

Managing Legacy/Classic Networks in Cisco Nexus Dashboard Fabric Controller

First Published: 2023-08-23

Last Modified: 2024-06-12

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

© Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1 **New and Changed Information** 1

New and Changed Information 1

CHAPTER 2 **Overview** 3

Understanding NDFC for Legacy/Classic Networks 3

Support for Brownfield Deployments 9

Guidelines and Limitations 10

Recommended Platform and Software Versions 12

Prerequisites 12

CHAPTER 3 **Day 0 Configurations** 15

Configure the Enhanced Classic LAN Fabric 15

Configure the External Connectivity Network Fabric 31

Create a Fabric Group 38

CHAPTER 4 **Day 1 Configurations** 43

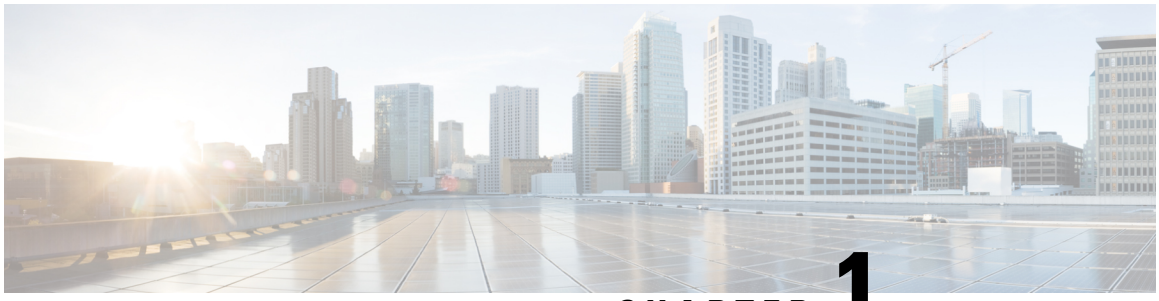
Create a Layer 2 Network 43

Create a Layer 3 Network 48

Configure the VRF Lite Extension 54

CHAPTER 5 **Day 2 Configurations** 61

Locating Information for Day 2 Configurations 61



CHAPTER 1

New and Changed Information

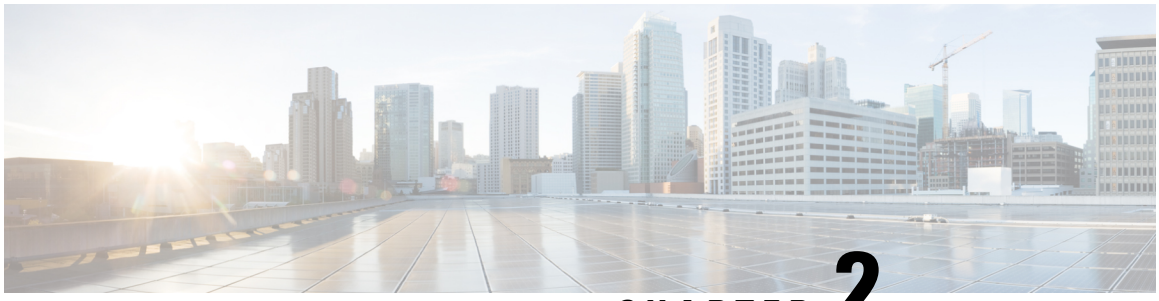
• New and Changed Information, on page 1

New and Changed Information

The following table provides an overview of the significant changes up to this current release. The table does not provide an exhaustive list of all changes or of the new features up to this release.

Release Version	Feature	Description
NDFC release 12.1.3	Initial release of this use case document.	Initial release of this use case document.

New and Changed Information



CHAPTER 2

Overview

- Understanding NDFC for Legacy/Classic Networks, on page 3
- Support for Brownfield Deployments, on page 9
- Guidelines and Limitations, on page 10
- Recommended Platform and Software Versions, on page 12
- Prerequisites, on page 12

Understanding NDFC for Legacy/Classic Networks

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

- The following sections describe the software components that are used in this process:
  - Understanding ND, on page 3
  - Understanding NDFC, on page 4
- Understanding Legacy/Classic Deployments, on page 4 describes information about standard legacy/classic deployments.
- Supported Legacy/Classic Network Topologies, on page 5 describes the two topology types that are supported when managing a legacy/classic network through NDFC.
- Understanding Access-Aggregation Device Attachments, on page 8 provides the different options for attaching the Access devices to the Aggregation devices.
- Understanding How NDFC Fabric Templates Are Used to Manage Legacy/Classic Networks, on page 9 describes how NDFC fabric templates are used to manage these types of standard legacy/classic deployments.

Understanding ND

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

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

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

Understanding NDFC

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

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

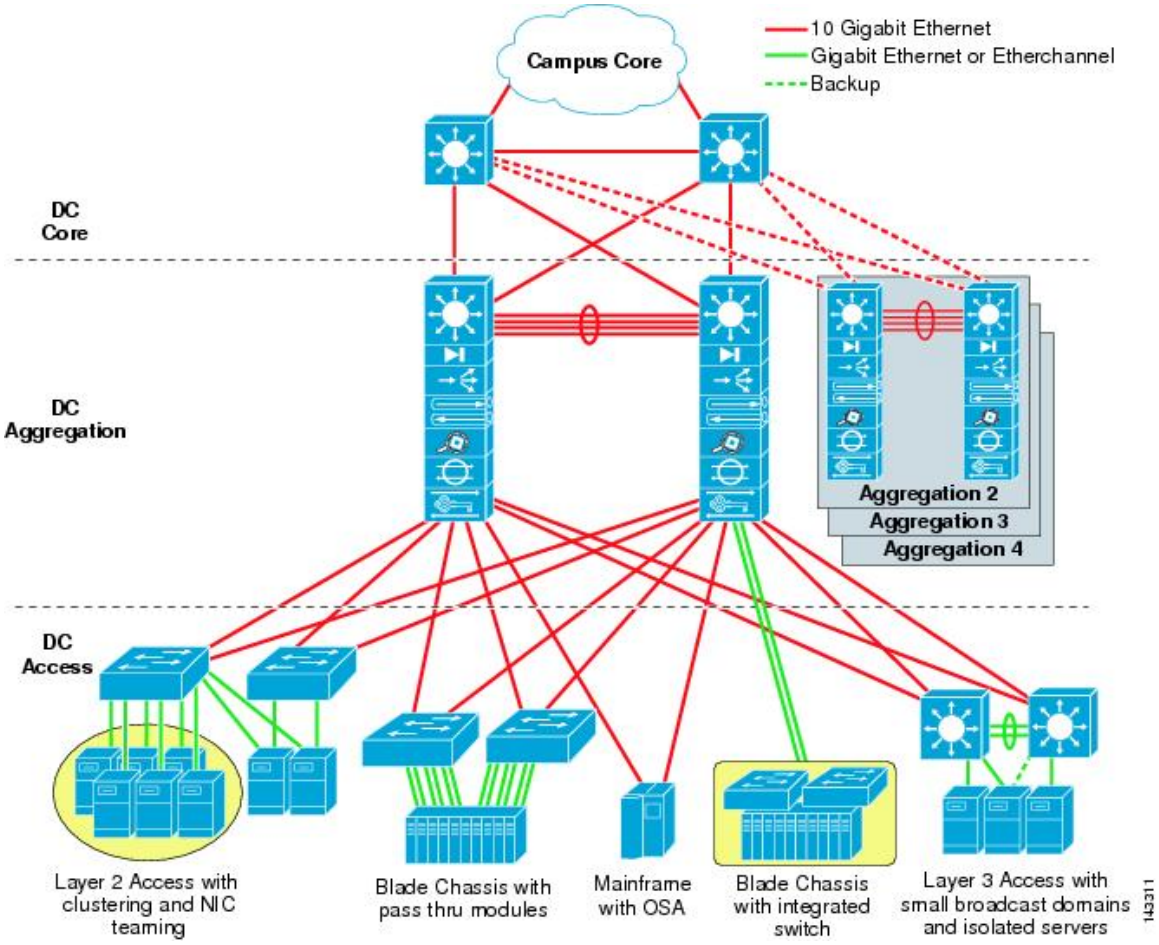
You can use the NDFC Capacity planning tool to determine the number of ND nodes required for your scale. When the ND cluster is formed and healthy, NDFC can be installed from the Cisco App store, which is directly linked to the Nexus Dashboard. On enabling the service, the cluster will intelligently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

Understanding Legacy/Classic Deployments

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

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

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



Supported Legacy/Classic Network Topologies

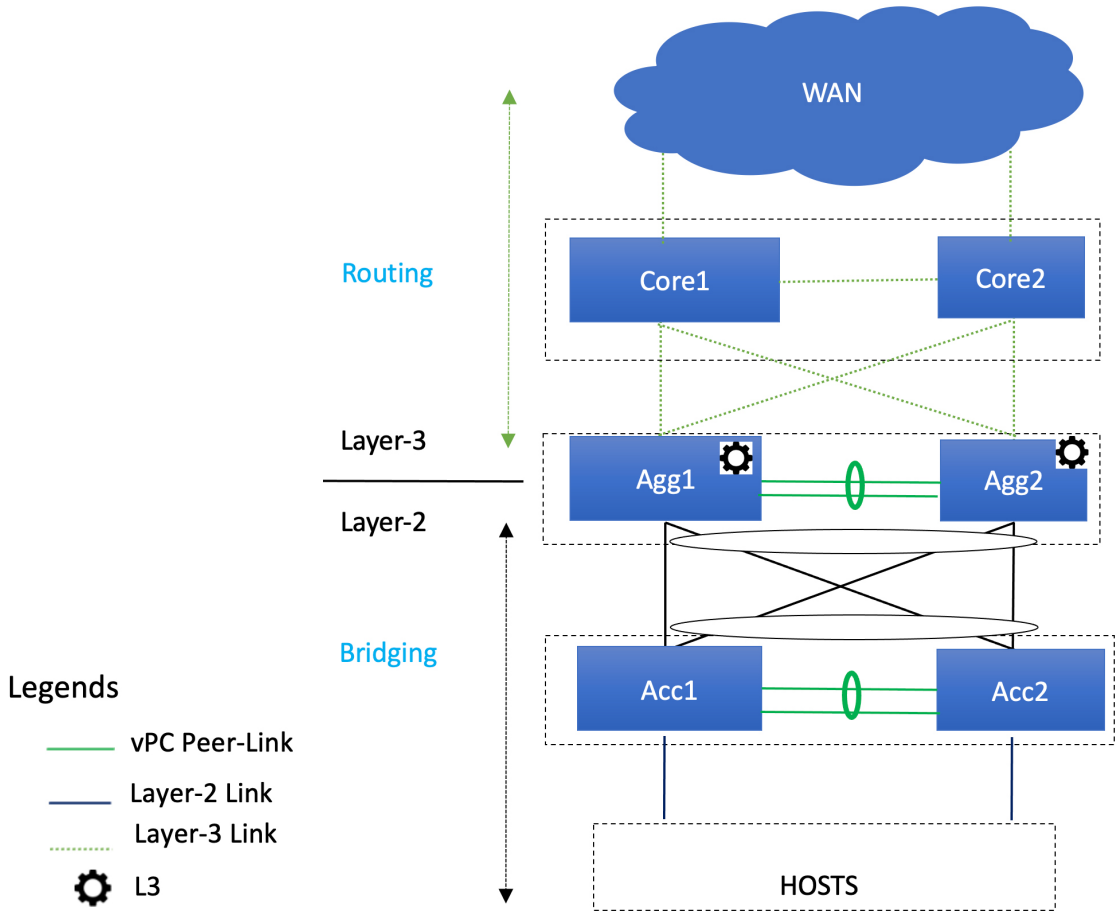
There are two types of legacy/classic network topologies that are supported, as described in the following sections.

Three-Tier Hierarchical Topology

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

- Access
- Aggregation
- Core

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



For this topology:

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

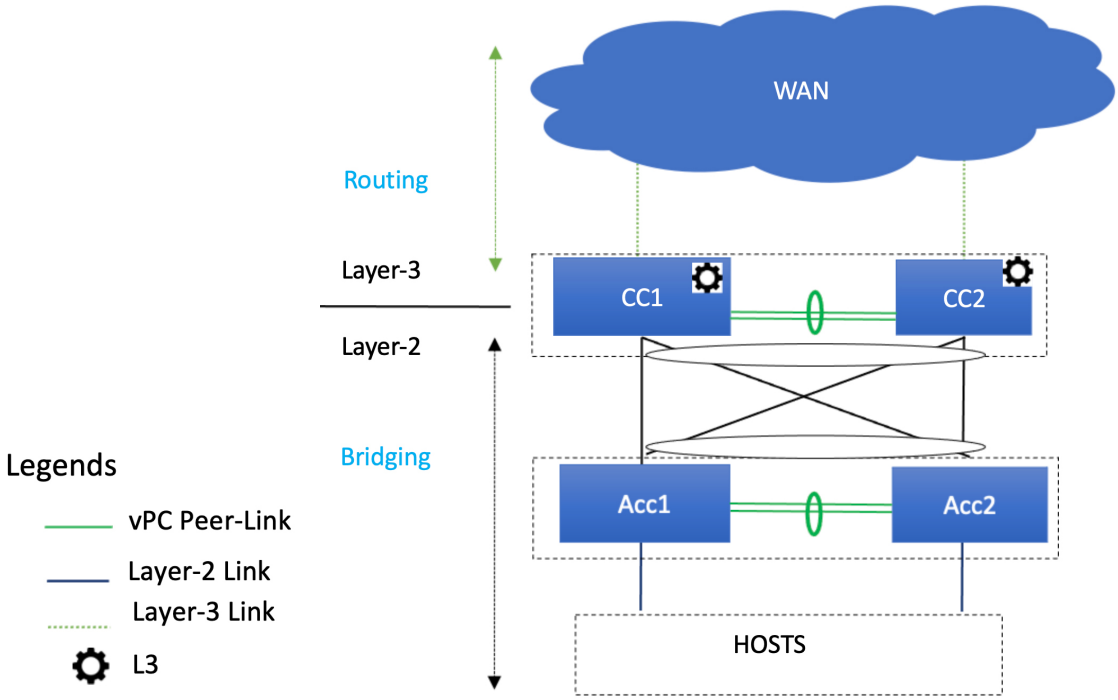
See Understanding How NDFC Fabric Templates Are Used to Manage Legacy/Classic Networks, on page 9 for more information on the two fabric template types.

