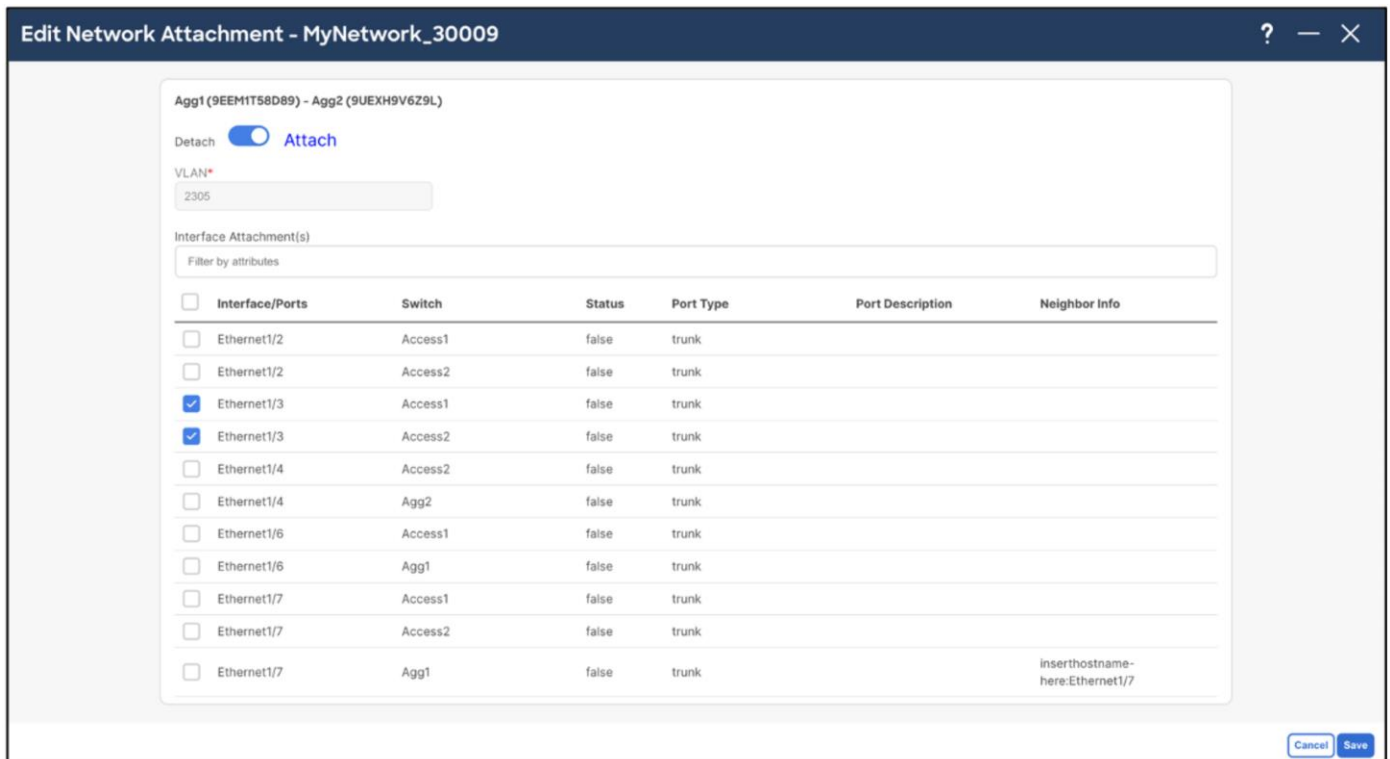


Step 2: Attach the Network



You attach the networks to the host-facing interfaces of the Access switches by selecting the Aggregation switches where the Access devices are connected. The Access switches will be connected to a unique pair of Aggregation devices, and the host VLANs must be allowed to the Aggregation. Hence, nNetwork attachments begin by selecting the respective Aggregation pair.

You must select the Access ports of interest for the network attachment as a next step. The uplink from Access and downlink from Aggregations will thereafter be auto-configured. NDFC handles various topologies, including B2B vPC.



Step 3: Review Pending Configurations on Access and Aggregation

This step includes allowing the VLAN on the host-facing port on the Access and the port channels between the Aggregation and Access layers.

The following screenshot shows the pending configurations for Access1:

```

Pending Config

vlan 2305
configure terminal
interface ethernet1/3
    switchport trunk allowed vlan add 2305
interface port-channel1
    switchport trunk allowed vlan add 2305
  
```

The following screenshot shows the pending configurations for Aggregation1:

Pending Config

```
vlan 2305
configure terminal
interface port-channel1
  switchport trunk allowed vlan add 2305
```

Step 4: Deploy the Configuration

Deploy Configuration - Classic-Demo

Filter by attributes

Network Name	Fabric Name	Switch Name	Serial Number	IP Address	Role	Network Status	Status Description	Progress
MyNetwork_30000	Classic-Demo	Agg1	92PV30FWOHR	10.30.12.15	aggregation	In-Sync	Config compliance sync completed	<div style="width: 100%;"></div>
MyNetwork_30000	Classic-Demo	Agg2	922L6WNL45G	10.30.12.18	aggregation	In-Sync	Config compliance sync completed	<div style="width: 100%;"></div>
MyNetwork_30000	Classic-Demo	Access2	9NUXUAWS5IO	10.30.12.17	access	In-Sync	Config compliance sync completed	<div style="width: 100%;"></div>
MyNetwork_30000	Classic-Demo	Access1	9L73CUD6EB9	10.30.12.16	access	In-Sync	Config compliance sync completed	<div style="width: 100%;"></div>

Layer 3 Network in the Default VRF Instance

A Layer 3 network can either be in a default or custom VRF instance. This section creates and attaches an Layer 3 network to a default VRF instance. You must define the v4/v6 gateway virtual IP address. The IP addresses for FHRP active and standby must be defined for the network under **General Parameters**. This time the gateway for the network is the Aggregation switch (or Collapsed Core) within the Enhanced Classic fabric.

You can choose the Aggregation switch that will be the FHRP Active during the network attachment. NDFC auto-selected a default value based on hashing.

You can customize FHRP settings under the **Advanced** tab. Based on the fabric settings, you can choose either HSRP or VRRP.

You can customize FHRP settings under **Advanced** tab. Based on the fabric settings, you can choose either HSRP or VRRP.

Step 1: Create the Network

The screenshot shows a 'Create Network' configuration window with the following fields and options:

- Network Name***: MyNetwork_30010
- Layer 2 Only**:
- VRF Name***: default (with a 'Create VRF' button)
- VLAN ID***: 2306 (with a 'Propose VLAN' button)
- Network Template***: Network_Classic >
- General Parameters** (selected) / **Advanced**
- IPv4 Gateway/NetMask**: 192.168.1.1/24 (Example: 192.0.2.1/24. Address for FHRP VIP)
- Interface IPv4 addr on active***: 192.168.1.10 (example: 192.0.2.2. Interface IP address on the active/master device)
- Interface IPv4 addr on standby***: 192.168.1.11 (example: 192.0.2.3. Interface IP address on the standby/backup device)
- IPv6 Gateway/NetMask**: (empty) (IPv6 address for VIP. For VRRPv3, this is the VRRP secondary global IPv6 address.)

Buttons at the bottom right: Close, Create

The advanced settings are the same as described in the [Layer 2 Network](#) section.

Step 2: Attach the Network and Choose the FHRP Master per Network

Edit Network Attachment - MyNetwork_30010

Agg1 (9EEM1T58D89) - Agg2 (9UEXH9V6Z9L)

Detach Attach

VLAN*
2306

FHRP Active
 Agg1 (9EEM1T58D89)
 Agg2 (9UEXH9V6Z9L)

Interface Attachment(s)
Filter by attributes

<input type="checkbox"/>	Interface/Ports	Switch	Status	Port Type
<input type="checkbox"/>	Ethernet1/2	Access1	false	trunk
<input type="checkbox"/>	Ethernet1/2	Access2	false	trunk
<input type="checkbox"/>	Ethernet1/3	Access1	false	trunk
<input type="checkbox"/>	Ethernet1/3	Access2	false	trunk
<input type="checkbox"/>	Ethernet1/4	Access2	false	trunk
<input type="checkbox"/>	Ethernet1/4	Agg2	false	trunk
<input checked="" type="checkbox"/>	Ethernet1/6	Access1	false	trunk
<input type="checkbox"/>	Ethernet1/6	Agg1	false	trunk
<input type="checkbox"/>	Ethernet1/7	Access1	false	trunk
<input checked="" type="checkbox"/>	Ethernet1/7	Access2	false	trunk

Step 3: Review Pending Configurations on the Access and Aggregation Layer

This includes allowing the VLAN on the host-facing port on the Access layer and the port channels between the Aggregation and Access layers. For Aggregation switches, in the case of a Layer 3 network, the configurations additionally include creating an SVI with the HSRP configurations, with lower priority for FHRP Active. The gateway and HSRP active/standby IP addresses used here are per-user inputs when creating a network.

The following screenshot shows the pending configurations for Access1:

Pending Config

```
vlan 2306
configure terminal
interface ethernet1/6
  switchport trunk allowed vlan add 2306
interface port-channell
  switchport trunk allowed vlan add 2306
```

The following screenshot shows the pending configurations for Aggregation1:

Pending Config

```
vlan 2306
interface Vlan2306
  ip address 192.168.1.11/24 tag 12345
  no ip redirects
  no ipv6 redirects
  no shutdown
  hsrp version 2
  hsrp 1
    ip 192.168.1.1
    preempt
exit
configure terminal
interface port-channell
  switchport trunk allowed vlan add 2306
```


Step 4: Deploy the Configuration

Deploy Configuration - Classic-Demo								
Filter by attributes								
Network Name	Fabric Name	Switch Name	Serial Number	IP Address	Role	Network Status	Status Description	Progress
MyNetwork_30001	Classic-Demo	Agg1	92PV30FWOHR	10.30.12.15	aggregation	In-Sync	Config compliance sync completed	<div style="width: 100%;"></div>
MyNetwork_30001	Classic-Demo	Agg2	922L6WNL45G	10.30.12.18	aggregation	In-Sync	Config compliance sync completed	<div style="width: 100%;"></div>
MyNetwork_30001	Classic-Demo	Access2	9NUXUAWS510	10.30.12.17	access	In-Sync	Config compliance sync completed	<div style="width: 100%;"></div>
MyNetwork_30001	Classic-Demo	Access1	9L73CUD6EB9	10.30.12.16	access	In-Sync	Config compliance sync completed	<div style="width: 100%;"></div>

Layer 3 Network with a Custom VRF Instance

As with the [Layer 3 Network in the Default VRF Instance](#) scenario, you can create a Layer 3 network for a custom VRF instance instead of a default VRF instance. The only extra step includes creating a new VRF instance using the VRF_Classic template. NDFC picks the routing protocol for VRF_Lite from the fabric settings. There is also a flag to control per VRF instance iBGP/OSPF peering between Aggregations. You can enter the IP addresses for this when attaching the VRF instance, as discussed in the [VRF-Lite Extension Between the Aggregation and Core/Edge Layers](#) section.

Step 1: Create the Network

Create Network

Network Name*

Layer 2 Only

VRF Name*

VLAN ID*

Network Template*
[Network_Classic >](#)

General Parameters [Advanced](#)

IPv4 Gateway/NetMask
 Example 192.0.2.1/24. Address for FHRP VIP

Interface IPv4 addr on active*
 example 192.0.2.2. Interface IP address on the active/master device

Interface IPv4 addr on standby*
 example 192.0.2.3. Interface IP address on the standby/backup device

Step 2: Create a VRF Instance to Link a Custom VRF Instance to This Layer 3 Network

Create VRF

VRF Name*
MyVRF_50003

VLAN ID*
2004 Propose VLAN

VRF Template*
VRF_Classic >

General Parameters **Advanced**

Routing Protocol
eBGP ▼ VRF Lite Agg-Core or Collapsed Core-WAN Peering Protocol (from Fabric Settings)

OSPF Process Tag OSPF Routing Process Tag (from Fabric Settings)

OSPF Area Id OSPF Area Id in IP address format

OSPFv3 Process Tag OSPFv3 Routing Process Tag (from Fabric Settings)

OSPFv3 Area Id OSPFv3 Area Id in IP address format

VRF VLAN Name if > 32 chars enable:system vlan long-name

VRF Interface Description

VRF Description

Enable Auto Peering Between VPC Aggs Flag to Control per VRF IBGP/OSPF Peering between Aggs. The protocol to use is based on vrf-lite routing protocol in Fabric

General Parameters includes the **Enable peering per VRF between Aggregations** option. We recommend this option as a backup path to reach the Core switches should the link between the active FHRP to the Core goes down.

The **Advanced** tab includes options for BGP authentication, route maps, and static 0/0 configurations. That is, you can configure a default (0/0) route toward the Core layer.