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

Two-Tier Collapsed Core Topology

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

- Access
- Collapsed Core



For this topology:

- In this topology, the Core and Aggregation tiers are collapsed into a single, combined tier, called the "Collapsed Core" tier. The Layer 2/Layer 3 boundary occurs at the Collapsed Core tier.
- You will use the Enhanced Classic LAN fabric template to configure the fabric that will be used for the Collapsed Core tiers. See Understanding How NDFC Fabric Templates Are Used to Manage Legacy/Classic Networks, on page 9 for more information.

Switches Used At Each Tier

In typical legacy/classic networks using Cisco equipment:

- Cisco Nexus 7000 Series switches are deployed at the aggregation and core layers, though these switches are used mainly at the aggregation layer
- Cisco Nexus Nexus 5000 and 6000 Series switches are deployed at the access layer

- You might also have Cisco Nexus 2000 Series Fabric Extender switches attached to either the aggregation or the access layers, though these switches are typically attached to the access layer
- Cisco Nexus 3000 and 9000 Series switches might be used at the core, aggregation or access layers

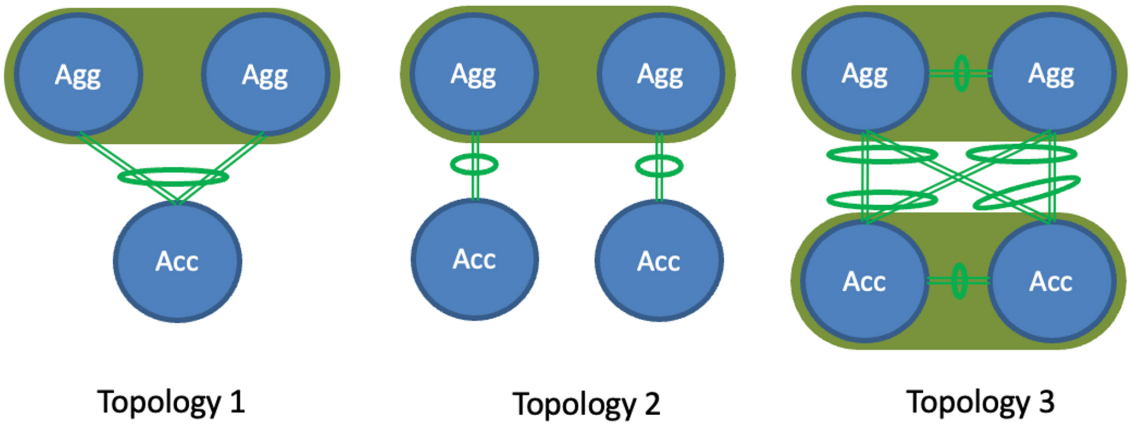
When using the new Enhanced Classic LAN fabric template in NDFC, Cisco Nexus 2000, 3000, 7000, and 9000 Series switches are supported, in these areas:

- Cisco Nexus 2000 Series Fabric Extender switches might be used at the access layer
- Cisco Nexus 7000 Series switches might be used at the aggregation layer
- Cisco Nexus 3000 and 9000 Series switches might be used at the aggregation or access layers

In addition, at the core layer, Cisco Nexus 3000, 7000, or 9000 Series switches, or possibly Catalyst 9000 or ASR 9000 Series switches, could be used because the Core role is present in the External fabric and not in the Enhanced Classic LAN fabric.

Understanding Access-Aggregation Device Attachments

For the three-tier hierarchical topology described in Supported Legacy/Classic Network Topologies, on page 5, the Access devices can be attached to the Aggregation devices using one of the following options:



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

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

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

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

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

- **Enhanced Classic LAN fabric template**: The Enhanced Classic LAN fabric template is used to configure the fabric for one of the following types of tiers, as described in Supported Legacy/Classic Network Topologies, on page 5:
  - For a three-tier topology, you will use the Enhanced Classic LAN fabric template to configure the fabric that will be used for the Access and Aggregation tiers.
  - For a two-tier (Collapsed Core) topology, you will use the Enhanced Classic LAN fabric template to configure the fabric that will be used for both tiers in that topology.
- **External Network Connectivity fabric template**: The External Network Connectivity fabric template is used to configure the fabric specifically for the Core tier in a three-tier topology, as described in Supported Legacy/Classic Network Topologies, on page 5.

Support for Brownfield Deployments

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

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

Make the following additional configurations for brownfield deployments:

- You must set the appropriate spanning tree related parameters in the fabric settings.
- Disable NX-API if it's not required, since by default these options are enabled in the fabric settings.
- Import the switches into this Enhanced Classic LAN fabric using the **Preserve Config= Yes** option.
- Set the role of the Aggregation devices because, by default, all roles will be set to Access.

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

- Aggregation devices must be configured as a vPC pair; otherwise, NDFC will return an error.

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

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

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

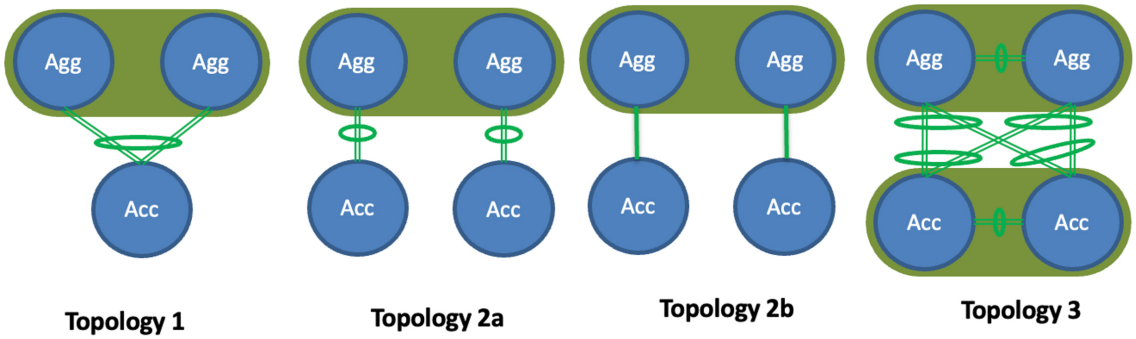
All vPC pairing related information (such as the vPC domain, the vPC peer KPA, and the vPC peer link) will be learned for the Aggregation devices and, if applicable, the Access layer devices. In addition, all interface-related configurations will be learned during the brownfield import, such as configurations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

Guidelines and Limitations

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

- The Enhanced Classic LAN fabric does not support admin VDC on Cisco Nexus 7000 Series switches. The fabric only supports default and user VDC.
- Support is available for configuring greenfield or brownfield legacy/classic networks with NDFC.
  - For brownfield legacy/classic networks, it is assumed that you have a fully-functional, running fabric before bringing that brownfield fabric into NDFC to manage.
  - A legacy/classic network configured with Cisco Fabric Path is not supported as a brownfield import into NDFC.
- Inband management and inband POAP is not supported for devices in the Enhanced Classic LAN fabric type.
- Configuring a Layer 3 port-channel as a Peer Keep Alive link is not supported. You can use any standalone Layer 3 front-port or management interface instead as a Peer Keep Alive link.
- You must have vPC configured at the following tiers:
  - At the Aggregation tier for the 3-tier topology
  - At the Collapsed Core tier for the 2-tier topology

- For greenfield deployments, the following guidelines and limitations apply:
  - You can have multiple Aggregation pairs configured at the Aggregation tier for the 3-tier topology; however, you must configure vPC for each Aggregation pair.
  - The following topologies are supported for a brownfield deployment, as described in Support for Brownfield Deployments, on page 9:



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

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

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

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

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

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

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

- The following topologies are supported:
  - 3-tier topology (Access/Aggregation/Core tiers)
  - 2-tier Collapsed Core topology, where:
    - The first tier (the Access tier) is the same as the Access tier in the 3-tier topology

- The second tier (the Collapsed Core tier) combines the Aggregation and Core tiers from the 3-tier topology into a single unified layer

Recommended Platform and Software Versions

Following are the recommended versions for the platforms and software:

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

Prerequisites

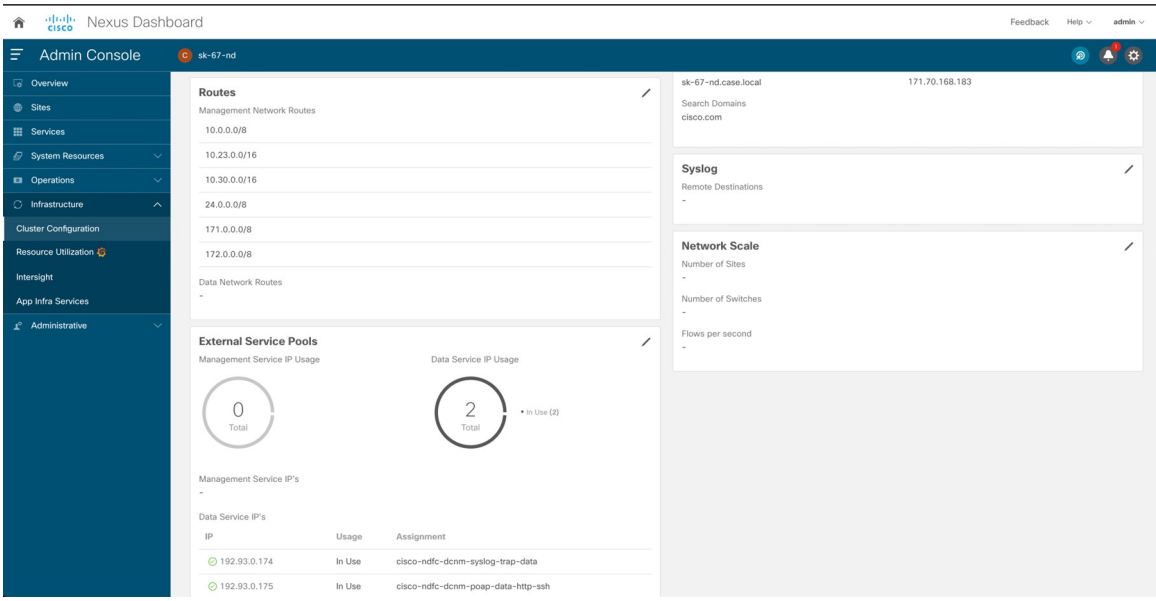
Before configuring NDFC with legacy/classic networks, complete the following prerequisite tasks to configure the Nexus Dashboard cluster and verify that the NDFC service is healthy:

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

   Deploy the Nexus Dashboard using the procedures provided in the *Cisco Nexus Dashboard Deployment Guide*, release 2.3.x or later.
2. Configure the Nexus Dashboard Fabric Controller (NDFC).

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

3. Verify the reachability between the NDFC service and the switches that you will be managing through NDFC.
   - In Nexus Dashboard, navigate to **Infrastructure** > **Cluster Configuration** to define the appropriate routes for the reachability of the switches from the ND cluster.

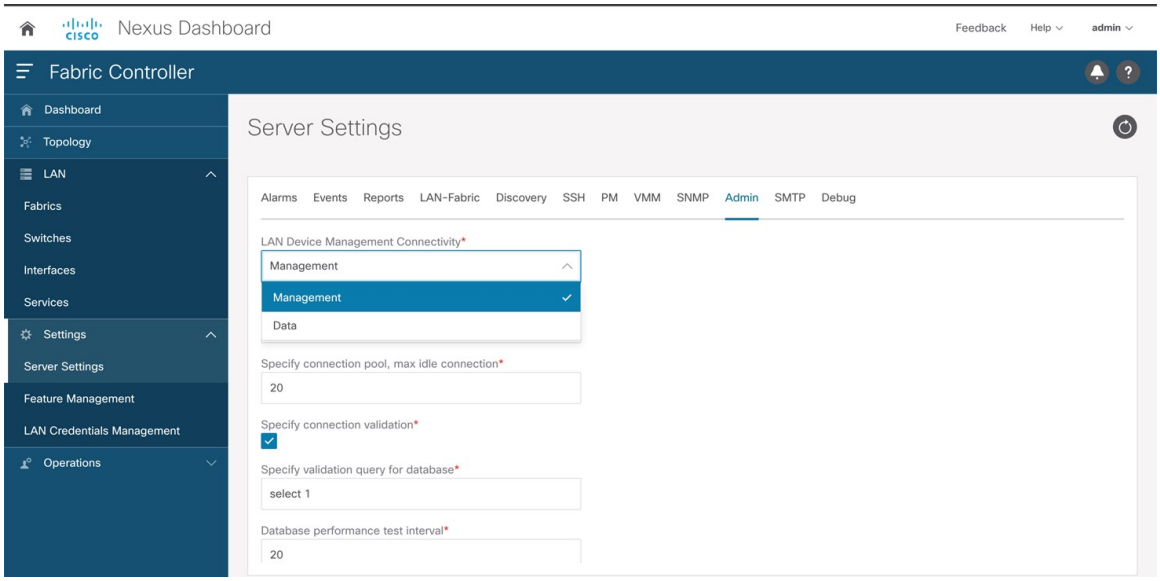


   - Define the External Service Pools for SNMP and POAP over the management or data subnet.

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

     a. In Nexus Dashboard Fabric Controller, navigate to **Settings** > **Server Settings**, then click the **Admin** tab.

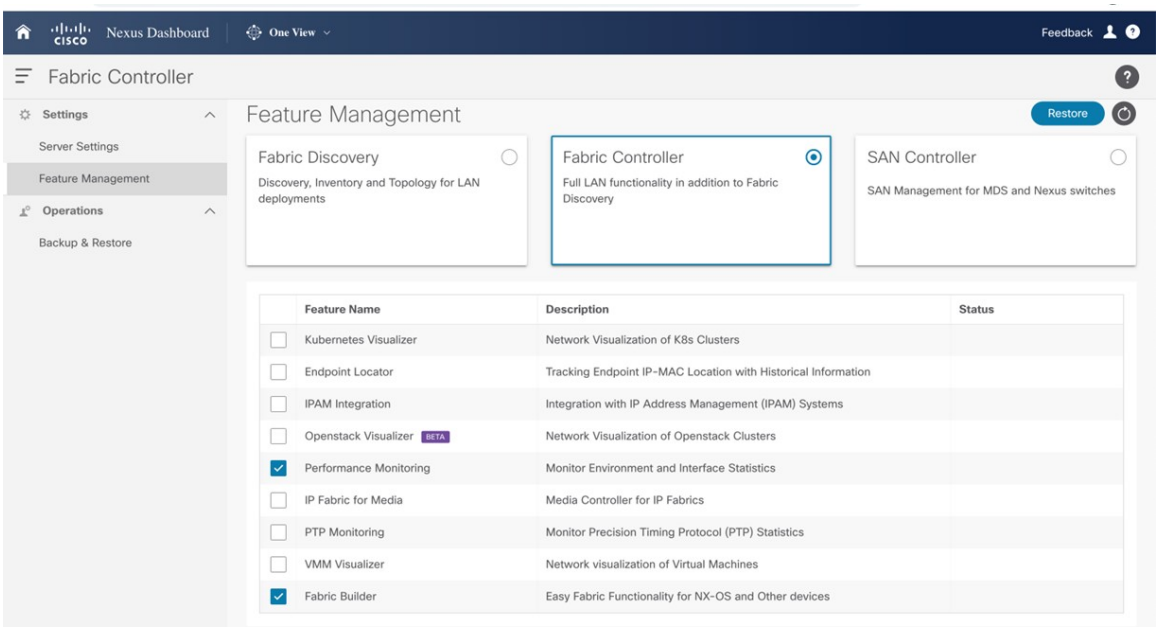
     b. In the **LAN Device Management Connectivity** field, choose **Management**.

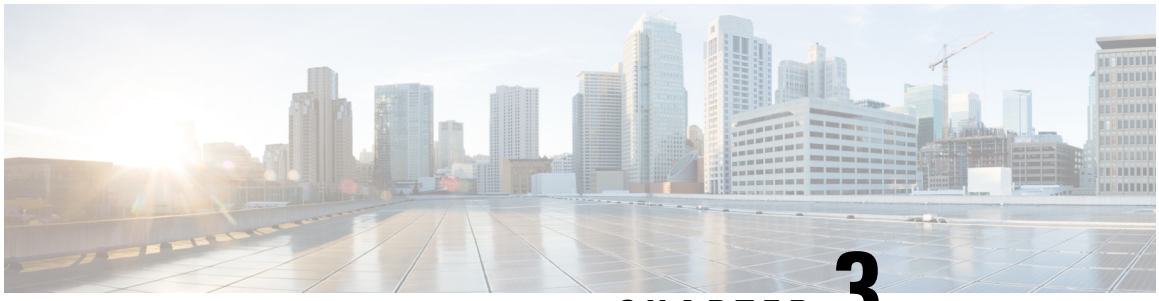


4. Configure the necessary features in NDFC.

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

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





CHAPTER 3

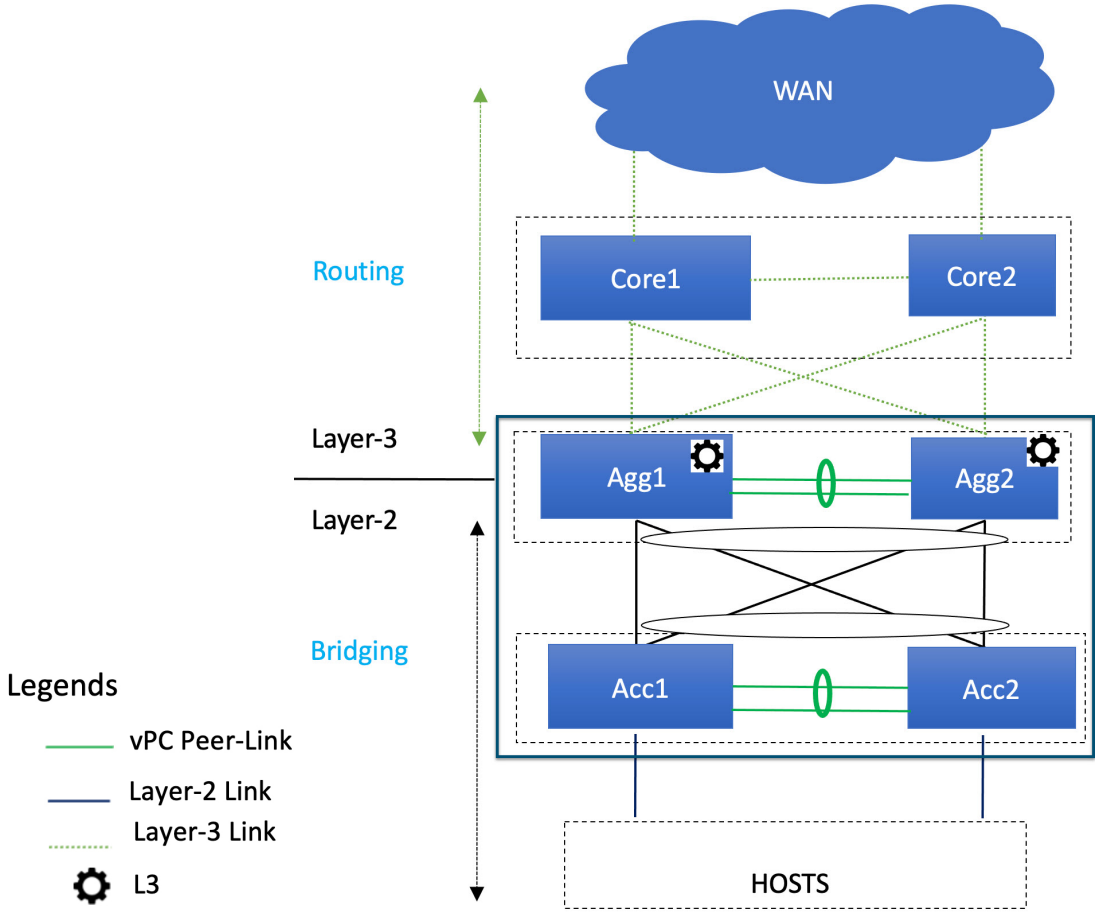
Day 0 Configurations

- Configure the Enhanced Classic LAN Fabric, on page 15
- Configure the External Connectivity Network Fabric, on page 31
- Create a Fabric Group, on page 38

Configure the Enhanced Classic LAN Fabric

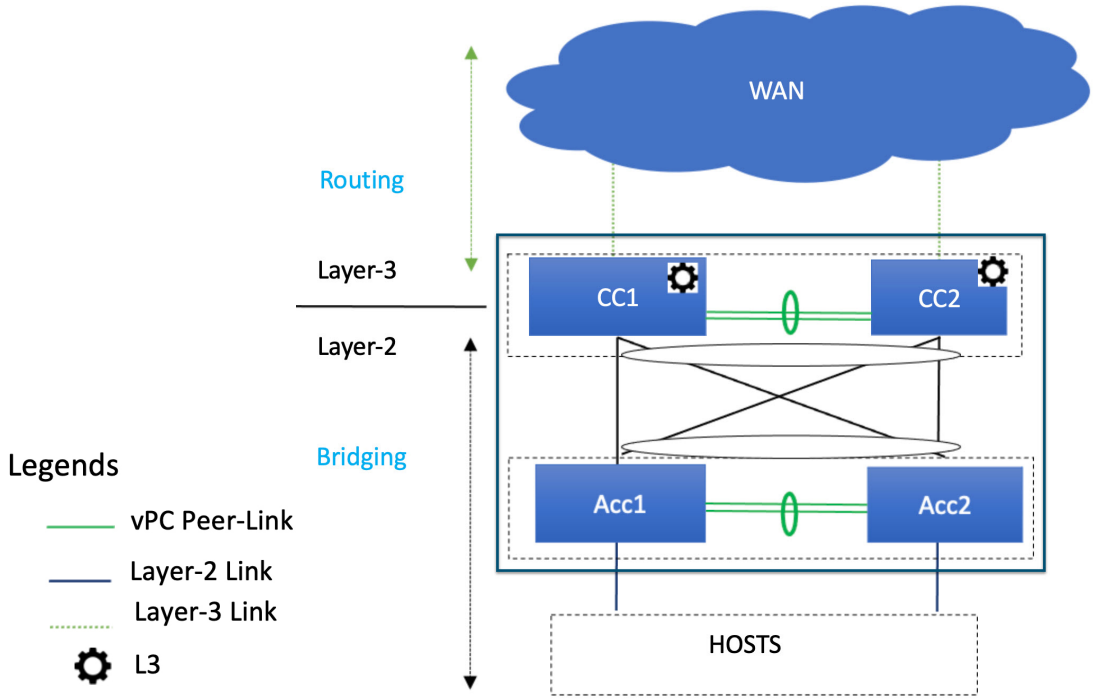
In these procedures, you will be configuring the Enhanced Classic LAN fabric that will be used to manage part of your legacy/classic network through NDFC.

- If your existing legacy/classic network falls into the three-tier topology scenario, you will be configuring a fabric that will be used for the Access and Aggregation tiers in that three-tier topology.



You will be configuring a separate External Connectivity Network fabric for the Core tier in a later section in this case.

- If your existing legacy/classic network falls into the two-tier (Collapsed Core) topology scenario, you will be configuring a fabric that will be used for both tiers in that two-tier topology.



In this case, the Core and Aggregation tiers are collapsed into a single, combined tier called the "Collapsed Core" tier.

You will be using the Enhanced Classic LAN fabric template for these configurations. For more information on that fabric template, see *Enhanced Classic LAN*.

Before you begin

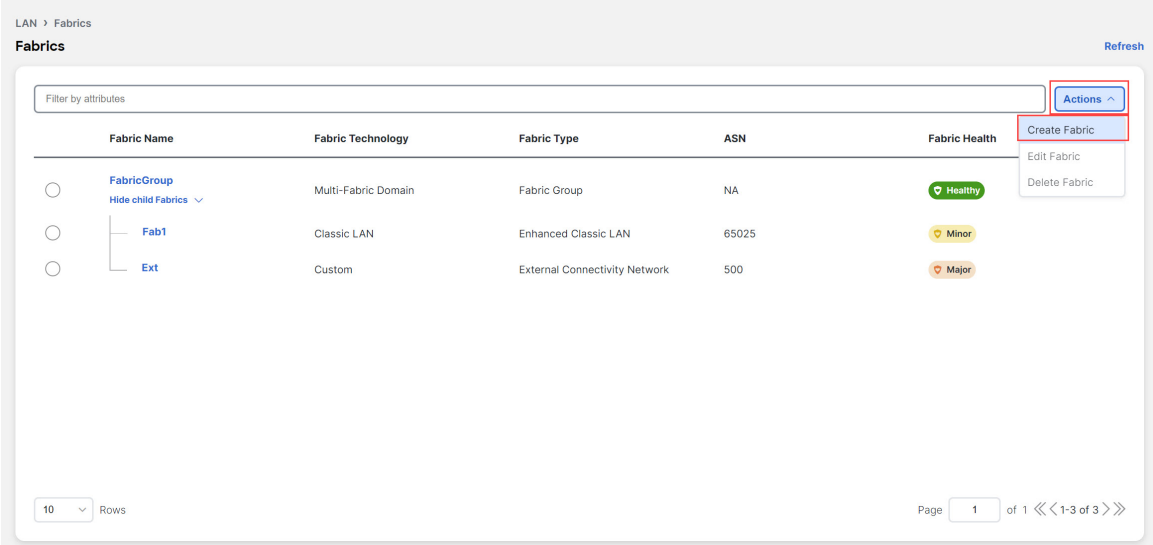
Verify that you have completed all of the necessary tasks provided in Prerequisites, on page 12 before beginning these procedures.

Procedure

Step 1 In NDFC, navigate to **LAN** > **Fabrics**.

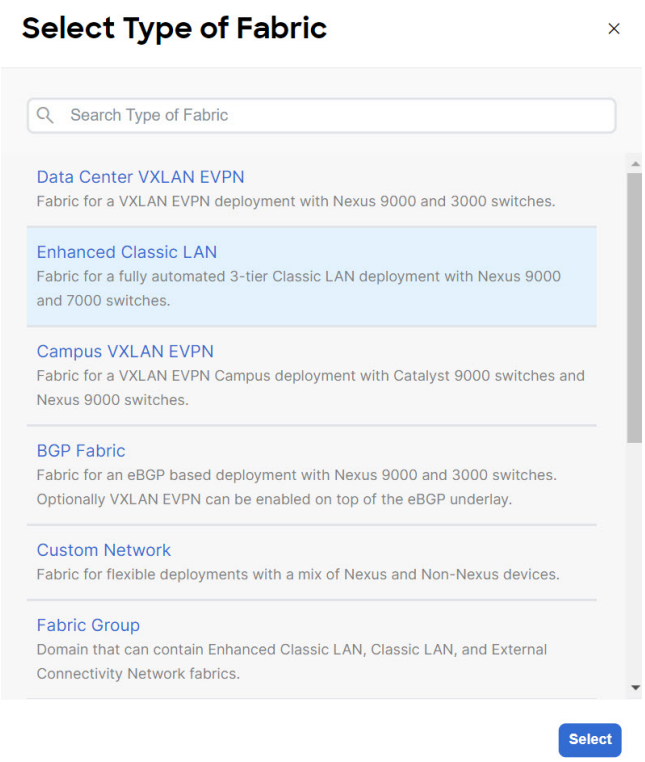
A page showing all of the configured fabrics appears.

Step 2 Click **Actions** > **Create Fabric**.



Step 3 In the **Create Fabric** screen, enter a name for the Enhanced Classic LAN fabric (for example, `Access-Agg-Fab`), then click **Choose Template**.

Step 4 Choose the `Enhanced Classic LAN` template, then click **Select**.



Step 5 In the **General Parameters** tab in the `Enhanced Classic LAN` template, make the necessary configurations.

The following configurations are mandatory in the **General Parameters** page:

- In the **Routing Protocol** field, choose the routing protocol between this Enhanced Classic LAN fabric and the External Connectivity Network fabric that will be configured in the next section, if necessary.