General Parameters **Advanced**

Redistribute Direct Route Map

FABRIC-RMAP-REDIST-SUBNET

Max BGP Paths

1

1-64

Config Static 0/0 Route

Flag to Control Static Default Route Configuration

Enable BGP Authentication

BGP Password Key Encryption Type

Select an Option



VRF Lite BGP Key Encryption Type: 3 - 3DES, 7 - Cisco

BGP Neighbor Password

VRF Lite BGP neighbor password (Hex String)

Enable OSPF Authentication

Applicable to OSPF only, can only be enabled if OSPF Process Tag is configured

OSPF Authentication Key ID

(Min:0, Max:255)

OSPF Authentication Key

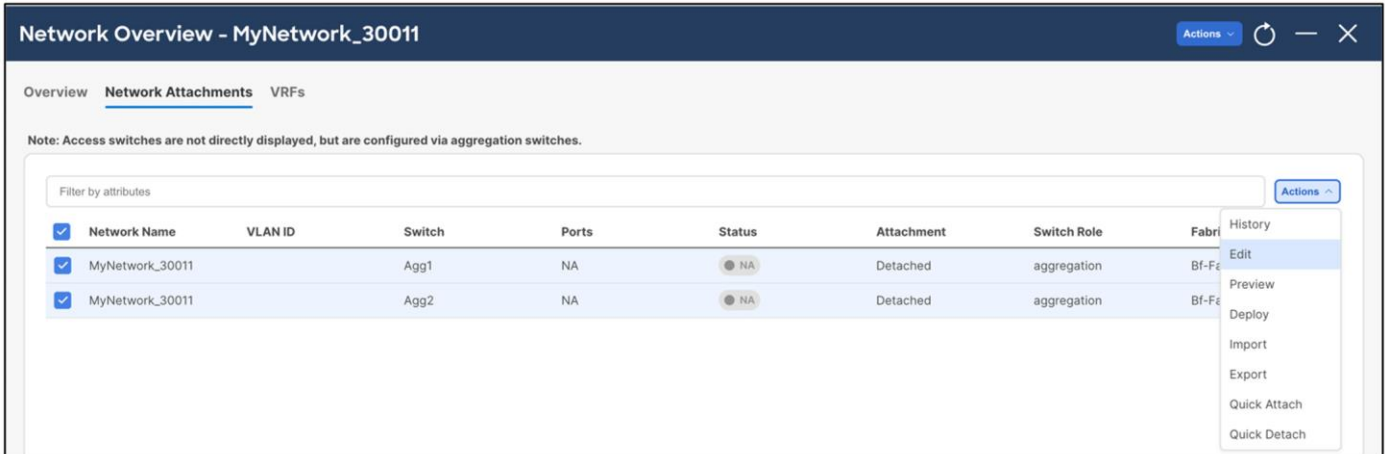
3DES Encrypted

Enable Netflow

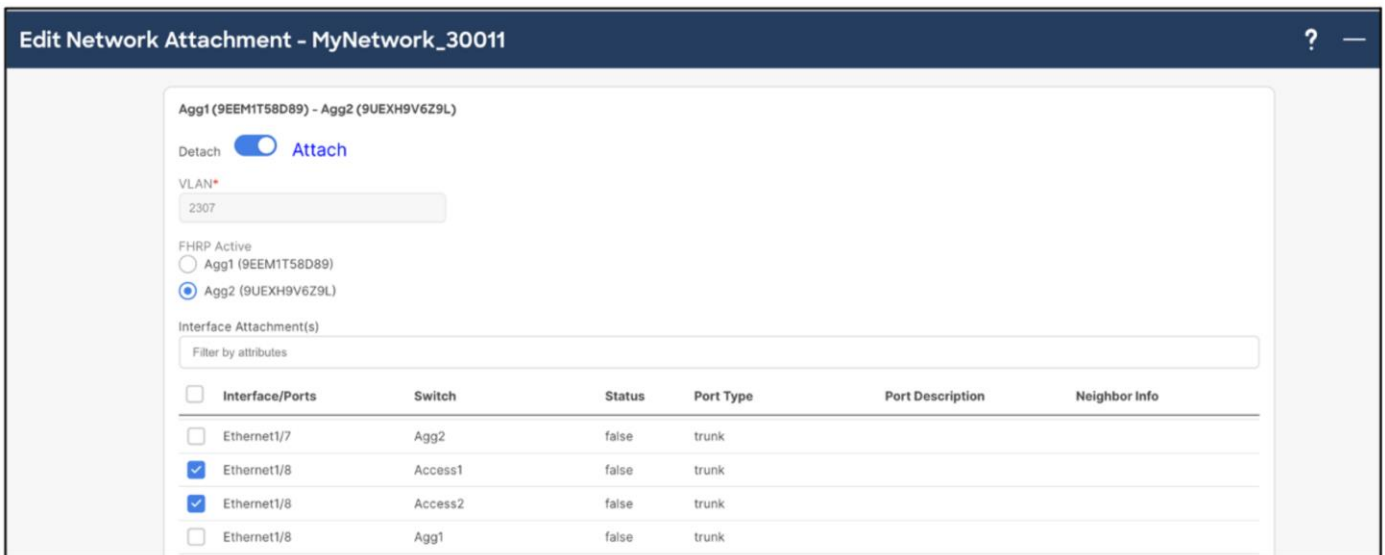
For netflow on VRF-LITE Sub-interface. Supported only if netflow is enabled on fabric

NetFlow Monitor

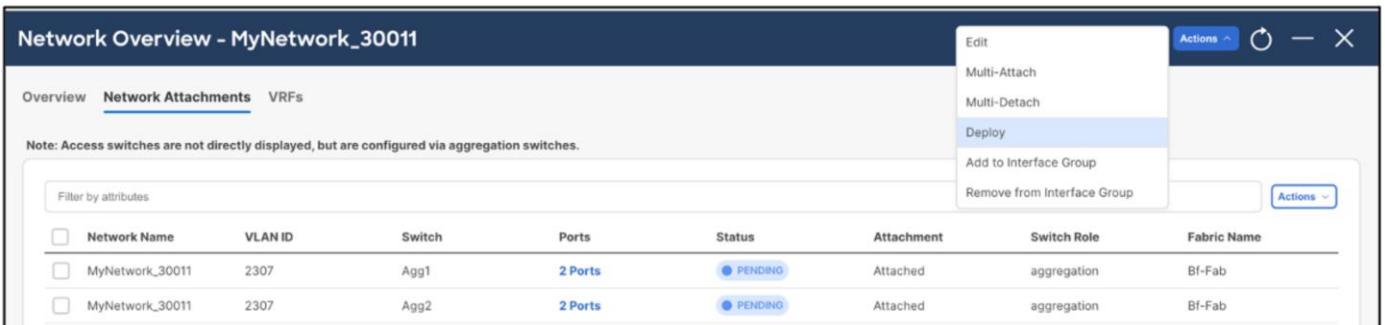
Step 3: Attach the Network



Choose the FHRP active:



Step 4: Preview and Deploy the Configuration



The following screenshot shows access configurations to allow VLAN on host-facing ports and port channels between the Access and Aggregation layers:

Pending Config

```
vlan 2307
configure terminal
interface ethernet1/8
    switchport trunk allowed vlan add 2307
interface port-channel1
    switchport trunk allowed vlan add 2307
```

Configuration at the Aggregation layer includes:

- Creating the VRF instance
- Creating an SVI for the VRF instance
- Instantiating a BGP session with route maps for the VRF instance that will be used for routing between the Aggregation and Core layers and also for between the Aggregation layers
- Creating an SVI for the gateway with relevant FHRP configurations
- Allowing the VLAN on port channel between the Access and Aggregation layers

Note: The route map is configured in the presence of SVI on Aggregations and for connectivity between Aggregations and Core. So, the Core has a reverse path to the Aggregation, the subnet needs to be advertised from the Aggregations to the Core. 'redistribute direct' is done because the subnets configured is always with a direct route, and we want to control which subnets to advertise. The route map thus matches the tag (12345), which is editable.

NDFC also has a knob in the fabric settings under the **Advanced** tab that enables you to disable the default route maps, and you can optionally use user-provided route maps.

Create Route-map fabric-rmap-redist-subnet



This route-map matches tag 12345

```

vrf context myvrf_50003
  address-family ipv4 unicast
  address-family ipv6 unicast
exit
router bgp 65535
  vrf myvrf_50003
    address-family ipv4 unicast
      redistribute direct route-map fabric-rmap-redist-subnet
    exit
    address-family ipv6 unicast
      redistribute direct route-map fabric-rmap-redist-subnet
configure terminal
vlan 2307
interface Vlan2307
  vrf member myvrf_50003
  ip address 192.169.1.10/24 tag 12345
  no ip redirects
  no ipv6 redirects
  no shutdown
  hsrp version 2
  hsrp 1
    ip 192.169.1.1
    priority 120
    preempt
exit
configure terminal
interface port-channel1
  switchport trunk allowed vlan add 2307

```

VRF-Lite Extension Between the Aggregation and Core/Edge Layers

NDFC supports Auto and Manual VRF-Lite between the Aggregation and Core/Edge layers. This document discusses Auto VRF-Lite configurations. NDFC supports the auto option with Cisco Nexus switches used for Core or Edge. The VRF Lite IP address version can be IPv4, IPv6, or IPv4 and IPv6.

Step 1: Fabric Settings

Under **Resources** in the fabric settings, set **Agg-Core/Edge Connectivity** to **Auto** and put a check in the **Auto Generate VRF_Lite Configurations on Aggregation and Core/Edge** check box. These are set by default and you can disable them. However, if the Core/Edge are Cisco Nexus switches, we recommend that you keep the default option of auto-generating the configurations, as this option will generate all the configurations you must deploy without you having to define VRF_Lite configurations manually.

General Parameters Spanning Tree VPC Protocols Advanced **Resources** Manageability Bootstrap Configuration Backup Flow Monitor

Network VLAN Range
 Per Switch Overlay Network VLAN Range (Min:2, Max:4094)

VRF VLAN Range
 Per Switch Overlay VRF VLAN Range (Min:2, Max:4094)

Subinterface Dot1q Range
 Per Border Dot1q Range For Agg-Core VRF Lite Connectivity (Min:2, Max:4093)

Agg-Core/Agg-Edge Connectivity
 VRF Lite Agg-Core and Agg-Edge Router Inter-Fabric Connection Options

Auto Generate VRF Lite Configuration on Agg and Core/Edge
 Flag that controls auto generation of VRF LITE sub-interface and peering configuration on Agg & Core/Edge devices. If set, auto created VRF Lite links will have 'Auto Generate Flag' enabled.

VRF Lite IP Version
 Choice of IPv4, IPv6 or both.

IPv4 VRF Lite Subnet IP Range*
 IPv4 Address range to assign P2P Agg-Core Connections

IPv4 VRF Lite Subnet Mask Length*

If you navigate to the fabric and check the links between the Aggregation and Core layers, the fabric must have the right template type attached to it with all parameters such as the source and destination Interfaces and BGP ASN all auto-populated.

Link Management - Edit Link : LINK-UUID-4420

Link Type*

Link Sub-Type*

Link Template*
[ext_fabric_setup >](#)

Source Fabric

Destination Fabric

Source Device*

Destination Device*

Source Interface*

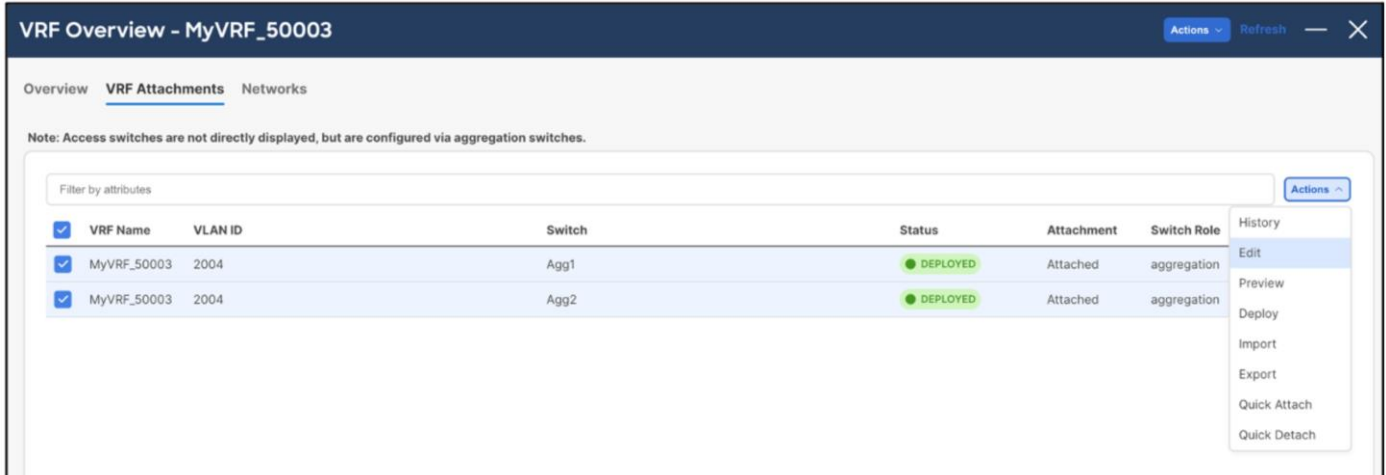
Destination Interface*

General Parameters	Advanced	Default VRF
Source BGP ASN*	<input type="text" value="65535"/>	BGP Autonomous System Number in Source Fabric
Source IP Address/Mask	<input type="text" value="10.33.0.1/30"/>	IP address for sub-interface in each VRF in Source Fabric
Destination IP Address*	<input type="text" value="10.33.0.2"/>	IP address for sub-interface in each VRF in Destination Fabric
Source IPv6 Address/Mask	<input type="text"/>	IPv6 address for sub-interface in each VRF in Source Fabric
Destination IPv6 Address	<input type="text"/>	IPv6 address for sub-interface in each VRF in Destination Fabric
Destination BGP ASN*	<input type="text" value="65534"/>	BGP Autonomous System Number in Destination Fabric
Link MTU	<input type="text" value="9216"/>	Interface MTU on both ends of VRF Lite IFC
Auto Generate Configuration for Peer	<input checked="" type="checkbox"/>	If enabled, auto generate VRF Lite configuration for managed NX-OS neighbor devices

As seen in the **General Parameters** of the link, the **Auto Generate Configuration for Peer** box has a check, which will generate configurations for Core/Edge without you having to do so manually.

Step 2: VRF Attachments

When you edit the VRF attachments of the VRF instance that you want to extend using VRF-Lite, NDFC shows a list of Aggregation to Core attachments that NDFC auto-detected with the respective VRF_Lite policy attached (for the auto generation of configurations) as shown in the following screenshot:

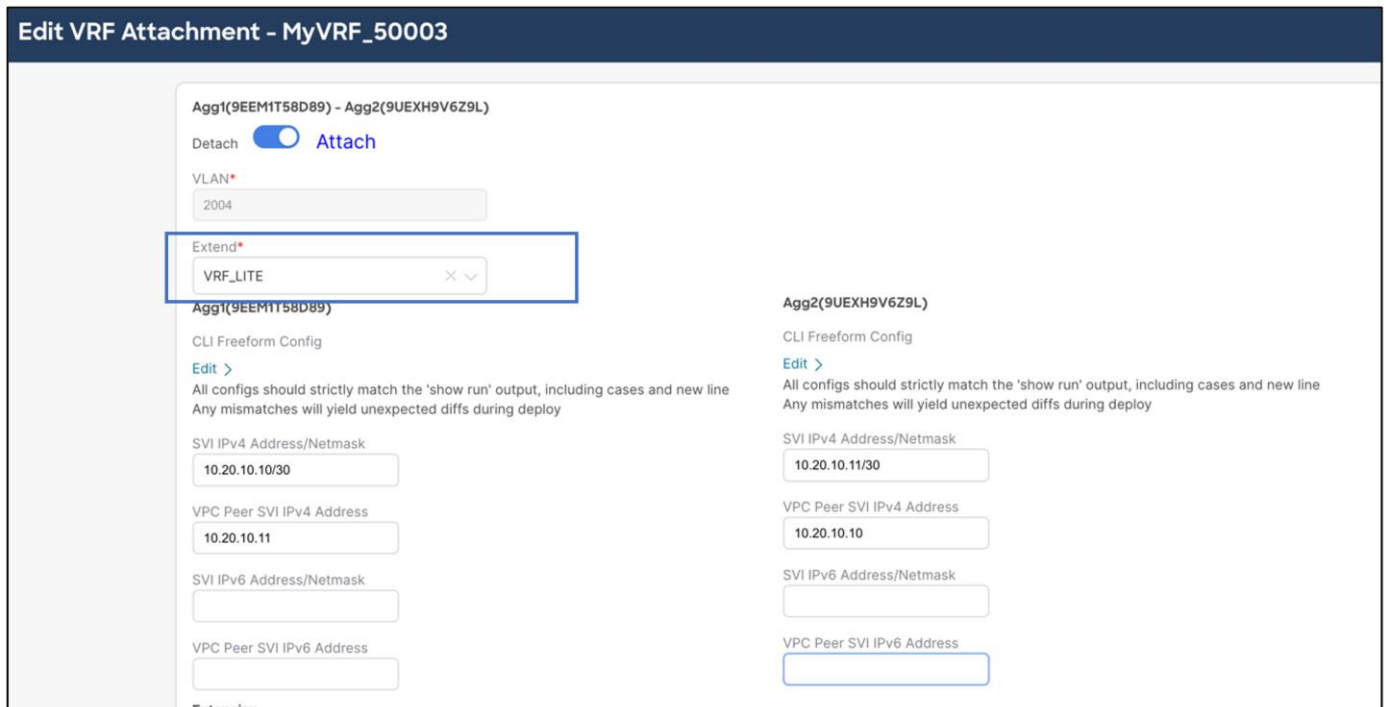


You can also see routing configurations between Aggregations on this screen.

Note: VRF instances have been already deployed on the Aggregation devices as result of the network provisioning that you previously performed.

You should set the **Extend** option to VRF Lite.

NDFC automatically picks the IP addresses to be used for peering between Aggregations from the IP address pool, which is auto-populated under **Fabric Settings > Resources**. You can customize this pool. This peering establishes routing between Aggregations to provide a backup route in case the links between FHRP Active (Aggregation) and Core goes down.



Next is establishing VRF_Lite between the Aggregation and Core/Edge layers. You can do an **Attach All** or click **edit** to selective extend specific connections.

Extension												
Filter by attributes											Attach-All	Detach-All
Action	Attached	Source Switch	Type	IF_NAME	Dest. Switch	Dest. Interface	DOT1Q_ID	IP_MASK	IP_TAG	NEIGHBO...	NEIGHBO...	I
Edit	Attached	Agg1	VRF_LITE	Ethernet1/2	Core1	Ethernet1/2	3	10.33.0.1/30		10.33.0.2	65534	
Edit	Attached	Agg1	VRF_LITE	Ethernet1/3	Core2	Ethernet1/3	3	10.33.0.6/30		10.33.0.5	65534	
Edit	Attached	Agg2	VRF_LITE	Ethernet1/3	Core1	Ethernet1/3	3	10.33.0.10/30		10.33.0.9	65534	
Edit	Attached	Agg2	VRF_LITE	Ethernet1/2	Core2	Ethernet1/2	3	10.33.0.14/30		10.33.0.13	65534	

After NDFC saves the intent to extend the VRF instance, you can perform a deploy operation first on the Aggregations in the Enhanced Classic LAN Fabric and then on the Core fabric for the Core.

Step 3: Deploy on Aggregations

Overview <u>Switches</u> Links Interfaces Interface Groups Policies Networks VRFs Event Analytics History Resources Virtual Infrastructure Metrics											
Filter by attributes											Actions
Switch	IP Address	Role	Serial Number	Mode	Config Status	Oper Status	Discovery Status	Model	VPC Role		
<input type="checkbox"/> Access1	192.18.0.13	Access	9HMPQ6NTXEV	Normal	In-Sync	Healthy	Ok	N9K-C9300v	Secondary		
<input type="checkbox"/> Access2	192.18.0.14	Access	9MZSTD2N4ML	Normal	In-Sync	Healthy	Ok	N9K-C9300v	Primary		
<input checked="" type="checkbox"/> Agg1	192.18.0.16	Aggregation	9EEM1T58D89	Normal	Pending	Healthy	Ok	N9K-C9300v	Primary		
<input checked="" type="checkbox"/> Agg2	192.18.0.17	Aggregation	9UEXH9V6Z9L	Normal	Pending	Healthy	Ok	N9K-C9300v	Secondary		

The **Pending Configuration** shows how NDFC creates sub-interfaces that are members of respective VRF instances, as well as the SVI for the VRF instance with the IP addresses entered during VRF instance extension (used for iBGP peering between Aggregations) and establishes peering between the Aggregation and Core layers for these subinterfaces.

Pending Config Side-by-Side Comparison

```
vlan 2004
interface Vlan2004
  vrf member myvrf_50003
  ip address 10.20.10.10/30
  mtu 9216
  no ip redirects
  no ipv6 redirects
  no shutdown
exit
router bgp 65535
  vrf myvrf_50003
    neighbor 10.33.0.2
      remote-as 65534
      address-family ipv4 unicast
        send-community both
      exit
    exit
  neighbor 10.33.0.5
    remote-as 65534
    address-family ipv4 unicast
      send-community both
    exit
  exit
  neighbor 10.20.10.11
    remote-as 65535
    address-family ipv4 unicast
      send-community both
```

```
interface ethernet1/2.3
  encapsulation dot1q 3
  mtu 9216
  vrf member myvrf_50003
  ip address 10.33.0.1/30
  no shutdown
interface ethernet1/3.3
  encapsulation dot1q 3
  mtu 9216
  vrf member myvrf_50003
  ip address 10.33.0.6/30
  no shutdown
```

Step 4: Deploy on Core

You must now perform the same deploy operation for the Core-Fabric to provision the pending configurations for the respective VRF-lite to the Core.

Fabric Overview - Core-Fabric

Overview **Switches** Links Interfaces Policies Event Analytics History Resources Virtual Infrastructure Metrics

Filter by attributes

<input checked="" type="checkbox"/>	Switch	IP Address	Role	Serial Number	Mode	Config Status	Oper Status	Discovery Status	Model	VPC Role
<input checked="" type="checkbox"/>	Core1	192.18.0.22	Core Router	9FBR20MC75J	Normal	Pending	Healthy	OK	N9K-C9300v	
<input checked="" type="checkbox"/>	Core2	192.18.0.23	Core Router	92NWQIKMBHH	Normal	Pending	Healthy	OK	N9K-C9300v	

Actions ^

- Add Switches
- Preview
- Deploy
- Discovery >
- Set Role
- vPC Pairing
- ToR/Access Pairing
- vPC Overview
- More >

VRF-Lite extension between Collapsed Core and WAN

In case of Collapsed Core, you might need to use VRF-Lite between the Aggregation Role (Collapsed Core) and the WAN Router. The Aggregation switch will be a Cisco Nexus 9000. The WAN router can be managed in NDFC and can either be a Cisco Nexus or non-Cisco Nexus device. The options are as follows:

- WAN router is a Cisco Nexus platform: When the WAN router is a Nexus device, the process remains exactly the same as when we have a Core switch. The WAN router should be discovered and managed in the external connectivity network fabric with its role listed as "CORE." NDFC auto-generates the VRF-Lite configurations. All the steps for VRF-Lite between the Aggregation and the WAN router remain the same as the [VRF-Lite Extension Between the Aggregation and Core/Edge Layers](#) section.
- WAN router is a non-Nexus (Cisco or Non-Cisco) platform: When the WAN router is a non-Nexus device, the router should still be discovered and managed in the external connectivity network fabric with the role of "CORE." However, the VRF-Lite configurations will not be auto-generated. You must to add the VRF-Lite Jython policy on the WAN router manually. See the following procedures in the *Cisco NDFC-Fabric Controller Configuration Guide, Release 12.1.1e*:

[VRF Lite between Cisco Nexus 9000 based Border and Non-Cisco device](#)

[VRF Lite between Cisco Nexus 9000 based Border and Non-Nexus device](#)

- VRF-Lite to WAN router using a routed interface or port channel: Whether the WAN router is a Nexus or non-Nexus device, there may be scenarios where you want to use VRF-Lite using SVIs or Routed Port/Port-Channels. NDFC VRF-lite auto-workflow supports VRF-Lite extension only for sub-interfaces. Users can deploy policies manually when VRF-Lite between the Collapsed Core and WAN is required using routed interfaces or SVIs.

The policies are as follows:

- Ext_VRF_Lite_SVI (for VRF-Lite using an SVI)
- Ext_VRF_Lite_Routed (for VRF-Lite using routed ports or port channels)

Day 2 for Classic LAN

All maintenance and operational features listed below are supported equally for Classic LAN networks and VXLAN fabrics:

- Image Management: upgrades, downgrades, EPLDs, RPMs, and SMUs
- Change Management and Rollback
- Inventory View
- Event Analytics
- Deployment History and Audit Logs
- Backup and Restore
- Performance Metrics, Link and Interface stats, and Protocol Views
- Programmable Reports
- Virtual infrastructure (VMM, K8s, OpenStack) Visibility

These features are agnostic to the fabric type. For more information about these features, see the Cisco NDFC Configuration Guide.

Note:

- To use alarms and get immediate notification of link/interfaces/switch down, you must configure NDFC as a trap destination.
- For syslogs, NDFC by default is not a syslog receiver. You must configure NDFC to be a syslog receiver, and thereafter you can define policies to capture syslog messages of interest and trigger the appropriate alarms. Performance monitoring does not require this, as performance monitoring is an SNMPv3 poll from NDFC to the switch.
- SCP is required for image management, NX-API certificate installation, NDI functionality, and POAP.
- SNMP is used for device discovery.
- Both SCP and SNMP pods are always enabled by default. You must have a minimum of 2 persistent IP addresses in Nexus Dashboard when enabling the fabric controller.

Integration of Classic LAN with Services

It is common for services such as firewalls and load balancers to be connected to Aggregations or Collapsed Core (a switch serving as Layer 2/Layer 3 boundary), with traffic being redirected to these services for security or traffic optimizations.

In scenarios with services such as a firewall attached to an Enhanced Classic LAN fabric, you must manually provision the respective configurations to the service devices. NDFC will not push any configurations to these.

You can achieve connectivity to the service device using the following options:

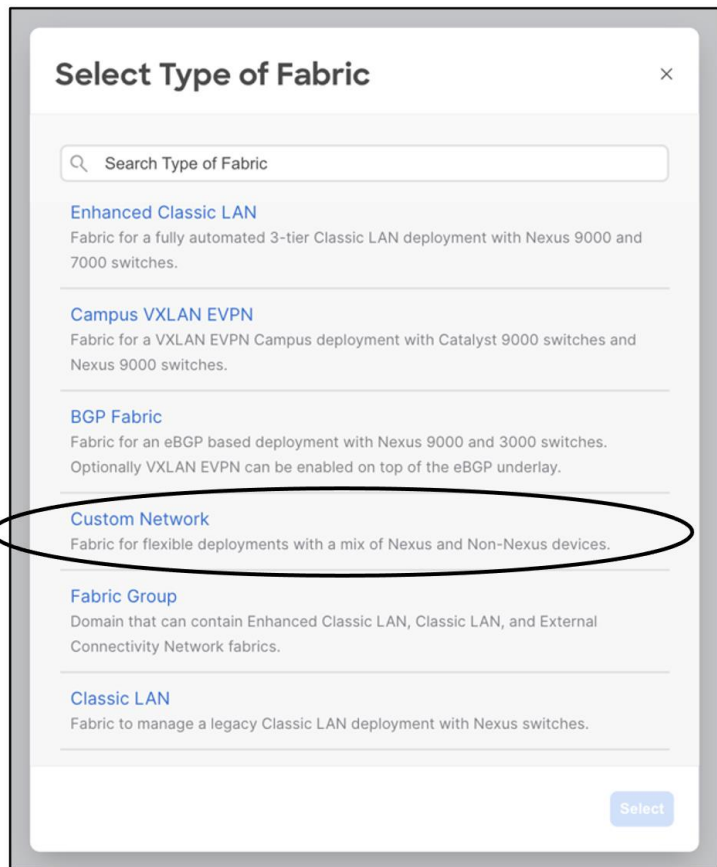
- [VRF-Lite Using Subinterfaces](#)
- [VRF-Lite Using SVIs](#)
- [VRF-Lite Using Routed Interfaces or Port Channels](#)

Note: The most common option is to use VRF-Lite with SVIs, as service nodes are typically deployed as a cluster (active/active or active/standby), and the cluster requires Layer 2 adjacency. You can use the subinterfaces and routed interface or port channels option in the case of a standalone service node (that is, you do not have a cluster).