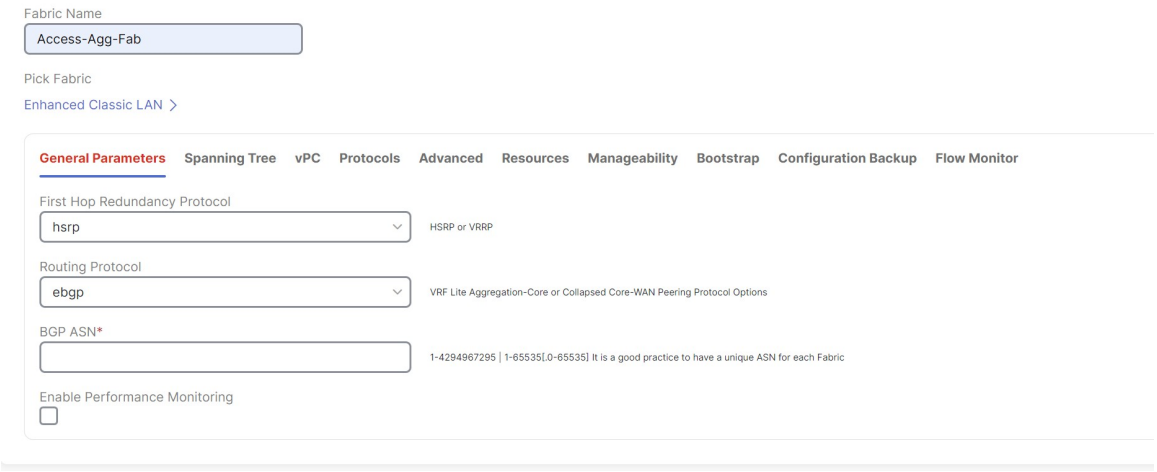
  Options are:

  - **ebgp**
  - **ospf**
  - **none**: NDFC does not configure the peering protocol if the **none** option is selected. You must manually configure the peering protocol with this option, if necessary.

- If the **Routing Protocol** is ebgp, in the **BGP ASN** field, enter a unique, non-overlapping value.

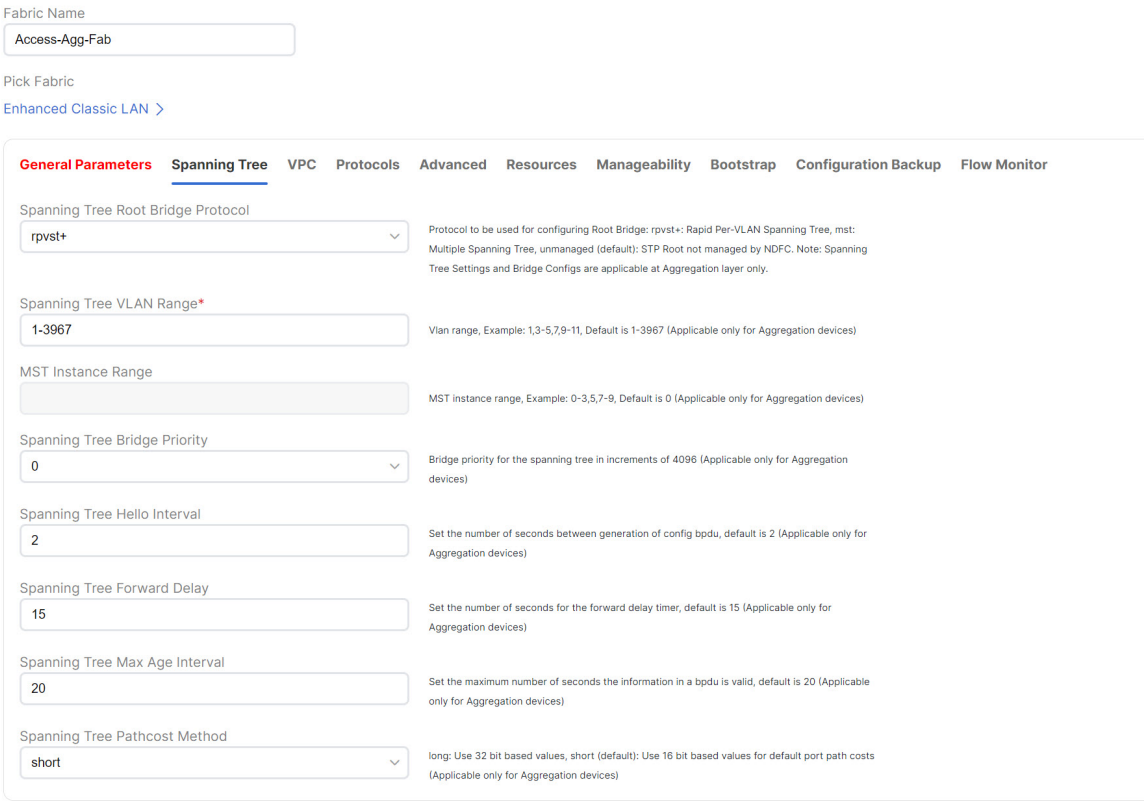
  For this use case, we will use `65535` as the BGP ASN number that will be assigned to the Enhanced Classic LAN fabric.

The remaining fields are optional. They have been automatically configured based on Cisco best practices, but you can change them if necessary. For example, in the **General Parameters** page.



Step 6 In the **Spanning Tree** tab, make the necessary configurations for your setup.

For example, in the **Spanning-tree Root Bridge Protocol** field, the `rpvst+` (Rapid Per-VLAN Spanning Tree) option is selected by default, but the `mst` (Multiple Spanning Tree) and `unmanaged` (STP root not managed by NDFC) options are also acceptable.



Step 7 In the **Bootstrap** tab, determine how you want NDFC to discover the switches in the Enhanced Classic LAN fabric.

You can use either of these two methods to discover the switches in the fabric:

- By manually entering the necessary information to allow NDFC to discover those switches. This option is applicable if you have already configured certain parameters, such as the out-of-band management IP addresses, on the switches that need to be discovered.

  If you decide to use this method, that step is provided later in these procedures (Step 10, on page 23).

- By using the Power On Auto Provisioning (POAP) feature in NDFC. This option is useful if you do not already have certain parameters, such as the management IP address, default route, and start up configurations, already configured on the switches that need to be discovered.

  POAP automates the process of installing configuration files on devices that are deployed on the network for the first time. POAP allows devices to be brought up without performing any manual configuration. When a POAP feature-enabled device boots and does not find the startup configuration, the device enters POAP mode, locates a DHCP server, and bootstraps itself with its interface IP address, gateway, and DNS server IP addresses. The device obtains the IP address of a TFTP server and downloads a configuration script that enables the switch to download and install the appropriate software image and configuration file.

  Note the following:

  - Only out-of-band POAP is supported for switches in the Enhanced Classic LAN Fabric type.
  - NDFC can be configured as the local DHCP server, handing out the IP addresses when requested by the switch while in the POAP phase. Once the switch fetches the IP address, a default route for the reachability

and desired startup configuration (optionally an image used to boot the switch) is pushed to the switch. Alternatively, an external DHCP server is also supported.

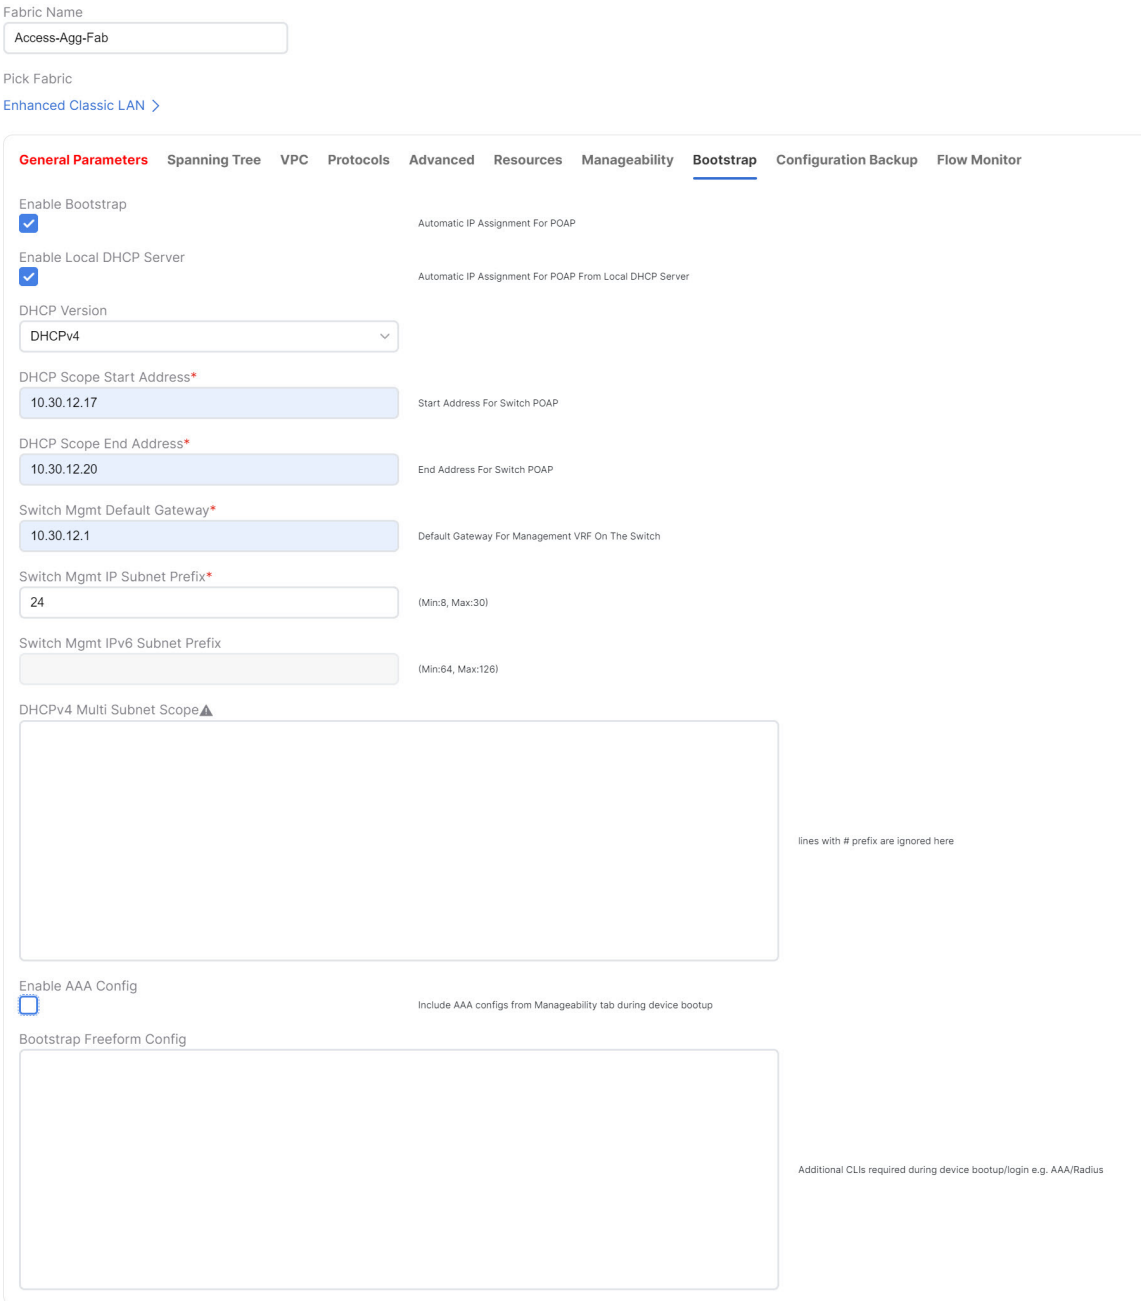
If you decide to use this method, follow these procedures to configure the POAP feature:

a. Click the **Bootstrap** tab in this page to make the necessary configurations in the **Bootstrap** area.

b. Check the box in the **Enable Bootstrap** field.

c. (Optional) Check the box in the **Enable Local DHCP Server** field (NDFC as DHCP server).

d. Define the subnet scope and default gateway that will be sent to the switch as soon as the switch is pre-provisioned while in the POAP loop.



Step 8 Make any additional configurations to the template for the Enhanced Classic LAN fabric in the remaining tabs, if necessary.

Updates to any of the values in the following tabs are all optional.

- **VPC**: Modify vPC default values.
- **Advanced**: Modify entries for AAA, NXAPI, or templates to be used for sub-operations and CoPP profile, as well as group freeform configurations for Access and/or Aggregation switches.
- **Resources**: Modify entries for the default IP and subnet ranges.

- **Manageability**: Modify entries for DNS, NTP, and syslog server settings.
- **Configuration Backup**: Modify entries to define the cadence of automatic fabric level backups.
- **Flow Monitor**: Modify entries to enable Netflow.

For more information on these fields in the Enhanced Classic LAN fabric template, see *xref2:Enhanced Classic LAN KB article*.

Step 9 When you have completed the necessary configurations to the Enhanced Classic LAN fabric template, click **Save**.

The **LAN Fabric** page appears again, with the newly created Enhanced Classic LAN fabric added to the list of configured fabrics.

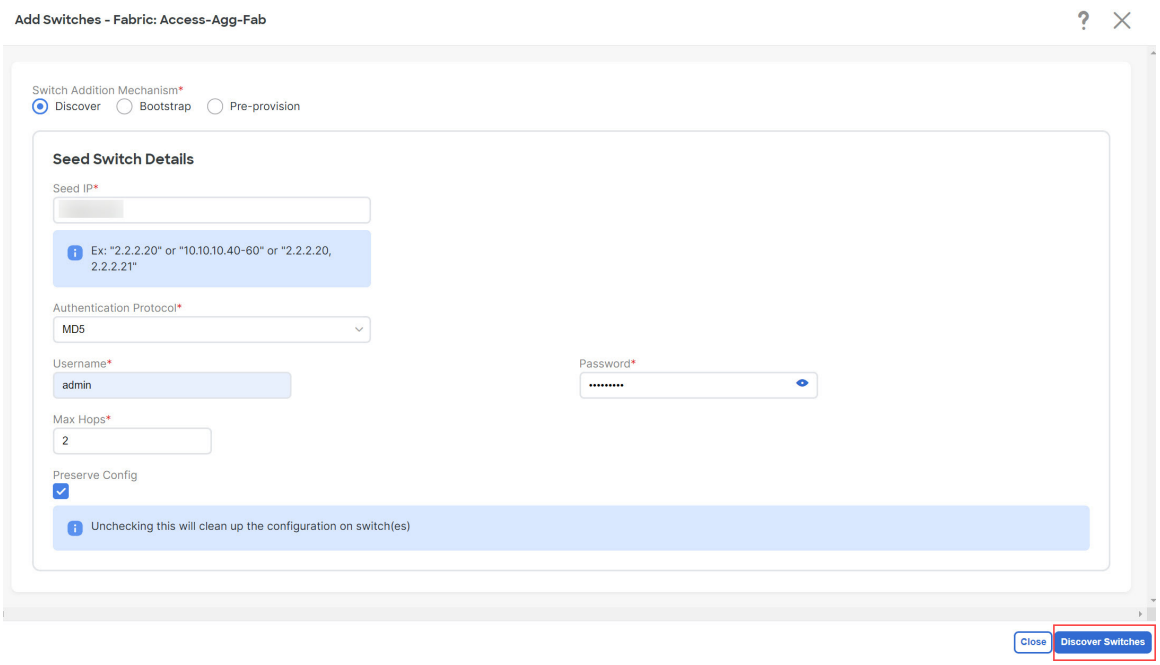
Step 10 If necessary, enter the necessary information to allow NDFC to discover the switches in the Enhanced Classic LAN fabric.

Note
You do not have to go through the procedures in this step if you **Enabled Bootstrap** through the **Bootstrap** area in NDFC in the previous step.

Verify that reachability exists between NDFC and these switches before proceeding with this step.

a) In the **Overview** page for the Enhanced Classic LAN fabric, click the **Switches** tab, then click **Action** > **Add Switches**.
b) In the **Add Switches** screen, verify that the **Discover** option is selected, then add the necessary information to discover the switches.
   - In the **Seed IP** field, enter the management IP addresses of the switches. Only out-of-band management of switches is supported.
   - In the **Preserve Config** field, make the appropriate selection.

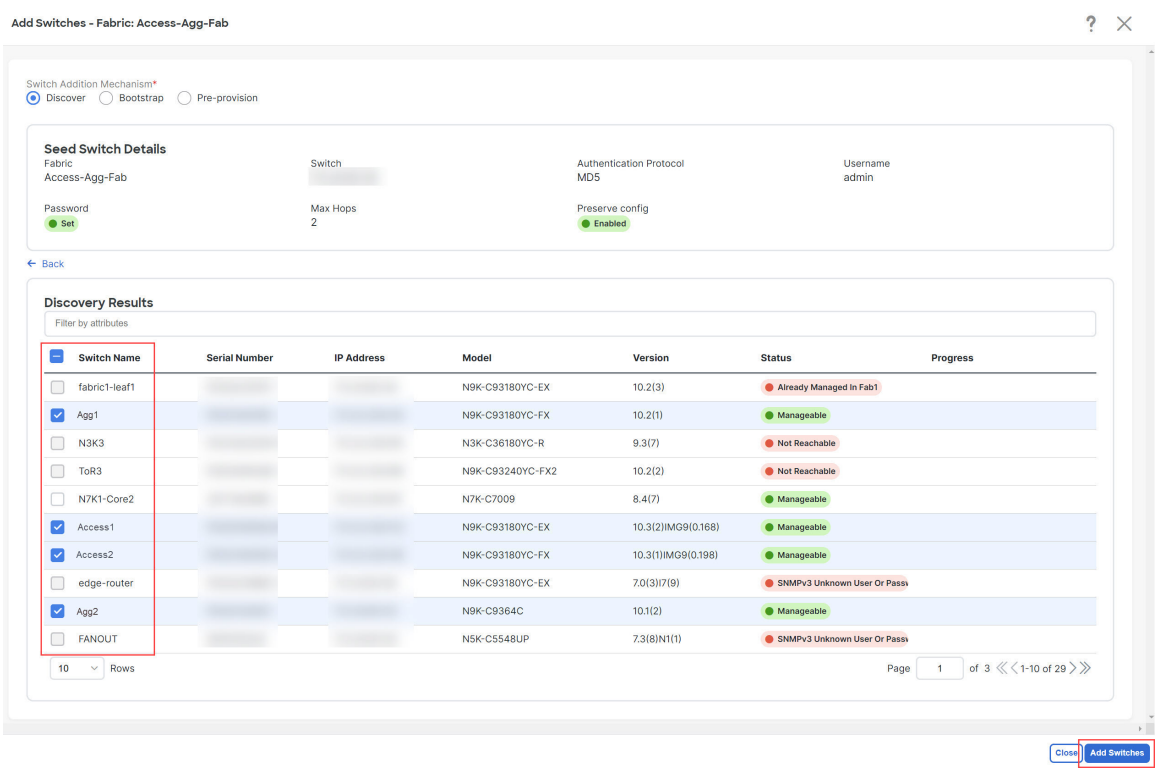
     If you remove the check from the box in the **Preserve Config** field, all existing configurations except the management IP address, the default gateway, and the boot variables will be erased so that a fresh configuration will be pushed out from NDFC.
     - If you are bringing in a greenfield deployment for NDFC to manage, remove the check from the box so that any existing configurations are not preserved on the switches.
     - If you are bringing in a brownfield deployment for NDFC to manage, check the box if you want existing configurations on the switches preserved; otherwise, remove the check from the box if you do not want the existing configurations on the switches preserved.
   - Enter any remaining information in the **Add Switches** screen that is necessary to discover the switches.



c) Click **Discover Switches**.

Click **Confirm** in the confirmation popup window that appears.

d) In the **Discovery Results** screen, check the check box next to the switches that will be imported into the Enhanced Classic LAN fabric and click **Add Switches**.



The status will change to **Switch Added** when the process is complete. Click **Close** to close out of this window.

Step 11 Define the roles for the switches in the Enhanced Classic LAN fabric.

Once the switches are discovered in the Enhanced Classic LAN fabric, the next step is to define the roles, or the intent, for those switches. Based on the roles that you assign to the switches, the appropriate configuration will be generated and pushed to the switches by NDFC.

The following roles are available for the switches in the Enhanced Classic LAN fabric in this step:

- **Access role**: Normally, you would manually assign this role to the switches in the Access tier in either the three-tier hierarchical network topology or the two-tier Collapsed Core topology.

  However, by default, all Nexus 9000 switches that are discovered in a fabric that uses the `Enhanced Classic LAN` template are automatically assigned an **Access** role. This is because the majority of switches that are used in this type of fabric are normally used at the Access tier, so you won't have to manually assign the **Access** role to those particular Nexus 9000 switches in this case.

- **Aggregation role**: You will assign this role to the following switches:
  - For the three-tier hierarchical network topology, you will assign the **Aggregation** role to the switches in the Aggregation tier as the Layer 2/Layer 3 demarcation. Aggregation also acts as the Spanning Tree Bridge and a gateway with the relevant FHRP configurations.
  - For the two-tier Collapsed Core topology, you will assign the **Aggregation** role to the switches in the Collapsed Core tier, where the Core and Aggregation levels are unified on the same switch. These switches serve as a Layer 2/Layer 3 demarcation, a bridge, and a gateway, and will also connect to the WAN (optionally using VRF-Lite, which is fully supported in the Aggregation layer).

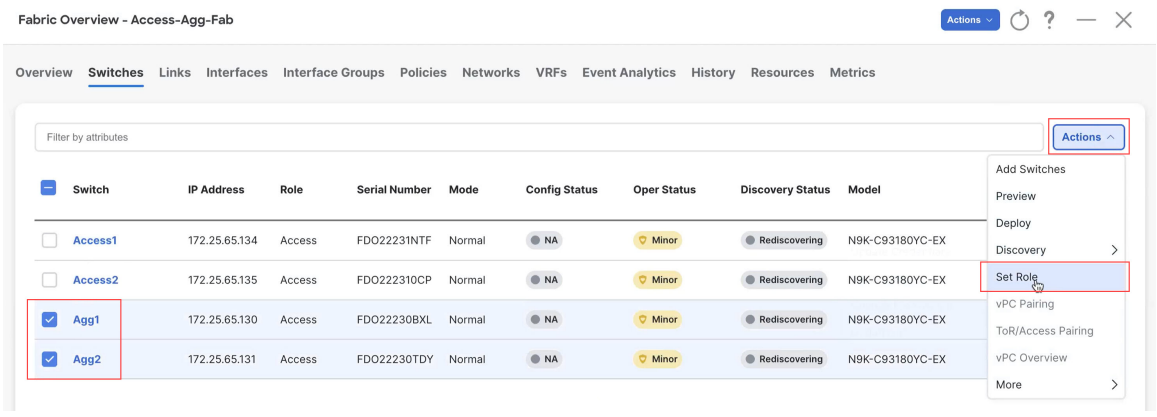
a) After the discovery process is finished for the switches, navigate back to the **Switches** tab, if necessary.

b) Verify that the roles for all of the Nexus 9000 switches discovered in this fabric are automatically set to **Access**.

As described previously, all Nexus 9000 switches in this type of fabric are automatically assigned an **Access** role so that you don't have to manually assign this role to all of the switches in the Access tier of a three-tier or two-tier Collapsed Core topology.

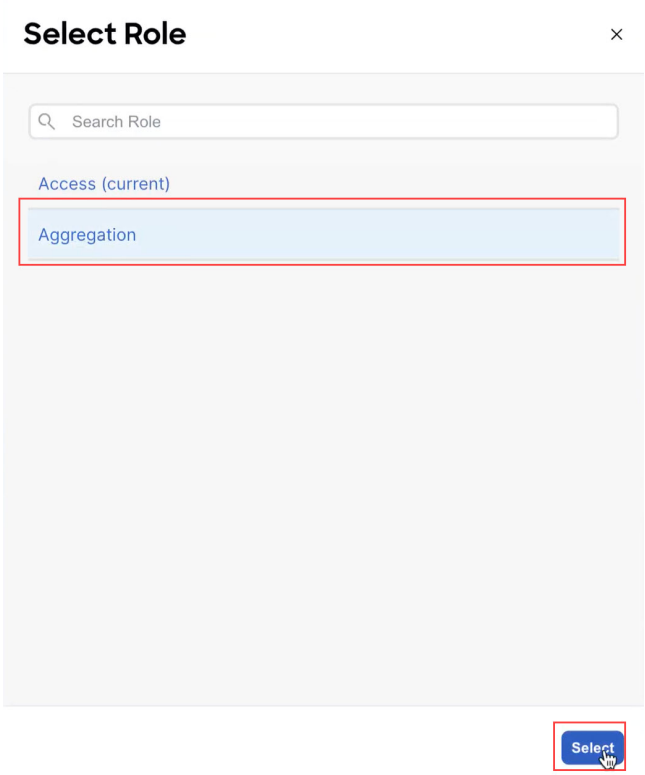
In addition, all the switches will show NA in the Config Status column, because no configurations have been pushed out to the switches at this point in the process.

c) Click the boxes next to the switches that will be assigned the **Aggregation** role.



d) Click **Actions** > **Set Role**.

e) Choose **Aggregation** from the list of roles for the switches, then click **Select**.



Click **Ok** in the warning popup that appears.

Step 12 Configure vPC pairing for the switches.

Once the roles have been defined, you can configure vPC pairing for the switches that have been assigned **Access** or **Aggregation** roles.

- For switches that are defined with the **Access** role, vPC is recommended, but is not mandatory.
- For the switches that are defined with the **Aggregation** role, vPC peering is mandatory for fabrics that are configured using the `Enhanced Classic LAN` template, as it is recommended based on Cisco best practices.

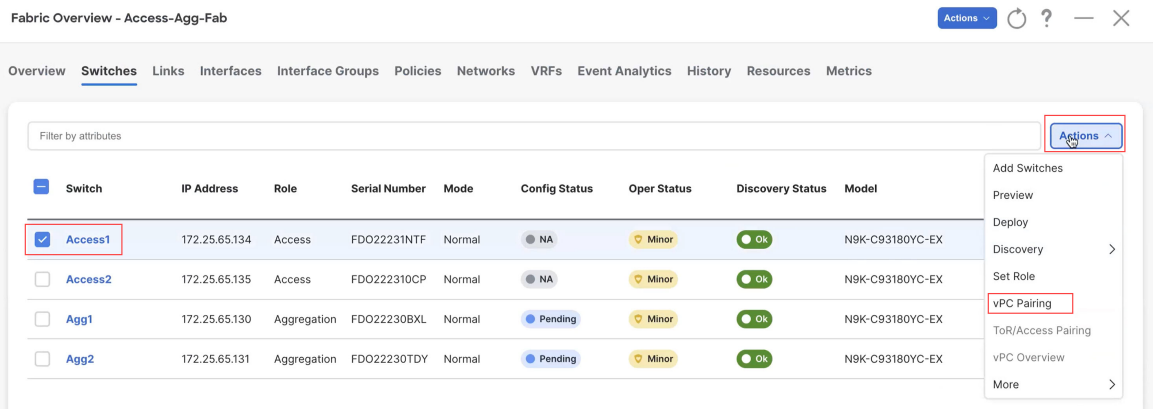
For both types of switches, a related setting in the **Advanced** tab of **Fabric Settings** is enabled by default, which allows NDFC to automatically detect and pair Access or Aggregation switches for optimal traffic engineering. You can disable this feature, if desired, in the **Enable Agg/Access Auto Pairing** field in the **Advanced** tab.