The succeeding subsections discuss these workflows. For the configurations, assume the Aggregation switch is connected to the firewall.

VRF-Lite Using Subinterfaces

You must create a separate fabric using the "Custom" fabric type for the service device.



Create Fabric

Fabric Name

Pick Fabric
[Custom Network >](#)

General Parameters | Advanced | Resources | Configuration Backup | Bootstrap | Flow Monitor

BGP AS #*
 1-4294967295 | 1-65535[.0-65535] It is a good practice to have a unique ASN for each Fabric.

Fabric Monitor Mode
 If enabled, fabric is only monitored. No configuration will be deployed

Enable Performance Monitoring (For NX-OS Switches Only)

The firewalls that connect to Enhanced Classic LAN fabric type do not support the existing Layer 4 to Layer 7 services workflow in NDFC as of release 12.1.3. Hence, these firewalls will be added as a meta device. From a routing or VRF-lite perspective, devices these are considered to be unmanaged.

You must browse to the Links under the Enhanced Classic LAN fabric and choose **Actions > Create**.

Overview | Switches | **Links** | Interfaces | Interface Groups | Policies | Networks | VRFs | Event Analytics | History | Resources | Virtual Infrastructure | Metrics

Links

Protocol View

Filter by attributes

Fabric Name	Name	Policy	Info	Admin State	Oper State
<input type="checkbox"/> Bf-Fab	Agg1-mgmt0---fanout-L2-GigabitEthernet1/0		Neighbor Present	↑ Up	↑ Up
<input type="checkbox"/> Bf-Fab	Agg2-mgmt0---fanout-L2-GigabitEthernet1/1		Neighbor Present	↑ Up	↑ Up
<input type="checkbox"/> Bf-Fab	Agg1-Ethernet1/7---inserthostname-here-Ethernet1/7		Neighbor Missing	↑ Up	↑ Up
<input type="checkbox"/> Bf-Fab	Agg2-Ethernet1/8---inserthostname-here-Ethernet1/8		Neighbor Missing	↑ Up	↑ Up

Actions ...

- Create
- Edit
- Delete
- Import
- Export

In the **Link Management - Create Link** dialog, perform the following procedure:

1. For **Link Type**, choose **Inter-fabric**.
2. For **Link Sub-Type**, choose **VRF-Lite**.
3. For **Source Fabric**, choose the Enhanced Classic LAN fabric that you created.
4. For **Destination Fabric**, choose **Services**.
5. For **Source Device**, choose the switch that the firewall is attached to, such as the Aggregation switch.
6. For **Destination Device**, choose the name of the firewall.
7. For **Source Interface**, choose the interface, such as "Ethernet1/27" on the Aggregation switch.
8. For **Destination Interface**, choose an interface.
9. For **Source BGP ASN**, enter the BGP autonomous system number that is in the source fabric.

10. For **Source IP Address/Mask**, enter the IP address and mask for the Ethernet1/27 subinterface, which is the source interface of the inter-fabric connection (IFC). NDFC creates a subinterface for each VRF instance that is extended over this IFC and assigns a unique 802.1Q ID to the subinterface. NDFC uses the IP address/mask that you enter, along with the value of the **BGP Neighbor IP** field, as the default values for the subinterface that NDFC creates at the VRF extension. You can overwrite the default values.
11. For **Destination IPv6 Address**, enter the BGP neighbor IP address on the metadvice.
12. Click **Save**.

For more information, see the [Cisco NDFC-Fabric Controller Configuration Guide, Release 12.1.2e](#).

Link Management - Create Link

Link Type*
Inter-Fabric

Link Sub-Type*
VRF_LITE

Link Template*
ext_fabric_setup >

Source Fabric*
BF-Fab

Destination Fabric*
Services

Source Device*
Agg1

Destination Device*
FirewallDevice

Source Interface*
Ethernet1/27

Destination Interface*
Ethernet1/2

General Parameters | Advanced | Default VRF

Source BGP ASN*
65535
BGP Autonomous System Number in Source Fabric

Source IP Address/Mask
15.15.10.1/30
IP address for sub-interface in each VRF in Source Fabric

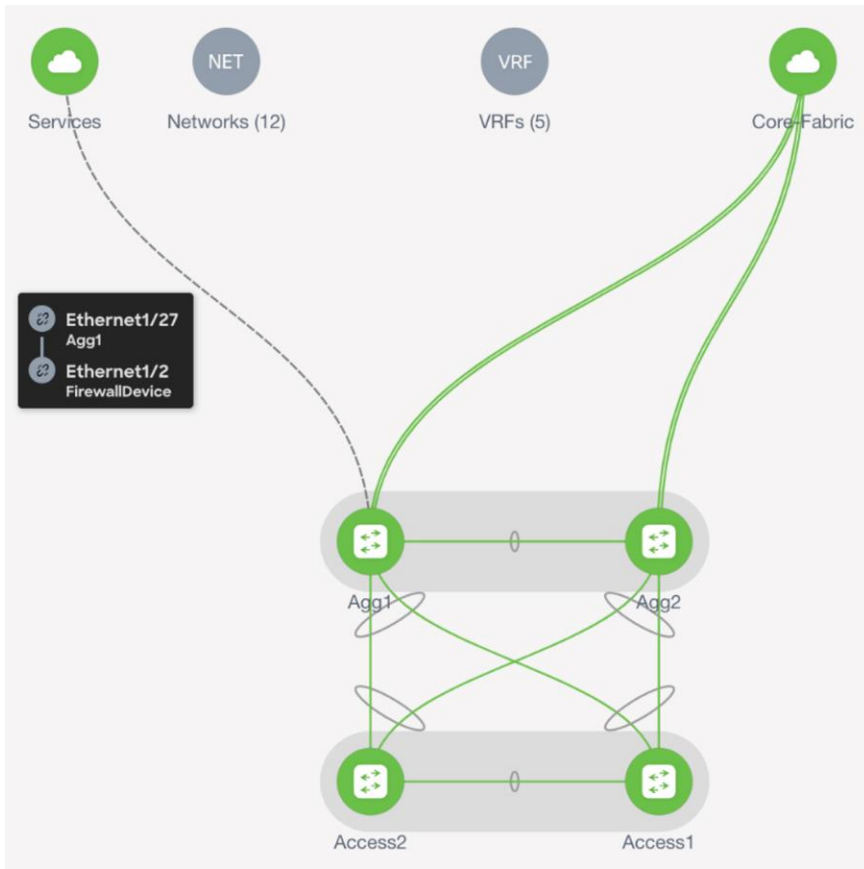
Destination IP Address*
15.15.10.2
IP address for sub-interface in each VRF in Destination Fabric

Source IPv6 Address/Mask
IPv6 address for sub-interface in each VRF in Source Fabric

Destination IPv6 Address
IPv6 address for sub-interface in each VRF in Destination Fabric

Cancel Save

After saving, NDFC creates the following link between the two fabrics, with the firewall as a metadvice:



Make sure to deploy the pending configurations on the Aggregation switch, which marks the source interface defined as a routed port.

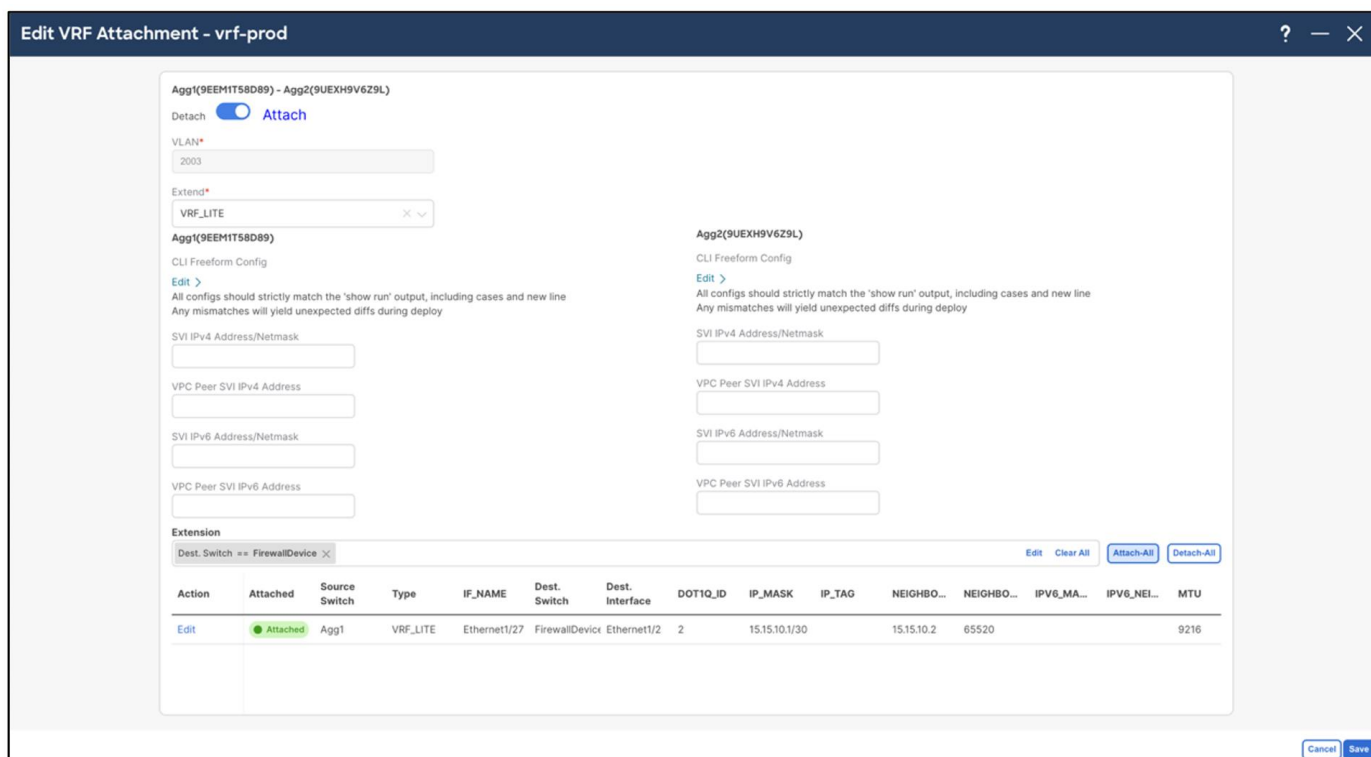
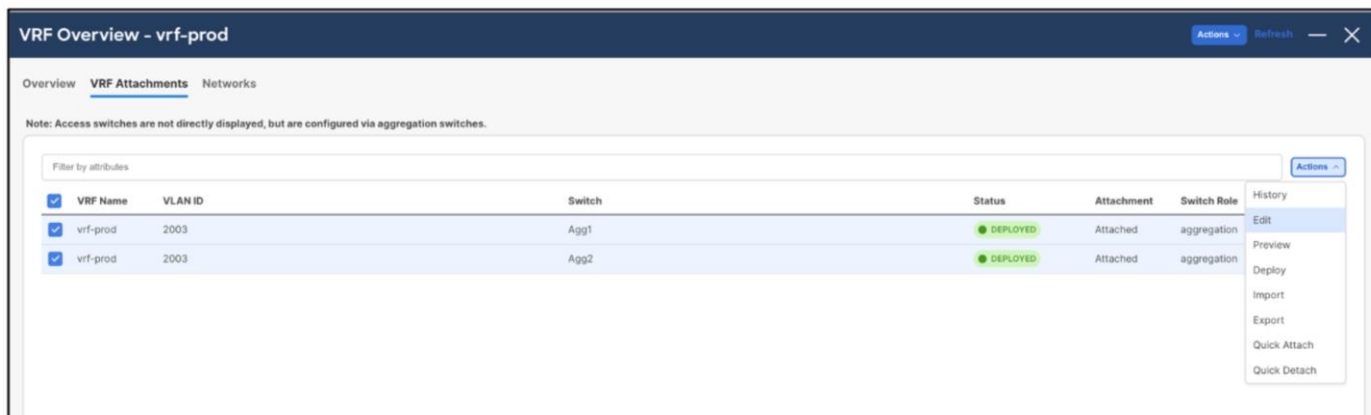
Switch	IP Address	Role	Serial Number	Mode	Config Status	Oper Status	Discovery Status	Model	VPC Role	VPC Peer	Software Ver
<input type="checkbox"/> Access1	192.18.0.13	Access	9HMPQ6NTXEV	Normal	In-Sync	Healthy	Ok	N9K-C9300v	Secondary	Access2	9.3(9)
<input type="checkbox"/> Access2	192.18.0.14	Access	9MZSTD2N4ML	Normal	In-Sync	Healthy	Ok	N9K-C9300v	Primary	Access1	9.3(9)
<input checked="" type="checkbox"/> Agg1	192.18.0.16	Aggregation	9EEMIT58D89	Normal	Pending	Healthy	Ok	N9K-C9300v	Primary	Agg2	9.3(9)
<input type="checkbox"/> Agg2	192.18.0.17	Aggregation	9UEXH9V6Z9L	Normal	In-Sync	Healthy	Ok	N9K-C9300v	Secondary	Agg1	9.3(9)

Pending Config Side-by-Side Comparison

```
interface ethernet1/27
  no switchport trunk allowed vlan none
  no switchport mode trunk
interface ethernet1/27
  mtu 9216
  no switchport
  no shutdown
```

The VRF-Lite workflow is identical to VRF-Lite between Aggregation and Core, as described in the [Day1 for Classic LAN](#) section.

The following screenshots summarize the process of configuring the VRF-Lite extension on a VRF instance that you created:



For the procedure to create a new VRF instance, see the [Layer 3 Network with a Custom VRF Instance](#) section.

You can enter the peer VRF instance by editing the extension. The IP addresses shown in the screenshot are what you entered when you created the links, and you can overwrite the IP addresses at this point. Click **Save** to save the extension.