The following vPC pairing options are supported:

- Back-to-back
- Port channel
- Trunk ports

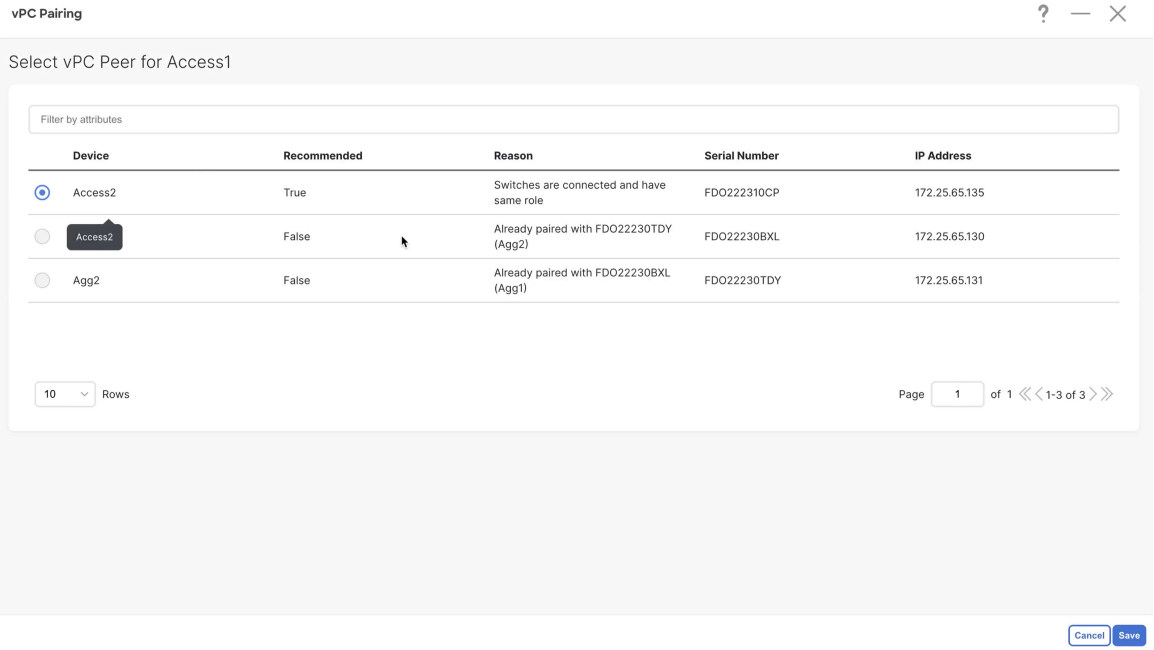
The following procedures apply, regardless of whether you are configuring a three-tier heirarchal network topology or a two-tier Collapsed Core topology:

a) To set vPC pairing for the **Access** switches, select a switch that is assigned with the **Access** role, then click **Actions** > **vPC Pairing**.

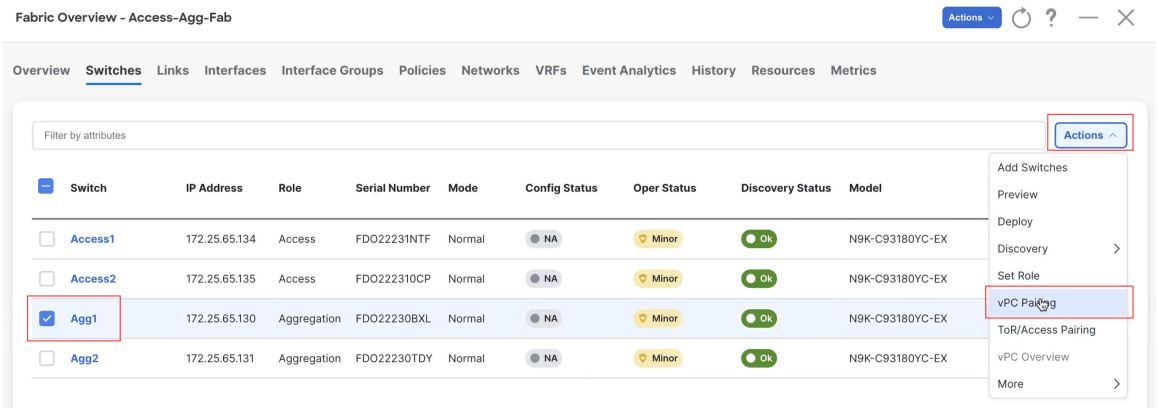


b) In the **Select vPC Peer** screen, select a second Access switch to use for the vPC pairing, then click **Save**.

The additional switches that NDFC will recommend to be used for the vPC pairing for Access switches will be shown with the value **True** under the **Recommended** column.



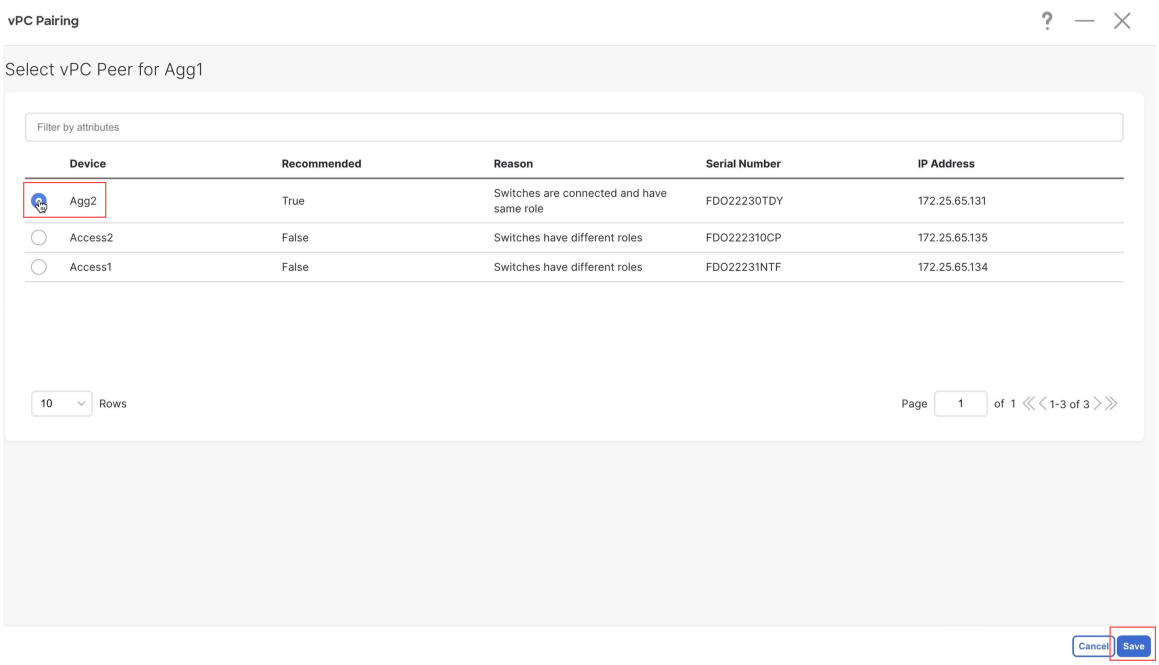
c) To set vPC pairing for the **Aggregation** switches, select a switch that is assigned with the **Aggregation** role, then click **Actions** > **vPC Pairing**.



d) In the **Select vPC Peer** screen, select a second Aggregation switch to use for the vPC pairing, then click **Save**.

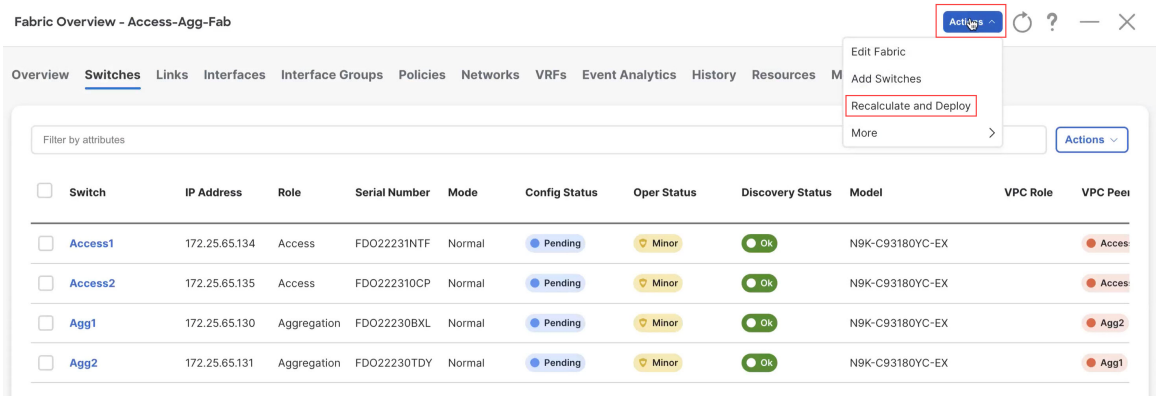
The additional switches that NDFC will recommend to be used for the vPC pairing for Aggregation switches will be shown with the value **True** under the **Recommended** column.

Note
vPCs are automatically discovered for brownfield deployments.

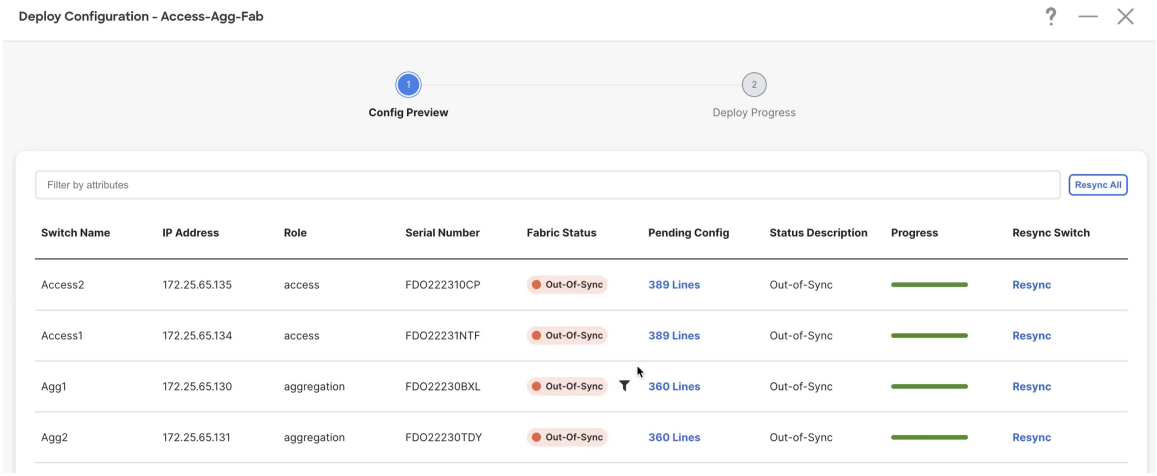


Step 13 Recalculate and deploy.

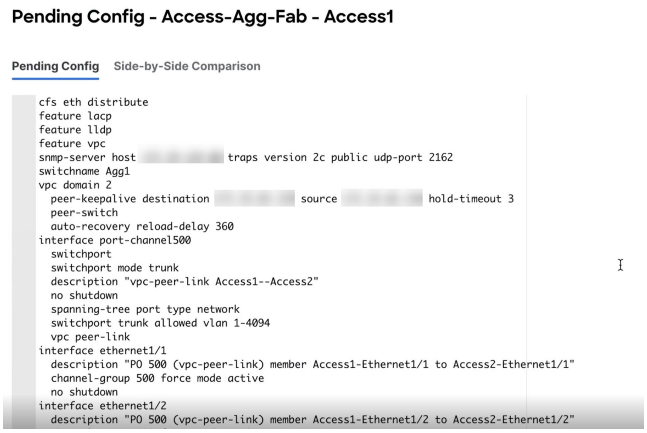
a) At the top of the page, click **Actions** > **Recalculate and Deploy**.



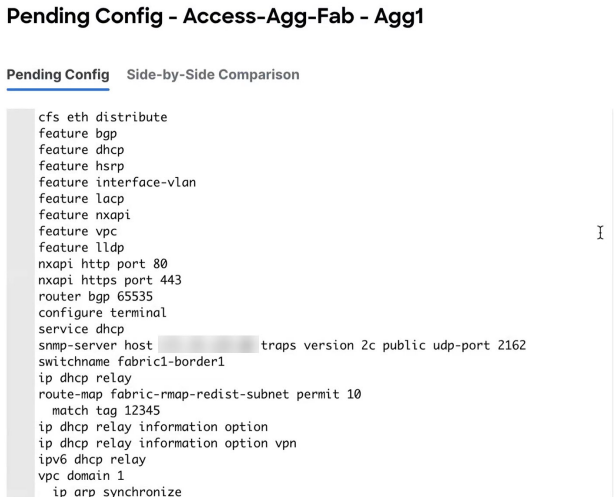
b) Preview the configuration updates as the recalculation process progresses.



You can click on the blue link in the **Pending Config** column to get additional information on the changes that are being configured for the switches. For example, if you were to click on the blue link in the **Pending Config** column for one of the Access switches, you might see information similar to the following.

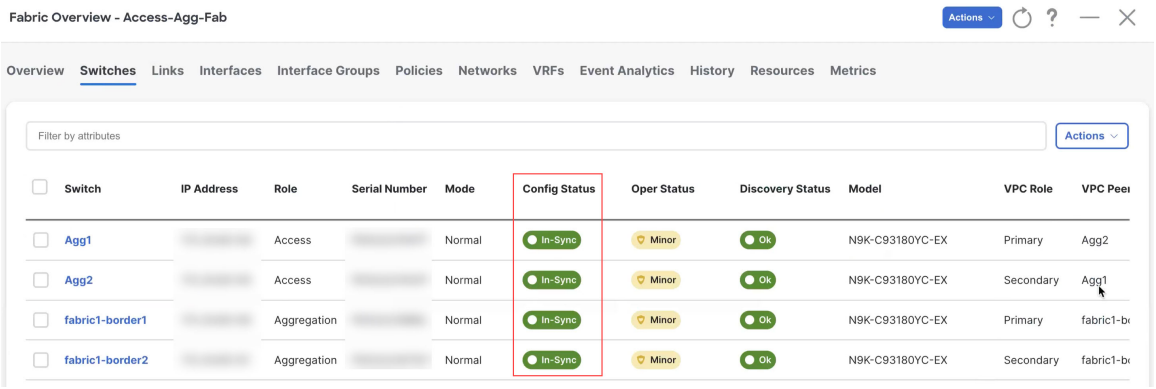


Similarly, you might see information similar to the following if you were to click on the blue link in the **Pending Config** column for one of the Aggregation switches.



c) When the recalculation process is completed, click **Deploy All**, then click **Close** when you see **Success** and **Deployment Completed** in the **Deploy Configuration** window.

d) In the **Fabric Overview** window, verify that the status shown in the **Config Status** column shows as **In-Sync**.