Edit Extension Details

DOTIQ_ID*

2

IP_MASK

15.15.10.1/30

IP_TAG

NEIGHBOR_IP

15.15.10.2

NEIGHBOR_ASN

65520

IPV6_MASK

IPV6_NEIGHBOR

MTU

9216

ENABLE_IFC_NETFLOW

AUTO_VRF_LITE_FLAG

PEER_VRF_NAME

prod

Cancel Save

Choose **Actions > Deploy** of the VRF instance.

VRF Overview - vrf-prod

Overview **VRF Attachments** Networks

Note: Access switches are not directly displayed, but are configured via aggregation switches.

Filter by attributes

VRF Name	VLAN ID	Switch	Status	Attachment	Switch Role
<input checked="" type="checkbox"/> vrf-prod	2003	Agg1	PENDING	Attached	aggregation
<input checked="" type="checkbox"/> vrf-prod	2003	Agg2	PENDING	Attached	aggregation

Actions

- History
- Edit
- Preview
- Deploy
- Import
- Export
- Quick Attach
- Quick Detach

The following configuration is generated for the Aggregation switch. You must provision the firewall/service device with equivalent configurations for the peering to be up.

```
router bgp 65535
  vrf vrf-prod
    neighbor 15.15.10.2
      remote-as 65520
    address-family ipv4 unicast
      send-community both
```

```

configure terminal
interface ethernet1/27.2
  encapsulation dot1q 2
  mtu 9216
  vrf member vrf-prod
  ip address 15.15.10.1/30
  no shutdown

```

VRF-Lite Using SVIs

In cases where firewalls do not support subinterfaces, you can use SVIs for VRF-Lite with eBGP. You must do this on the Aggregation layer using the Ext_VRF_lite_SVI policy that is included with NDFC.

For the procedure on how to add a policy in NDFC, see the [Cisco NDFC-Fabric Controller Configuration Guide, Release 12.1.2e](#). The following screenshots summarize the process:

The screenshot shows the 'Policies' page in the NDFC interface. The table lists several policies with columns for Policy ID, Switch, IP Address, Template, Description, Entity Name, Entity Type, Source, Priority, Content Type, Serial Number, and Editable. An 'Actions' menu is visible on the right side of the table.

Policy ID	Switch	IP Address	Template	Description	Entity Name	Entity Type	Source	Priority	Content Type	Serial Number	Editable
<input type="checkbox"/> POLICY-395080	Access1	192.18.0.13	switch_role_simul		SWITCH	SWITCH		10	PYTHON	9HMPQ6NTXI	true
<input type="checkbox"/> POLICY-395490	Access2	192.18.0.14	switch_role_simul		SWITCH	SWITCH		10	PYTHON	9MZSTD2N4A	true
<input type="checkbox"/> POLICY-394580	Agg1	192.18.0.16	switch_role_simul		SWITCH	SWITCH		10	PYTHON	9EEMIT58D8I	true
<input type="checkbox"/> POLICY-394830	Agg2	192.18.0.17	switch_role_simul		SWITCH	SWITCH		10	PYTHON	9UEXH9V6Z9	false
<input type="checkbox"/> POLICY-391940	Access1	192.18.0.13	host_1LJ		SWITCH	SWITCH		50	TEMPLATE_CLI	9HMPQ6NTXI	true
<input type="checkbox"/> POLICY-391950	Access1	192.18.0.13	pre_config		SWITCH	SWITCH	UNDERLAY	50	TEMPLATE_CLI	9HMPQ6NTXI	false

Select Switches

Search Switches

Select All Show Selected

- Access1
9HMPQ6NTXEV 192.18.0.13 access
- Access2
9MZSTD2N4ML 192.18.0.14 access
- Agg1
9EEM1T58D89 192.18.0.16 aggregation
- Agg2
9UEXH9V6Z9L 192.18.0.17 aggregation

Select (1)

Select Policy Template

Ext_VRF_lite_SVI

Ext_VRF_Lite_SVI
All

Input the parameters in the Ext_VRF_lite_SVI policy:

Create Policy ? - X

Switch List:

Priority*: 1-2000

Description:

Template Name:

General Parameters Advanced

VLAN ID*: 2-4094, value needs to be in switch allowed range

VRF Name: SVI VRF name, default VRF if not specified

SVI IPv4 Address/Netmask: For IPv4 VRF Lite peering

Neighbor IPv4 Address*: BGP Peer IPv4 Address

SVI IPv6 Address/Netmask: For IPv6 VRF Lite peering

Neighbor IPv6 Address: BGP Peer IPv6 Address

Neighbor ASN*: BGP ASN of IPv6/IPv6 Neighbor

Local ASN: The fabric ASN will be used if not specified

BGP Neighbor Password: Hex String

BGP Password Key Encryption Type:

To see the generated configuration, choose the policy, the **Actions -> Generated Config**. You can now push the policy to the Aggregation switch.

Overview **Switches** Links Interfaces Interface Groups **Policies** Networks VRFs Event Analytics History Resources Virtual Infrastructure Metrics

Filter by attributes

<input type="checkbox"/>	Policy ID	Switch	IP Address	Template	Description	Entity Name	Entity Type	Source	Priority	Content Type	Serial Number	Editable	
<input checked="" type="checkbox"/>	POLICY-442500	Agg1	192.18.0.16	Ext_VRF_Lite_SVI	VRF-Lite using SVI between Agg and Firewall	SWITCH	SWITCH		500	PYTHON	9EEMITS8DBI	true	Add Policy Edit Policy Delete Policy Generated Config Push Config
<input type="checkbox"/>	POLICY-392020	Access1	192.18.0.13	switch_freeform	Pre Interfaces Configuration	SWITCH	SWITCH		240	PYTHON	9HMPQ6NTXI	true	

The following screenshot shows the generated configuration for an SVI-based VRF-Lite:

Generated Config

```
#POLICY-442500#
vlan 55

interface Vlan55
 vrf member lab

interface Vlan55
 ip address 55.10.10.1/30

interface Vlan55
 mtu 9216

interface Vlan55
 no shutdown

vrf context lab

 address-family ipv4 unicast

router bgp 65535

 vrf lab

  address-family ipv4 unicast

  neighbor 55.10.10.2 remote-as 64400
  address-family ipv4 unicast
  send-community
  send-community extended
```

You must also provision the appropriate configurations on the firewall for the peering to come up.

VRF-Lite Using Routed Interfaces or Port Channels

VRF-Lite over eBGP can also be achieved using routed interfaces or port channels on the Aggregation switch. In this case, you must manually apply a policy using the add policy per switch workflow, as described in the [VRF-Lite Using SVIs](#) section.

Use the Ext_VRF_Lite_Routed NFDC policy for routed interfaces. The following screenshots show usage of the policy:

Switch List:

Priority*

1-2000

Description

Template Name

[Ext_VRF_Lite_Routed >](#)

General Parameters **Advanced**

Layer-3 Interface*

Physical or Port-channel interface (e.g. e1/14, Ethernet1/14)

VRF Name

VRF name, default VRF if not specified

IPv4 Address/Netmask*

For IPv4 VRF Lite peering

Neighbor IPv4 Address*

BGP Peer IPv4 Address

IPv6 Address/Netmask

For IPv6 VRF Lite peering

Neighbor IPv6 Address

BGP Peer IPv6 Address

Neighbor ASN*

BGP ASN of IPv4/IPv6 Neighbor

Local ASN

The fabric ASN will be used if not specified

BGP Neighbor Password

Hex String

BGP Password Key Encryption Type

BGP Key Encryption Type: 3 - 3DES, 7 - Cisco

MTU*

Layer-3 MTU (Min:576, Max:9216)

You must push the policy to the switches in Enhanced Classic LAN and you must provision the equivalent configurations on the firewalls.

Migration from Cisco Nexus 2000/5000/7000 Classic LAN networks to Cisco Nexus 7000/9000-based Classic LAN Networks

We recommend that you build newer classic Ethernet fabrics with Cisco Nexus 9000 switches in the Access, Aggregation, and Core layers, or that you have Cisco Nexus 9000 switches in the Access or Aggregation layer and use Cisco Nexus 7000 switches in the Core layer using the Enhanced Classic LAN fabric type. Cisco Nexus 7000 switches continue to be supported as a means to provide investment protection to customers who invested heavily in the platform and would like to continue to use the Cisco Nexus 7000 in the Core layer.

This section covers how networks comprising a mix of older platforms such as the Cisco Nexus 5000 and Cisco Nexus 6000 switches (which are not supported in an Enhanced Classic LAN) can co-exist with Cisco Nexus 7000 and 9000 switches in Enhanced Classic LAN. Use this option if you wish to refresh your older Cisco Nexus platforms with newer switches such as the Cisco Nexus 9000 switches. This can be a phased migration.

Note: To use an Enhanced Classic LAN Fabric, you must procure Cisco Nexus 9000 switches and have the cabling done for a 2 or 3 tier hierarchical network comprised of Cisco Nexus 9000 and 7000 switches, with vPCs at the Aggregation layer.

NDFC with Legacy Nexus Platforms

The following figures show topologies for use with legacy platforms:

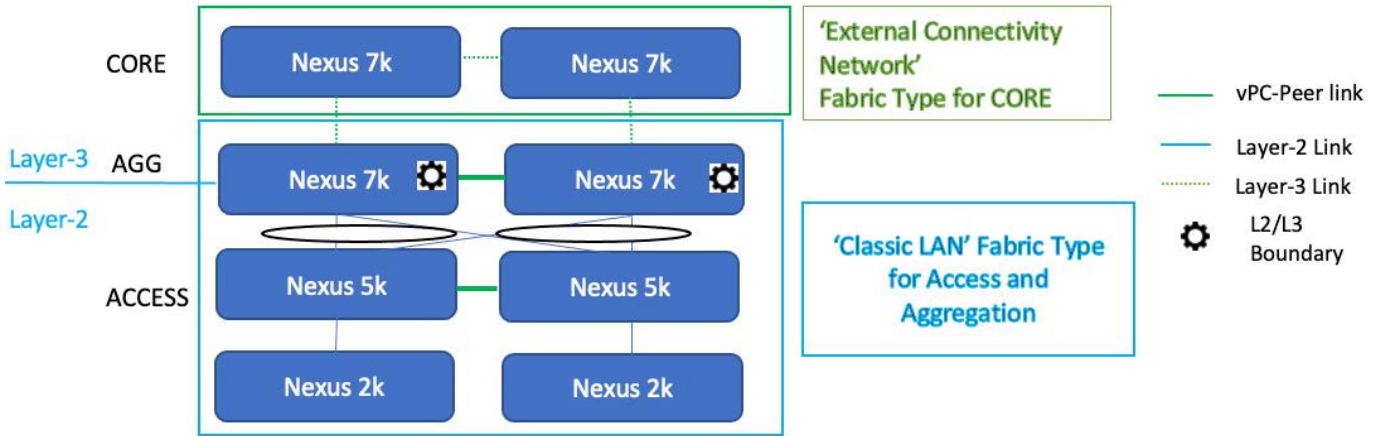


Figure 1: With a FEX attached to the Access layer

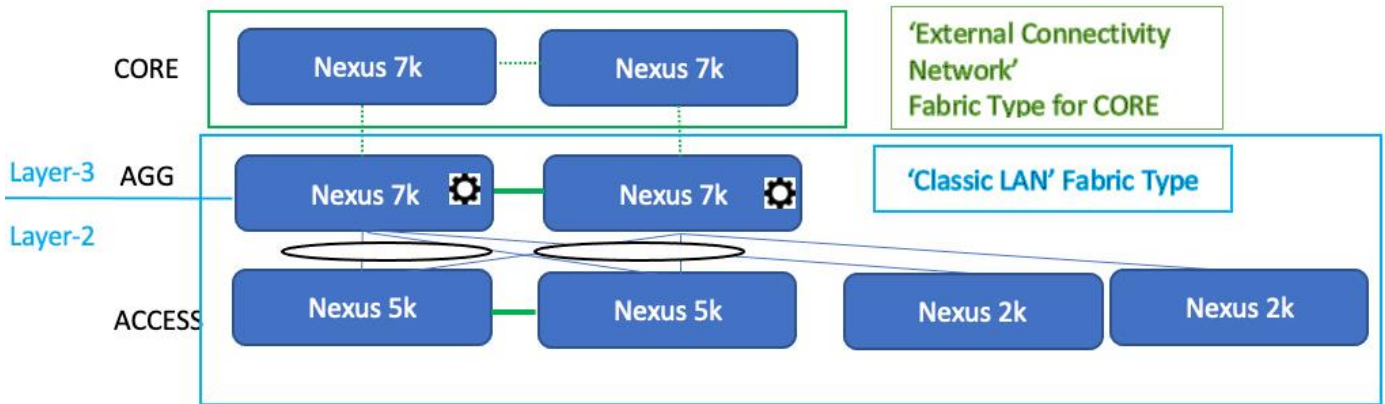


Figure 2: With a FEX attached to the Aggregation layer




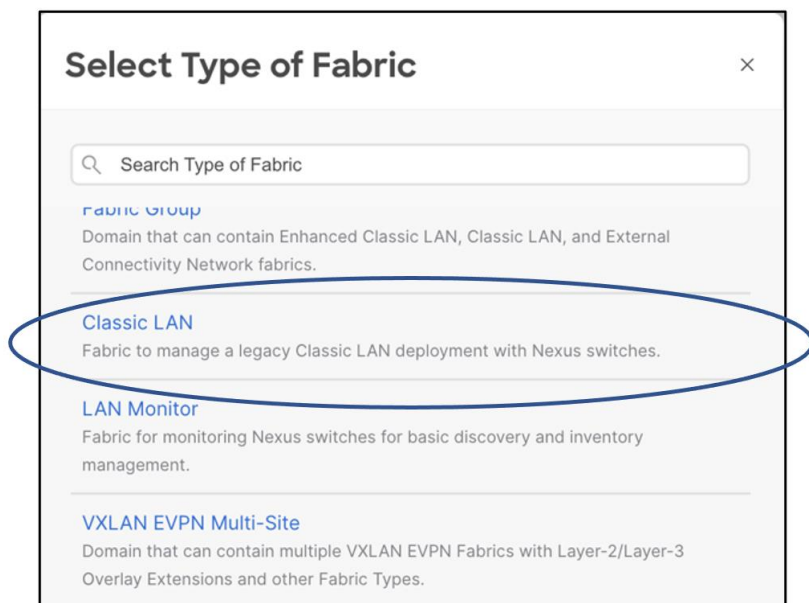
Figure 3: Without a FEX

This is the starting point where you do not have any Cisco Nexus 9000 switches. Both types of topologies--with and without FEXes--are supported. You can have several combinations, such as Cisco Nexus 9000 switches at the Core and Aggregation layers instead of having only Cisco Nexus 7000 switches as shown in the figures. These topologies use the Cisco Nexus 7000 switches with the assumption that you have not yet procured Cisco Nexus 9000 switches.

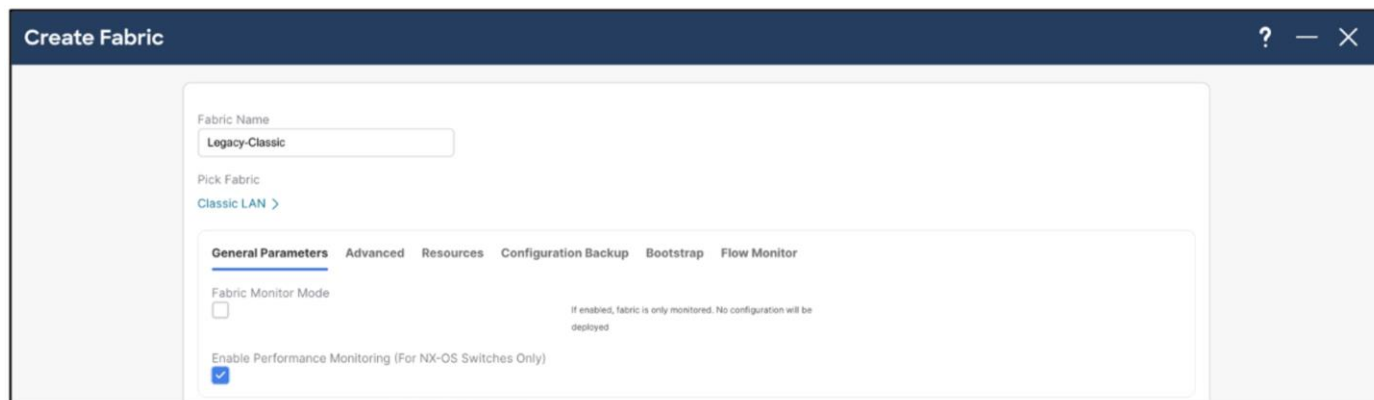
You must place Access and Aggregation switches in the "Classic LAN" fabric type, as Cisco Nexus 5000 switches are not supported in an Enhanced Classic LAN. This will be discussed in the following section.

For the Core layer with Cisco Nexus 7000 switches, follow the process in the [For the Core Layer](#) section and place the switches in the external connectivity network. After migration this fabric will be used as the Core layer.

Create the fabric with the "Classic LAN" fabric type for these topologies, as this type supports all Cisco Nexus platforms.



Make sure the **Fabric Monitor Mode** check box does not have a check.



Save the fabric after you have chosen all of the parameters.

Next, discover the switches. This step does not disturb any configurations on the switch. All configurations in the switch are kept as-is: NDFC does not add nor delete any configuration. NDFC only discovers the switches and imports them into the fabric.

Add Switches - Fabric: Legacy-Classic ? X

Switch Addition Mechanism*
 Discover

Seed Switch Details

Seed IP*

Ex: "2.2.2.20" or "10.10.10.40-60" or "2.2.2.20, 2.2.2.21"

Authentication Protocol*

Username*

Password*

Max Hops*

Because the Classic LAN fabric type does not support true brownfield, NDFC still does not have any intent that it learned from the switches while preserving all the configurations on the switch. NDFC has a feature called "Host Port Resync" that we recommend you use for switches that already have configurations. NDFC supports this feature for all Cisco NX-OS switches. Using Host Port Resync, NDFC learns all interface-related configurations and states for each switch, and these configurations become part of NDFC intent. After in the NDFC intent, you can incrementally manage the configurations from the controller using out-of-box or custom policies. Examples include vPC configurations, port channels, subinterfaces, loopbacks, routed ports, host-facing ports, VLANs, and FEXes on the Access layer, and SVIs + VLAN + HSRP configurations on the Aggregation layer.

Host Port Resynchronizing

1. Navigate to **Fabric > Policies > Actions > Add Policy**.

Overview Switches Links Interfaces Interface Groups **Policies** Networks VRFs Event Analytics History Resources Virtual Infrastructure Metrics


Filter by attributes

<input type="checkbox"/>	Policy ID	Switch	IP Address	Template	Description	Entity Name	Entity Type	Source	Priority	Content Type	Serial Num	
<input type="checkbox"/>	POLICY-362600	Access3	192.18.0.15	switch_role_si		SWITCH	SWITCH		10	PYTHON	9SHA	Add Policy
<input type="checkbox"/>	POLICY-363100	Access4	192.18.0.20	switch_role_si		SWITCH	SWITCH		10	PYTHON	9NTC	Edit Policy
<input type="checkbox"/>	POLICY-363400	Access5	192.18.0.21	switch_role_si		SWITCH	SWITCH		10	PYTHON	9FVA	Delete Policy
												Generated Config
												Push Config

2. Select the switches and deploy the template `host_port_resync`.

Switch List:

- 7K-SpineA-DDD37RU28-26
- 7K-SpineB-DDD37-RU25-23
- N5K-LeafA-DDD37RU30
- N5K-LeafB-DDD37RU29
- S1-Core1-DDD35-RU39



Pick a Template

Choose Template

Switch List:

- 7K-SpineA-DDD37RU28-26
- 7K-SpineB-DDD37-RU25-23
- N5K-LeafA-DDD37RU30
- N5K-LeafB-DDD37RU29
- S1-Core1-DDD35-RU39

Priority*

500

1-1000

Description

resync

Template Name

[host_port_resync >](#)

Interface Configuration Resync

Switch will be placed in Migration mode on clicking 'Save'.
Recalculate Config in the fabric must be performed to complete the interface configuration resync process.

After NDFC deploys the template, NDFC starts learning Interface-specific intent and starts mapping interface configurations to existing policies. After this process completes, each interface becomes associated with a policy. Clicking on a policy shows the exact configuration present on the switch that has now been imported into NDFC.

Fabric Overview - Classic-LAN

Overview Switches Links Interfaces Policies Event Analytics History Resources Virtual Infrastructure Metrics

Filter by attributes

Role	Serial Number	Config Status	Oper Status	Discovery Status	Model	VPC Role	VPC Peer	Mode	Software Versi
Spine	JPG192900BQ	NA	Healthy	Ok	N77-C7702	Primary	7K-SpineB-DDD37-RU25-23	Migration	8.4(5)
Spine	JPG1928004L	NA	Healthy			Secondary	7K-SpineA-DDD37RU28-26	Migration	8.4(5)
Leaf	FOC2025R0TQ	NA	Minor			Primary	N5K-LeafB-DDD37RU29	Migration	7.1(0)N1(1b)
Leaf	FOC2025R1C9	NA	Minor			Secondary	N5K-LeafA-DDD37RU30	Migration	7.1(0)N1(1b)
Core Router	FDO205012N9	NA	Minor	Ok	N9K-C92300YC			Migration	9.3(5)

Recalculating Config on Switches

10% - Host Port Resync - Gathering information from switches

Overview Switches Links Interfaces Policies Event Analytics History Resources Virtual Infrastructure Metrics

Filter by attributes

Interface	Admin Status	Oper. Status	Reason	Policies	Overlay Network	Sync Status
Vlan998	Up	Down	VLAN/BD is down	int_vlan	NA	In-Sync
Loopback0	Up	Up	ok	int_freeform	NA	In-Sync
Loopback1	Up	Up	ok	int_freeform	NA	In-Sync
Port-channel100	Up	Down	No operational members	int_port_channel_trunk_host	NA	In-Sync
Ethernet1/1	Up	Up	ok	int_routed_host	NA	In-Sync
Ethernet1/2	Up	Up	ok	int_routed_host	NA	In-Sync
Ethernet1/3	Up	Down	XCVR not inserted	int_routed_host	NA	In-Sync
Ethernet1/4	Up	Up	ok	int_routed_host	NA	In-Sync
Ethernet1/5	Up	Down	XCVR not inserted	int_routed_host	NA	In-Sync
Ethernet1/6	Up	Down	XCVR not inserted	int_routed_host	NA	In-Sync

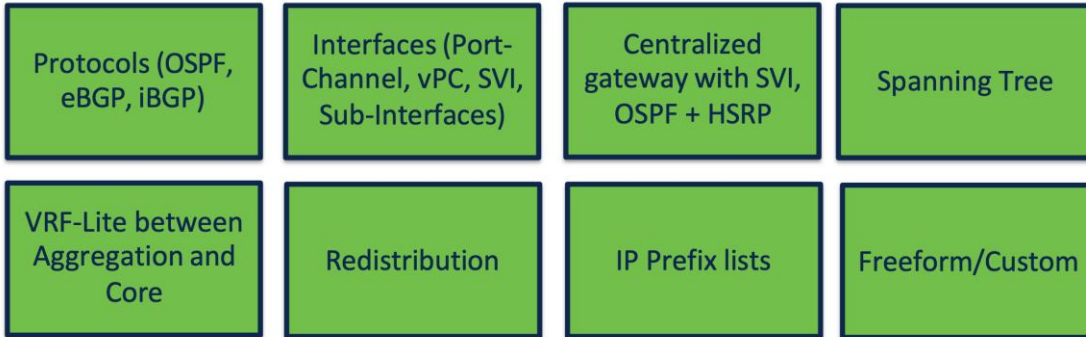
```

1 interface Ethernet1/1
2   no switchport
3   mtu 9216
4   description V,
5   service-policy type qos input QOS-BB-POLICY-IN
6   no shutdown
7

```

Hereafter, you can edit any of these Interface configurations and NDFC pushes out the changes.

For all the other global configurations, such as routing or spanning tree, because there was no support for a brownfield import, you must bring that intent manually into NDFC using out-of-box policies. There are a variety of policies available for Classic LAN use cases. NDFC has templates for all of the following user cases and more:



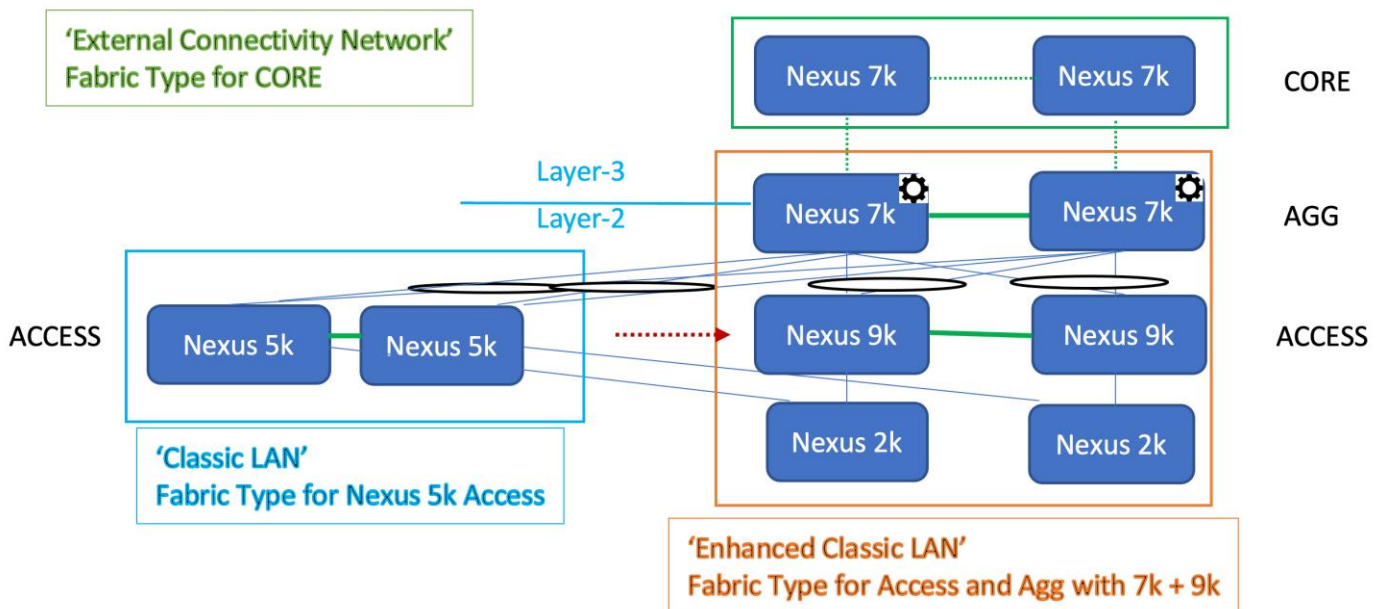
You must perform networks and VRF instance creation and the respective interface attachments using the appropriate policies.

This is how you can bring classic LAN networks with legacy Nexus platforms into NDFC and incrementally manage the networks using templates. All day 2 functionalities are available for the legacy classic LAN fabric as listed in the [Using Enhanced Classic LAN](#) section.

Note: Brownfield import plus end-to-end day 0 and day 1 automated workflows are available in the new Enhanced Classic LAN fabric type. However, the Cisco Nexus 2000, 5000, and 7000 platforms are EOL and EOS. Hence, we recommend that you refresh these platforms as soon as possible and start leveraging the benefits of this new fabric type in NDFC.