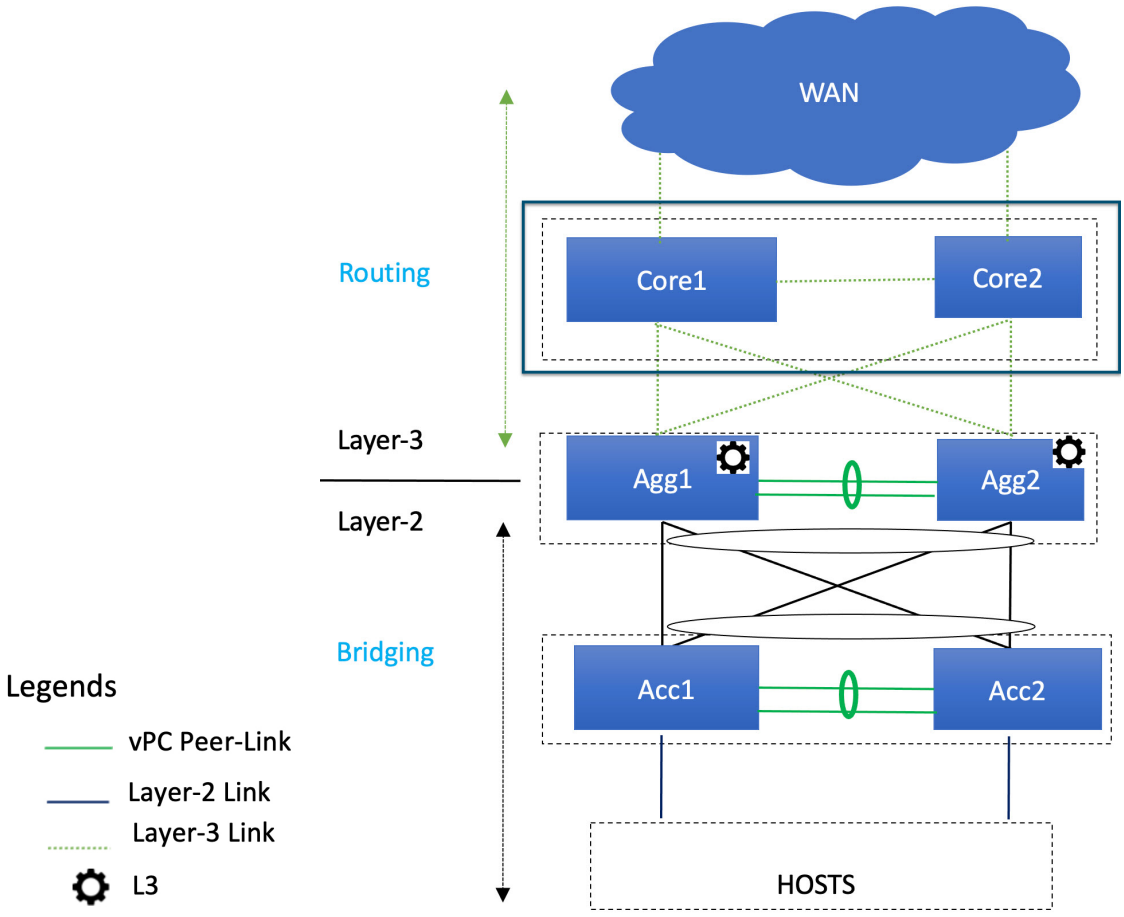
What to do next

- For a two-tier Collapsed Core topology, as described in Supported Legacy/Classic Network Topologies, on page 5, you have completed the Day 0 configurations because the Enhanced Classic LAN fabric that you configured in this topic covers both tiers in a two-tier Collapsed Core topology. You are now ready to begin the Day 1 configurations. Go to Day 1 Configurations, on page 43.
- For a three-tier hierarchical topology, as described in Supported Legacy/Classic Network Topologies, on page 5, you will configure the External Connectivity Network fabric that you will need for the Core tier next. Go to Configure the External Connectivity Network Fabric, on page 31.

Configure the External Connectivity Network Fabric

Note The procedures in this section apply only if you have a three-tier heirarchical topology, as described in Supported Legacy/Classic Network Topologies, on page 5. Do not follow these procedures if you have a two-tier, collapsed core topology, because you will have already configured the only fabric that you need for that two-tier topology in Configure the Enhanced Classic LAN Fabric, on page 15.

In these procedures, you will be configuring an External Connectivity Network fabric specifically for the core tier in a three-tier topology, as described in Supported Legacy/Classic Network Topologies, on page 5.



You already configured the Enhanced Classic LAN fabric for the access and aggregation tiers in a previous section. You will configure two separate fabrics because typical deployments use a shared core, which will reside in a separate, External Connectivity Network fabric that is shared by the first (access-aggregate) fabric that you configured in previous procedures.

You will be using the External Connectivity Network fabric template for these configurations. For more information on that fabric template, see *Enhanced Classic LAN*.

Before you begin

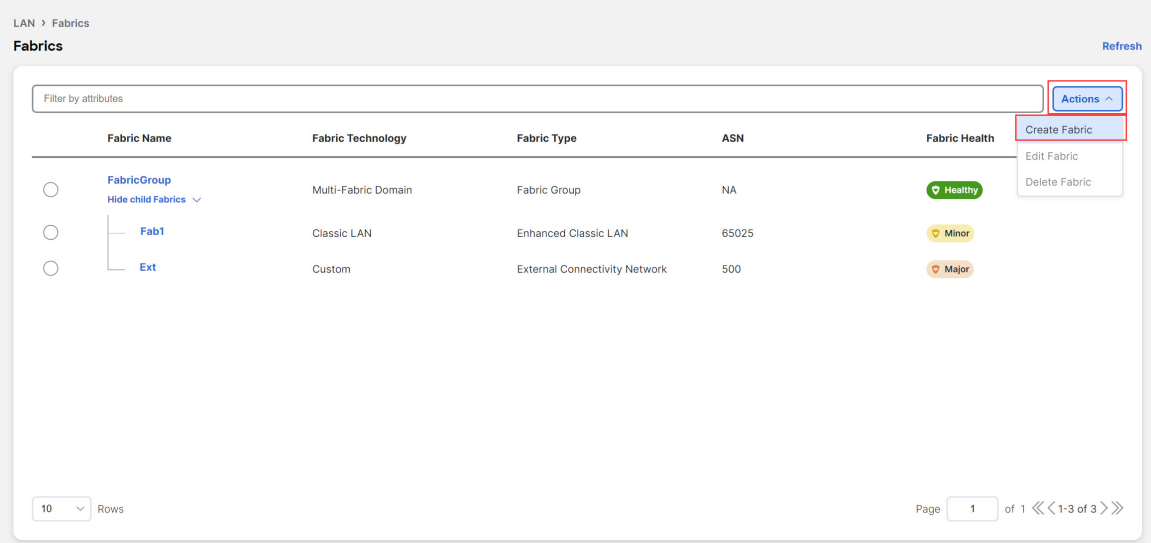
Verify that you have configured the Enhanced Classic LAN fabric using the procecures provided in Configure the Enhanced Classic LAN Fabric, on page 15.

Procedure

Step 1 In NDFC, navigate to **LAN** > **Fabrics**, if you are not there already.

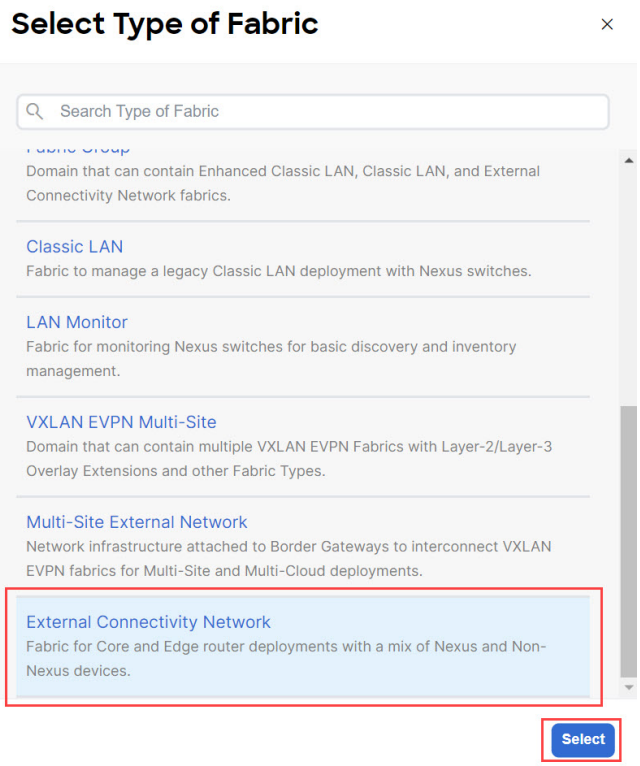
A page showing all of the configured fabrics appears.

Step 2 Click **Actions** > **Create Fabric**.



Step 3 In the **Create Fabric** screen, enter a name for the External Connectivity Network fabric (for example, `Core-Fab`), then click **Choose Template**.

Step 4 Choose the `External Connectivity Network` template, then click **Select**.

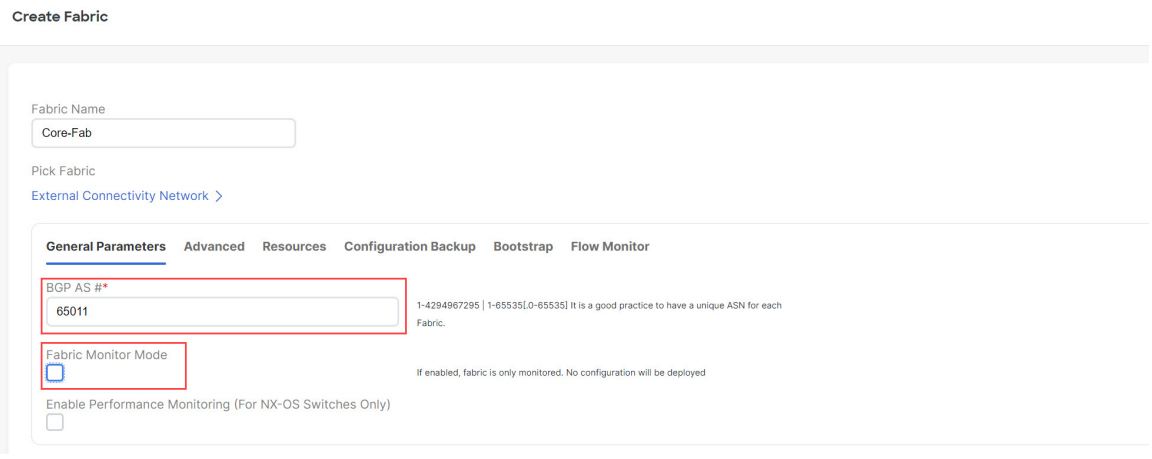


Step 5 In the **General Parameters** tab, make the necessary configuration specifically for this use case.

- In the **BGP ASN** field, enter a unique, non-overlapping value.

  For this use case, we will use `65011` as the BGP ASN number that will be assigned to the External Connectivity Network fabric.

- Uncheck the box next to the **Fabric Monitor Mode** field.



Step 6 Determine how you want NDFC to discover the switches in the External Connectivity Network fabric, if applicable.

The options available to you depend on the type of deployment that you are managing through NDFC:

- If you are managing a *greenfield* deployment, then, similar to the Enhanced Classic LAN fabric, you can use either of these two methods to discover the switches in the fabric:
  - By manually entering the necessary information to allow NDFC to discover those switches. This option is applicable if you have already configured certain parameters, such as the out-of-band management IP addresses, on the switches that need to be discovered.

    If you decide to use this method, that step is provided later in these procedures (Step 9, on page 34).
  - By using the Power On Auto Provisioning (POAP) feature in NDFC. This option is useful if you do not already have certain parameters, such as the management IP address, default route, and start up configurations, already configured on the switches that need to be discovered.

    If you decide to use this method, click the **Bootstrap** tab in this page to make the necessary configurations in the **Bootstrap** area using the same process as provided in the procedures for the Enhanced Classic LAN fabric.
- If you are managing a *brownfield* deployment, then the only option available for you is to manually enter the necessary information to allow NDFC to discover the switches; you cannot use the Power On Auto Provisioning (POAP) feature when managing a brownfield deployment. Go to Step 9, on page 34 for those procedures.

Step 7 Complete the remaining configurations for the External Connectivity Network fabric, if necessary.

The remaining parameters are optional. They have been automatically configured based on Cisco best practices, but you can change them if necessary.

Step 8 When you have completed the necessary configurations to the External Connectivity Network fabric template, click **Save**.

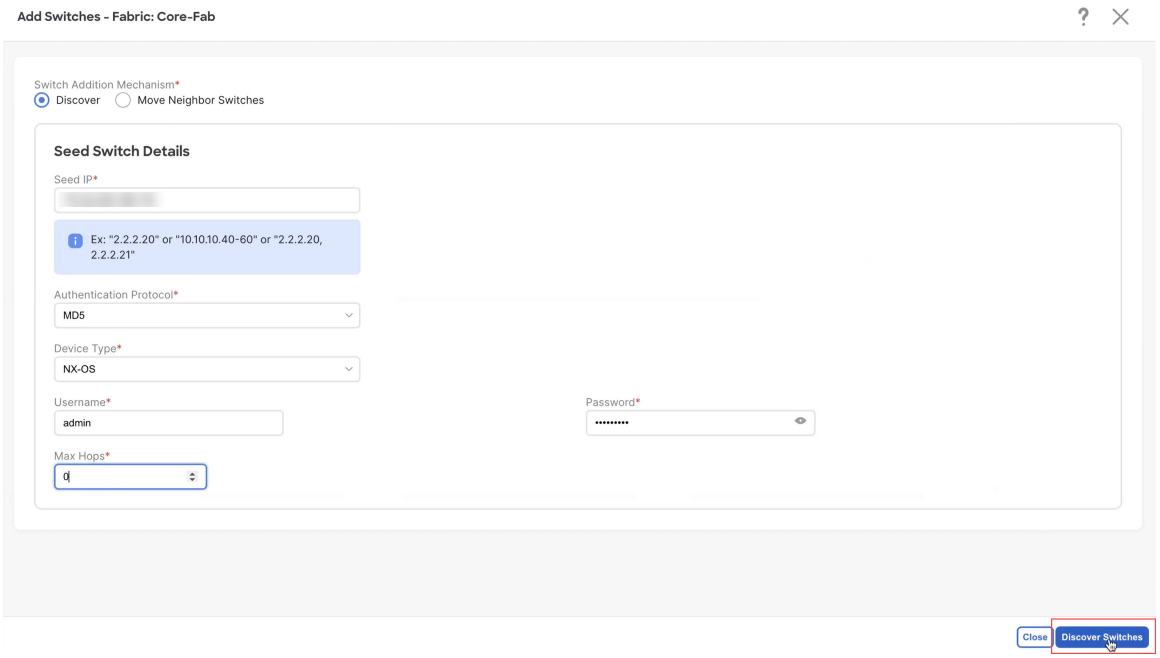
The **LAN Fabric** page appears again, with the newly created External Connectivity Network fabric added to the list of configured fabrics.