NDFC with Newer Cisco Nexus Platforms (Cisco Nexus 9000) Considering FEX

You can attach a FEX to an Access or Aggregation switch, or to both. You can transition all FEXes to Enhanced Classic LAN.



As you procure Cisco Nexus 9000 switches and you are ready to refresh the switches, you can build a best practice-based Classic Ethernet network from the ground up using the Enhanced Classic LAN fabric type with these Cisco Nexus 9000 switches at the Access layer (greenfield import) and Cisco Nexus 7000 switches at the Aggregation layer (brownfield Import). Alternatively, you can also replace the Cisco Nexus 7000 switches at the Aggregation layer with Cisco Nexus 9000 switches (greenfield import of the Cisco Nexus 9000 switches). You can also move FEXes to this new fabric and you can provision active-active/straight-through connectivity with the Cisco Nexus 9000 switches using NDFC.

You can leave the Cisco Nexus 5000 switches in the Classic LAN fabric type and manage the switches from there until they are ready to be retired. You can incrementally add any new Cisco Nexus 9000 switches to the Enhanced Classic LAN fabric in phases. NDFC provisions the configurations as long as you defined the role as discussed in the [Day 0 for Classic LAN](#) section. You can leave the Cisco Nexus 7000 Core switches in the External Connectivity fabric type to take advantage of auto VRF-Lite provisioning between the Aggregation and Core layers.

The process is summarized as follows:

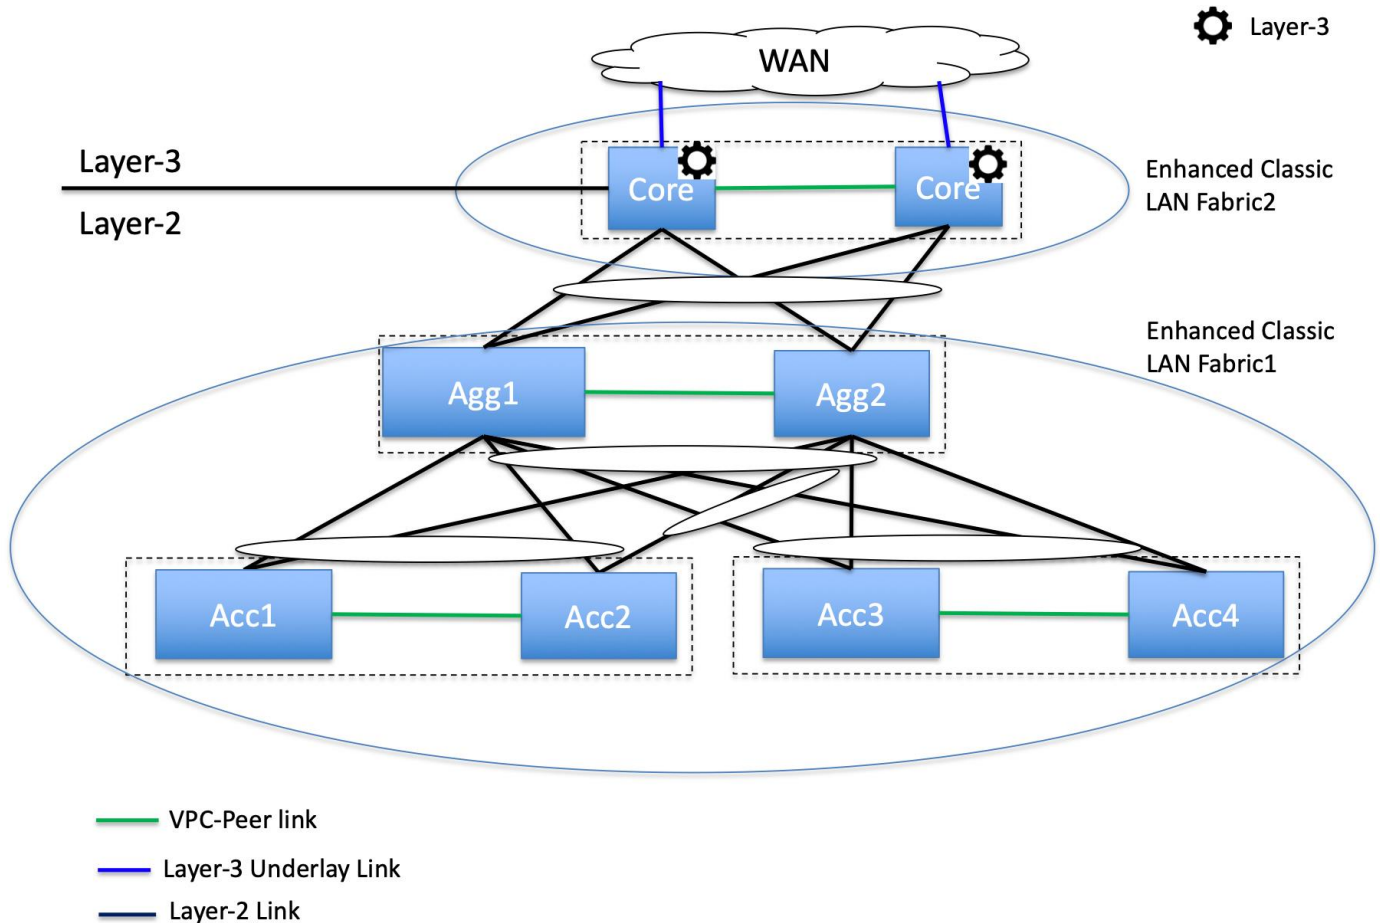
1. You can keep the Cisco Nexus 5000 switches in the Classic LAN fabric as before and incrementally manage the switches using policies until they are ready to be retired.
2. As and when you procure Cisco Nexus 9000 switches, you can discover the switches in a new fabric using the Enhanced Classic LAN fabric type. You must define the roles, which will be a greenfield import with " Preserve Config = NO" setting. One-click vPC pairing options are available. NDFC pushes all relevant configurations with respect to roles and vPCs to these Cisco Nexus 9000 switches.
3. You can move Cisco Nexus 7000 switches as vPC pairs at the Aggregation layer in the Enhanced Classic LAN fabric. Because these are existing switches with configurations already in place, you can use brownfield import with " Preserve Config = YES" to do a non-disruptive migration of the Cisco Nexus 7000 switches from the Classic LAN fabric to Enhanced Classic LAN fabric. This preserves the original connectivity of the Cisco Nexus 7000 switches with the Cisco Nexus 5000 switches even though these are in two different fabrics. Hereafter, for any new configurations between Cisco Nexus 7000 and 9000 switches, you can use the Enhanced Classic LAN workflows as described in the [Day 0 for Classic LAN](#) and [Day 1 for Classic LAN](#) sections.
4. You can move FEXes to the Enhanced Classic LAN fabric with a brownfield import. NDFC preserves all the configurations between the Cisco Nexus 5000 switches and FEXes. You can provision any new configurations between FEXes and Cisco Nexus 9000 switches from the Enhanced Classic LAN fabric. NDFC supports the active-active and straight-through options.
5. The Cisco Nexus 7000 Core switches stay in the External Connectivity fabric type. NDFC preserves all configurations. If there is no VRF-Lite between the Aggregation and Core switches, now would be a good time to configure VRF-Lite using Enhanced Classic LAN day 1 workflows to adhere by Cisco best practices.

The idea is to move all supported platforms to Enhanced Classic LAN to leverage the end-to-end automated workflows while eventually refreshing the older Cisco Nexus platforms with Cisco Nexus 9000 switches. For more information about Enhanced Classic LAN workflows, see the [Day 0 for Classic LAN](#), [Day 1 for Classic LAN](#), and [Day 2 for Classic LAN](#) sections.

Layer 2/Layer 3 Demarcation at the Core Layer

The topologies discussed in the paper so far consider Layer 2/Layer 3 demarcation at the Aggregation layer. This is per Cisco recommended best practice for a classical Ethernet network. However, if the network is designed so that the Layer 2/Layer 3 boundary is at the Core layer instead, with the Access and Aggregation layers being Layer 2, you can still use Enhanced Classic LAN to leverage the new features and automated workflows.

The following figure shows the first variation of the topology with all Cisco Nexus 9000 and Cisco Nexus 7000 at all the layers:



In this case, you can place the Access and Aggregation switches in "Enhanced Classic LAN Fabric1." You can set the options for FHRP and routing to NONE in the fabric settings and either keep everything else at the default values or customize the values. This will not provision any FHRP or routing configurations on both Access and Aggregation layers, making them all Layer 2. Set their roles appropriately as Access and Aggregation.

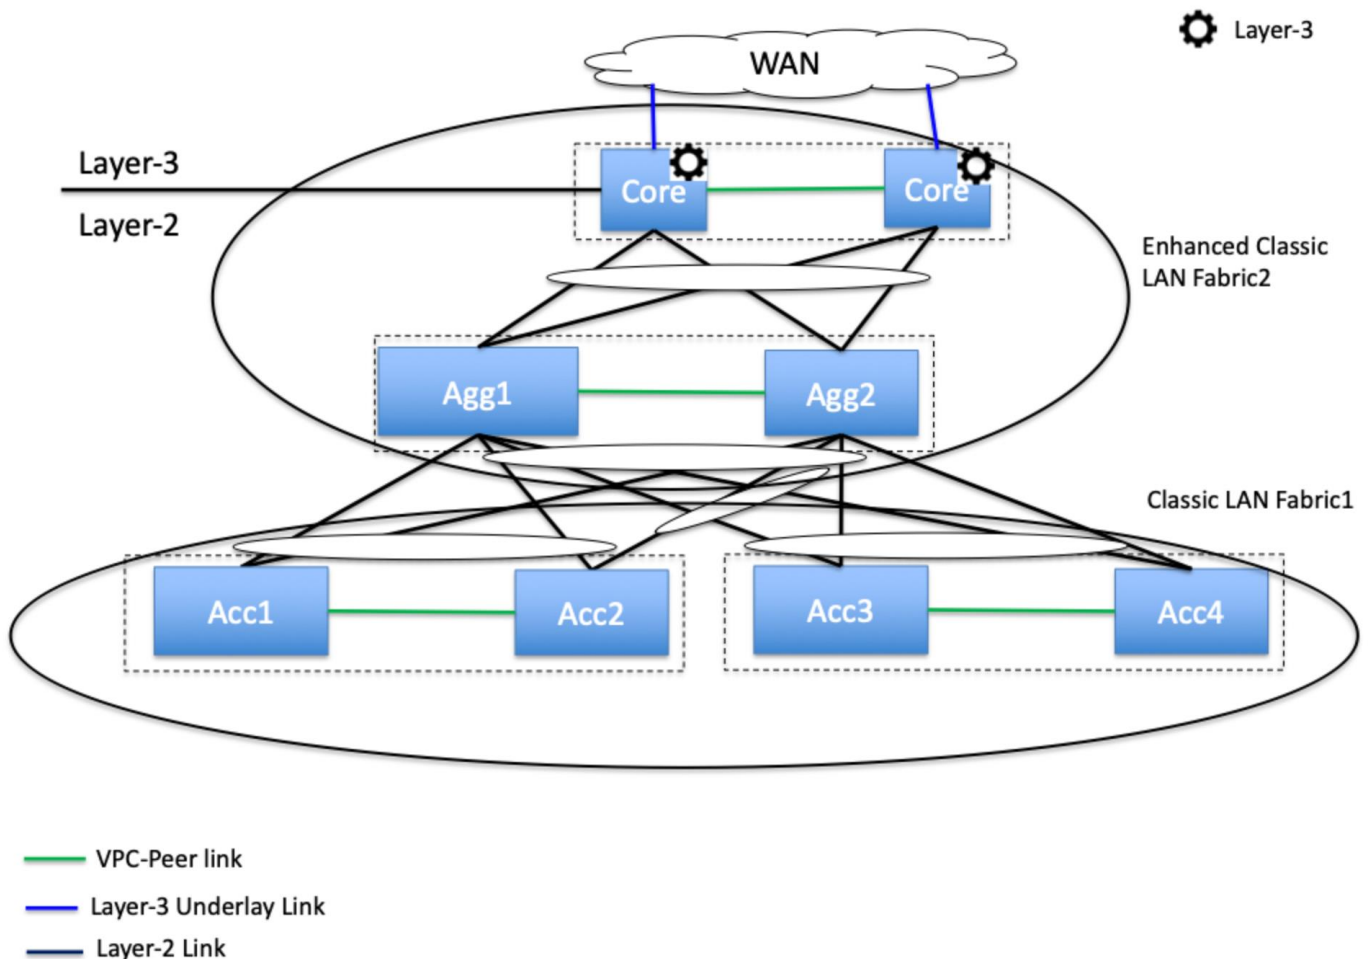
However, you must place the Core switches in "Enhanced Classic LAN Fabric2" with a role of "Aggregation." In the fabric settings, you can set the FHRP options to HSRP or VRRP. You can set the routing options based on the routing desired for external connectivity. NDFC considers this Aggregation switch as the Layer 2/Layer 3 boundary. You can place a WAN router as Core or Edge in a separate external connectivity network fabric to achieve VRF-Lite provisioning workflows with Fabric2 as discussed in the [VRF-Lite Extension Between the Aggregation and Core/Edge Layers](#) section.

A brownfield import works for both Fabrics 1 and 2. For Fabric 1, NDFC learns only Layer 2 networks, whereas for Fabric 2, NDFC learns Layer 3 Networks. You can perform any incremental provisioning using day 1 workflows from individual fabrics. You can also add both fabrics to a fabric group for better visualization.

For a greenfield, see the [Day 0 for Classic LAN](#), [Day 1 for Classic LAN](#), and [Day 2 for Classic LAN](#) sections.

Fabric 1 follows the Layer 2 network workflow. For more information, see the [Layer 2 Network](#) section. Fabric 2 follows the creation of Layer 3 network workflow with a custom or default VRF instance. For more information, see the [Layer 3 Network with a Custom VRF Instance](#) and [Layer 3 Network in Default VRF Instance](#) sections.

The following figure shows the second variation of the topology, when the Access layer consists of platforms unsupported in Enhanced Classic LAN (such as the Cisco Nexus 5000 switches), whereas the other layers consist of Cisco Nexus 9000 and 7000 switches:



In this case, you can place the Access switches in "Classic LAN Fabric1." You can incrementally manage these switches using policies until they are ready to be retired. You must create and attach networks and VRF instances, and configure vPCs and interface between Access and Aggregation switches using the appropriate policies.

You must place the Aggregation and Core switches in a separate "Enhanced Classic LAN Fabric2" with the Aggregation switch role set to "Access" and the Core switch role set to "Aggregation." In the fabric settings, you can set the FHRP options to HSRP or VRRP. You can set the routing options based on the routing desired for external connectivity. NDFC considers this Aggregation switch as the Layer 2/Layer 3 boundary. You can place a WAN router as Core or Edge in a separate external connectivity network fabric to achieve VRF-Lite provisioning workflows with Fabric2 as discussed in the [VRF-Lite Extension Between the Aggregation and Core/Edge Layers](#) section.

A brownfield import works for only Fabric 2. For Fabric 1, you can use a host port resync.

NDFC learns Layer 2 and Layer 3 networks for Fabric 2. You can perform any incremental provisioning using day 1 workflows from individual fabrics. You can also add both fabrics to a fabric group for better visualization.

For a greenfield, see the [Day 0 for Classic LAN](#), [Day 1 for Classic LAN](#), and [Day 2 for Classic LAN](#) sections.

Fabric 1 follows the creation of the Layer 2 network workflow using the policies in a Classic LAN. Fabric 2 follows the creation of a Layer 2 network at the Access layer. For more information, see the [Layer 2 Network](#) section. The Layer 3 network creation has a custom or default VRF instance at the Aggregation layer. For more information, see the [Layer 3 Network with a Custom VRF Instance](#) and [Layer 3 Network in Default VRF Instance](#) sections.

Migration from Classic LAN and VXLAN Networks

VXLAN as a technology has various benefits over traditional hierarchical Classical Ethernet networks. Along with being the industry-standard and widely adopted, VXLAN is:

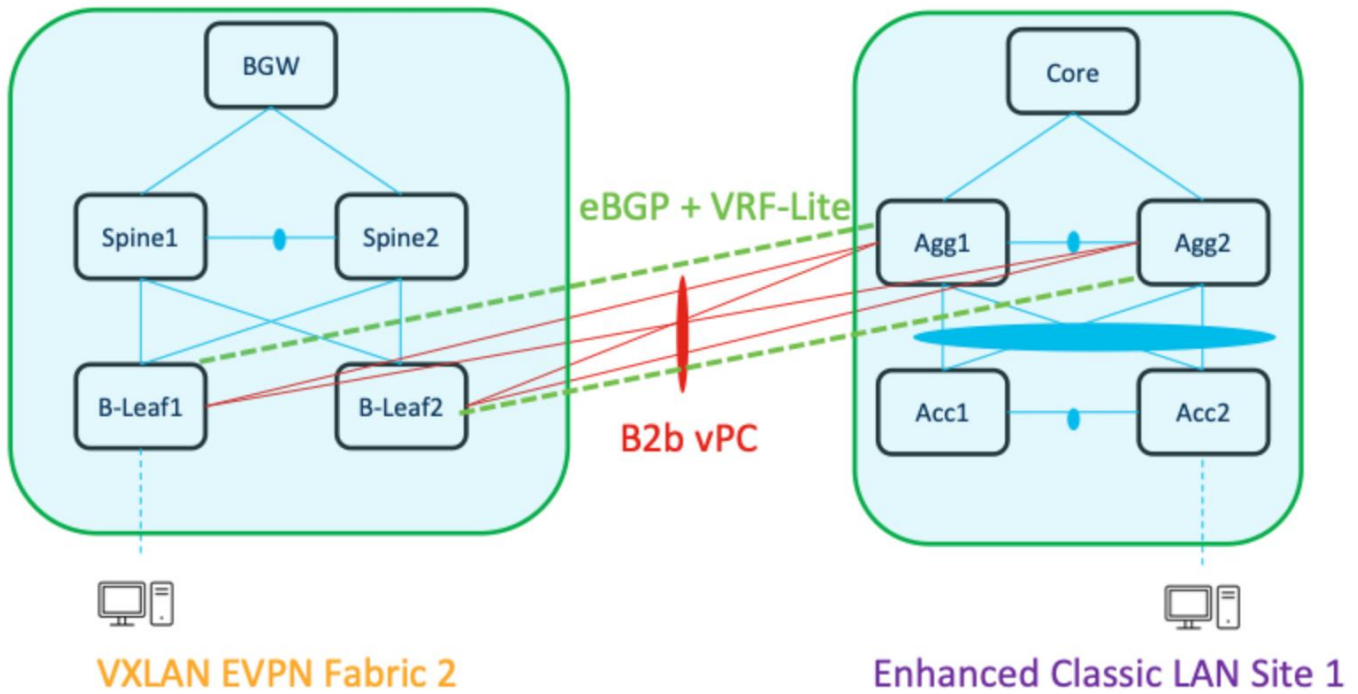
- Proven and scalable
- Improves network performance
- Increases network reliability
- Simplifies network management and IP address mobility
- Helps with segmentation and multi-tenancy

NDFC highly simplifies VXLAN greenfield (and brownfield) deployment with a few clicks, using the "Data center VXLAN EVPN" fabric type while adhering to best practices. You can find the end-to-end provisioning of a VXLAN using NDFC in the [Cisco NDFC-Fabric Controller Configuration Guide, Release 12.1.2e](#). As you plan for this adoption, it is imperative that you must be able to configure the Enhanced Classic LAN and VXLAN fabrics within the same NDFC instance until the Classic network is ready to be retired.

To plan for migration from Enhanced Classic LAN to VXLAN, you can consider the following topologies:


- Topology1: Layer 2 and Layer 3 connectivity between Aggregation switches in a brownfield Enhanced Classic LAN and border leaf switches in a greenfield VXLAN fabric.

This approach is viable if the legacy network and the new VXLAN fabric are geographically co-located.



- Topology 2: Layer 2 and Layer 3 connectivity between Aggregation switches in an Enhanced Classic LAN and a border gateway in a VXLAN fabric that extends to border gateways in a greenfield VXLAN fabric.

This approach is viable if the Legacy network and the new VXLAN Fabric are geographically dispersed.



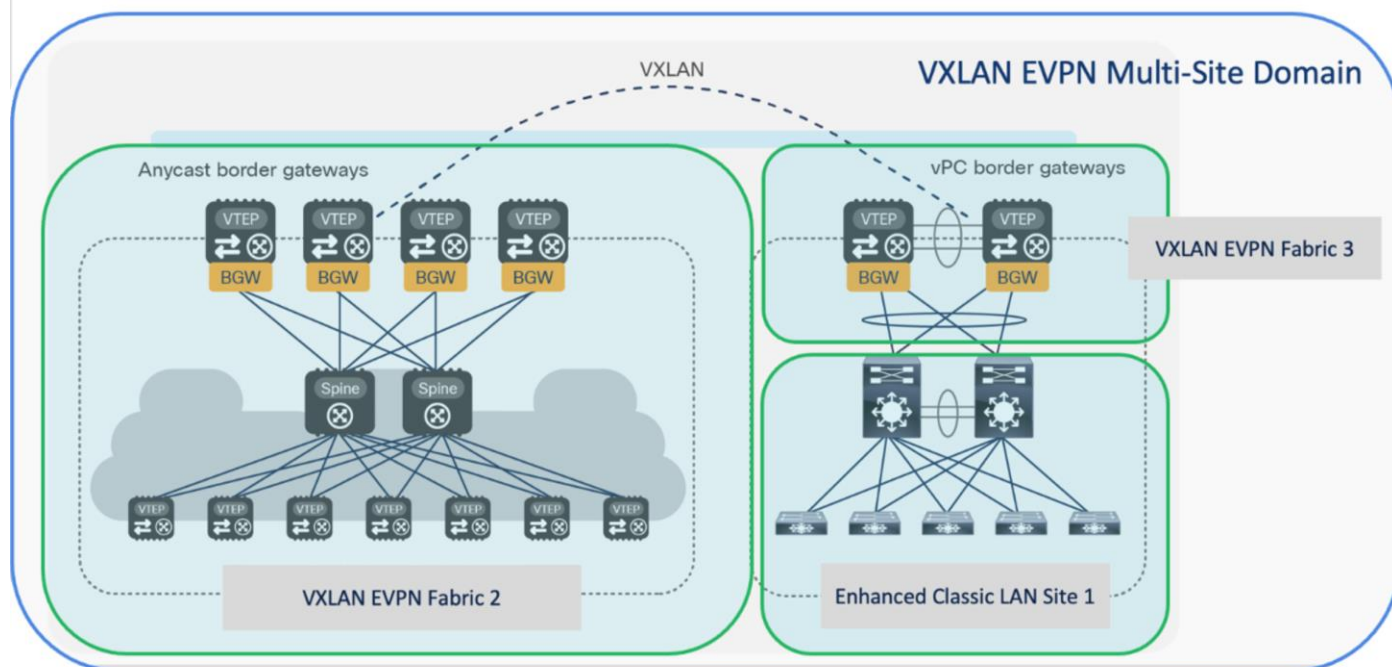
In both the topologies, you must make sure there is Layer 2 and Layer 3 connectivity between Aggregations in the brownfield site and border leaf switch or border gateway in the VXLAN site.

You must provision Layer 2 connectivity, such as using vPCs, port channels, and Layer 3 connectivity with VRF-Lite, using policies in NDFC. For information about using NDFC for provisioning Layer 2 connectivity, see:

- "[Adding Interfaces](#)" section in the *Cisco NDFC-Fabric Controller Configuration Guide, Release 12.1.2e*
- "[VRF-Lite](#)" chapter in the *Cisco NDFC-Fabric Controller Configuration Guide, Release 12.1.2e*

With NDFC release 12.1.3, you can make Enhanced Classic LAN (ECL) part of a VXLAN EVPN Multi-Site fabric along with VXLAN fabrics. This allows you to configure VXLAN EVPN Multi-Site for DCI between Enhanced Classic LAN and VXLAN fabrics. This helps with coexistence and migration to VXLAN.

The following figure shows a topology for DCI:



You can use the Data center VXLAN EVPN fabric type for a greenfield VXLAN fabric consisting of a leaf-spine-border Gateway (Fabric2) as well as border gateways (BGWs) (Fabric3) connecting the Aggregation switches (Enhanced Classic LAN Site1)—the existing brownfield Classic LAN legacy site. The border gateways allow for VXLAN EVPN Multi-Site.

Adding all 3 fabrics in a VXLAN EVPN Multi-Site domain allows for Layer 2 and Layer 3 extension of networks and VRF instances between:

- Border gateways in Fabric3 and Aggregation switches in the Enhanced Classic LAN Site1
- Border gateways in Fabric3 and border gateways in the greenfield VXLAN fabric Fabric1

You can use NDFC for:

- Importing and discovering existing brownfield classic Ethernet networks using Enhanced Classic LAN
- Setting up the greenfield VXLAN BGP EVPN fabric using POAP, bootstrap, or switch IP address discovery with border gateways
- VXLAN EVPN Multi-Site for Layer 2 and Layer 3 extension of networks between Classic and VXLAN fabrics

Enhanced Classic LAN uses the centralized gateway concept with FHRP, whereas VXLAN uses a distributed anycast gateway (DAG) concept. For coexistence of these two disparate types of gateways, that is keeping both running at the same time, Cisco NX-OS introduced a new feature starting in the 10.2(3) release as described in the "[Default Gateway Coexistence of HSRP and Anycast Gateway \(VXLAN EVPN\)](#)" chapter of the *Cisco Nexus 9000 Series NX-OS VXLAN Configuration Guide, Release 10.3(x)*. As long as the switches are running the NX-OS 10.2(3) release or later, both DAG and FHRP gateways can coexist and you can use the NDFC 12.1(3) release or later to provision the DAG and FHRP gateways. For the NX-OS 10.2(3) release and earlier, only one kind of gateway can exist. Hence, you must bring down

the gateway for workloads from the brownfield Enhanced Classic LAN network and move the gateway to the greenfield VXLAN fabric with anycast gateway.

The [Migrating Classic Ethernet Environments to VXLAN BGP EVPN](#) whitepaper provides all the details from connectivity to migration/coexistence and the configurations required, and how you can use NDFC for performing these procedures.

Hybrid Cloud Connectivity with NDFC

You can now achieve end-to-end provisioning and connectivity between on-premises data centers and cloud AWS or Azure sites. You can do this using the Hybrid Cloud Networking Solution powered by Cisco Nexus Dashboard Orchestrator (NDO) with a Cisco Nexus 9000 NX-OS-based fabric managed by Cisco Nexus Dashboard Fabric Controller (NDFC) and public cloud sites managed by Cisco Cloud Network Controller (CNC).

The prerequisites to use these solutions are to have an all Cisco Nexus 9000 switch-based VXLAN EVPN fabrics managed by NDFC. Hence, you must fully migrate from Classic LAN to VXLAN EVPN fabrics using the approach discussed in this document, with all Cisco Nexus 9000 switches.

For more information about this solution, see the [Hybrid Cloud Connectivity Deployment for Cisco NX-OS](#) document.

Conclusion

This whitepaper discusses the Enhanced Classic LAN in the NDFC 12.1.3 release, which is when this solution was first introduced. Although the Enhanced Classic LAN fabric type provides an end-to-end workflow for provisioning a 2 or 3 tier classical Ethernet network, we recommend that you migrate to a VXLAN-based overlay network due to all the benefits such a network provides over Classic Ethernet. As described in this document, if you are running networks with Cisco Nexus 5000 or 6000 switches, we highly recommend that you refresh these End-of-Life (EOL) platforms to Cisco Nexus 9000 switches to take advantage of the latest hardware and software features.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)