Step 9 If necessary, manually enter the necessary information to allow NDFC to discover the switches in the External Connectivity Network fabric.

Note
You do not have to go through the procedures in this step if you enabled POAP through the **Bootstrap** area in NDFC in the previous step.

Verify that reachability exists between NDFC and these switches before proceeding with this step.

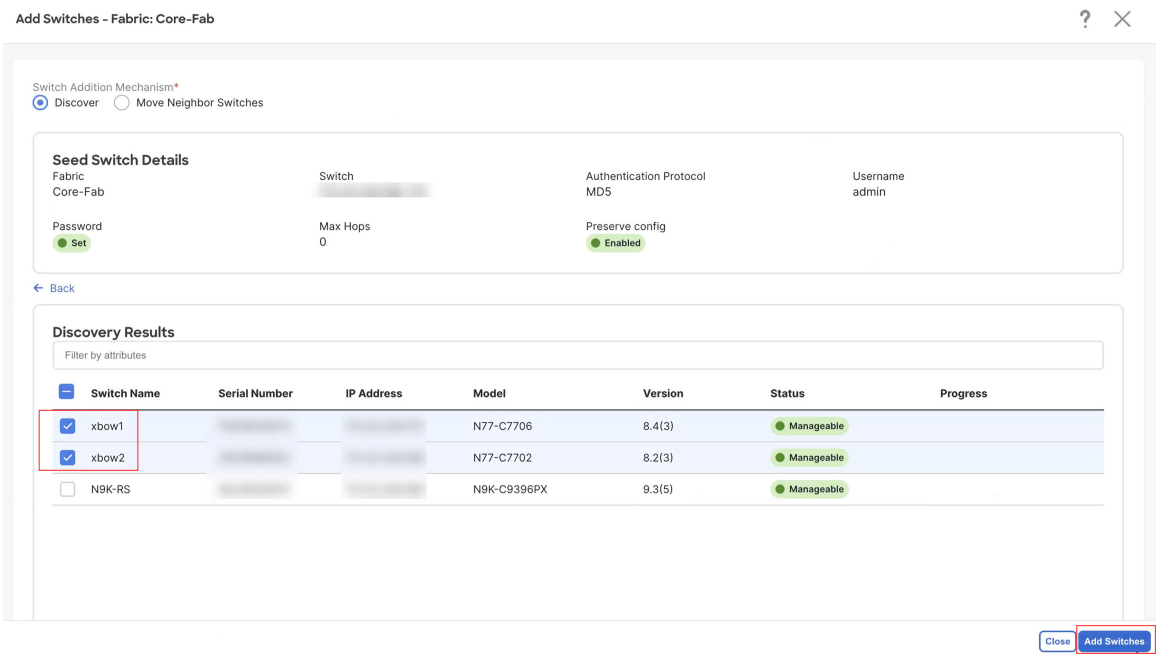
a) In the **Overview** page for the External Connectivity Network fabric, click the **Switches** tab, then click **Action** > **Add Switches**.
b) In the **Add Switches** screen, verify that the **Discover** option is selected, then add the necessary information to discover the switches.
   - In the **Seed IP** field, enter the management IP addresses of the switches. Only out-of-band management of switches is supported.
   - Enter any remaining information in the **Add Switches** screen that is necessary to discover the switches.



c) Click **Discover Switches**.

Click **Confirm** in the confirmation popup window that appears.

d) In the **Discovery Results** screen, check the check box next to the switches that will be imported into the External Connectivity Network fabric and click **Add Switches**.

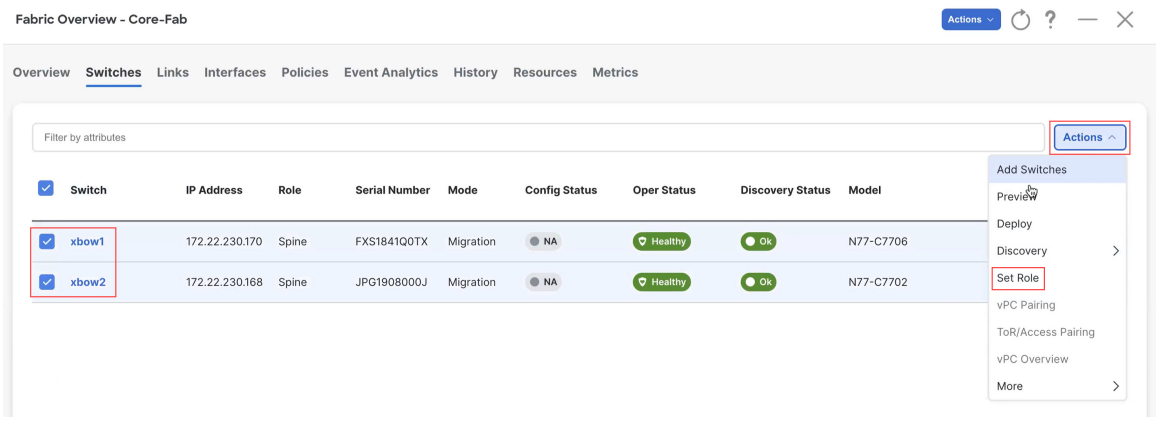


The status will change to **Switch Added** when the process is complete. Click **Close** to close out of this window.

Step 10 Define the role for the switches in the External Connectivity Network fabric.

Similar to the process in the Enhanced Classic LAN fabric, once the switches are discovered in the External Connectivity Network fabric, the next step is to define the roles for those switches. Based on the roles that you assign to the switches, the appropriate configuration will be generated and pushed to the switches by NDFC.

a) Click the box next to the switch for the External Connectivity Network fabric, then click **Actions** > **Set Role**.



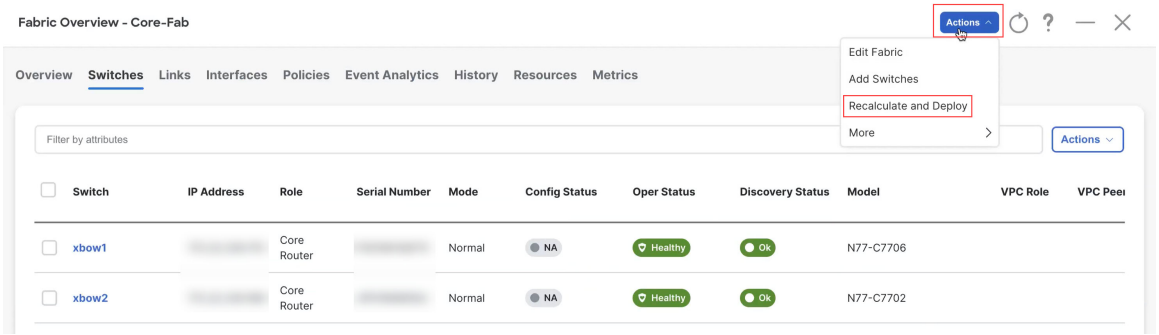
b) Determine what role you want to set for the switches in the External Connectivity Network.

As described in Supported Legacy/Classic Network Topologies, on page 5, you can set the roles for the switches in the External Connectivity Network to either the `Core Router` or `Edge Router` role.

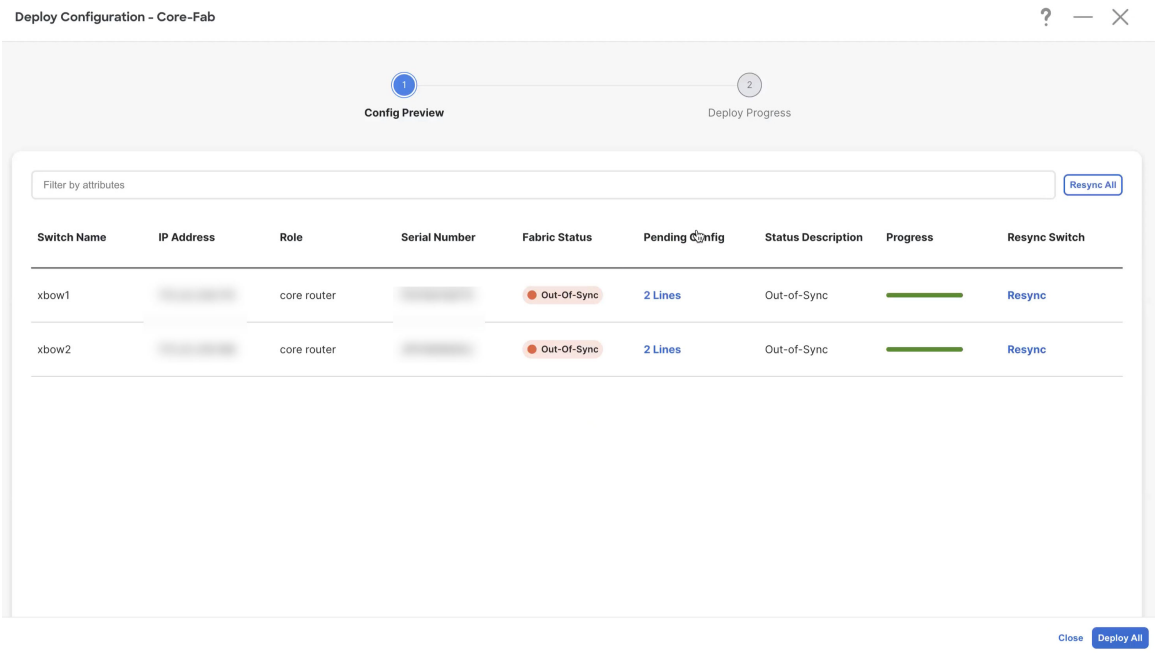
Locate and select either the `Core Router` or the `Edge Router` option in the **Select Role** list, then click **Select**.

Step 11 Recalculate and deploy.

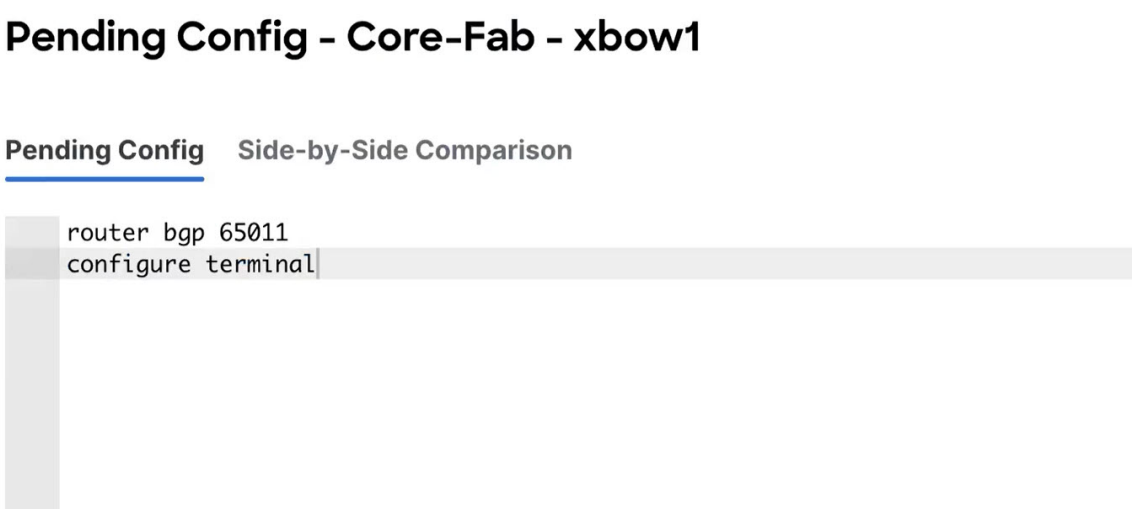
a) At the top of the page, click **Actions** > **Recalculate and Deploy**.



b) Preview the configuration updates as the recalculation process progresses.

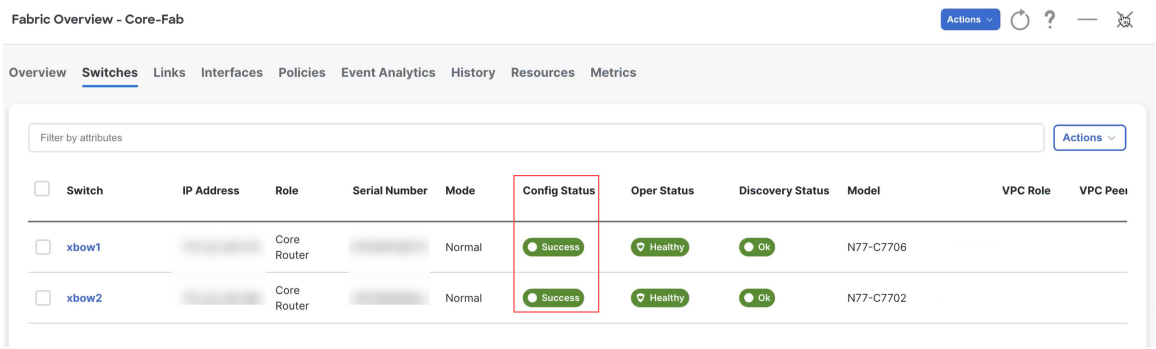


You can click on the blue link in the **Pending Config** column to get additional information on the changes that are being configured for the switches.



c) When the recalculation process is completed, click **Deploy All**, then click **Close** when you see **Success** and **Deployment Completed** in the **Deploy Configuration** window.

d) In the **Fabric Overview** window, verify that the status shown in the **Config Status** column shows as **Success**.



Step 12 Navigate to **LAN** > **Fabrics** and select the External Connectivity Network fabric that you just created.

The **Overview** page for this External Connectivity Network fabric appears.

Step 13 Click the **Switches** tab to verify that the switch that you just added for the second fabric appears correctly.

Step 14 Click the **X** at the top right corner of the window to exit out of this page.

What to do next

If you want to group the fabrics together using the Fabric Group fabric template, go to Create a Fabric Group, on page 38.

Create a Fabric Group

Note The procedures in this section apply only if you have a three-tier heirarchical topology, where you have two separate fabrics that you might want to show under a fabric group. You do not need these procedures if you have a two-tier, Collapsed Core topology, because there is only one fabric configured for that type of topology.

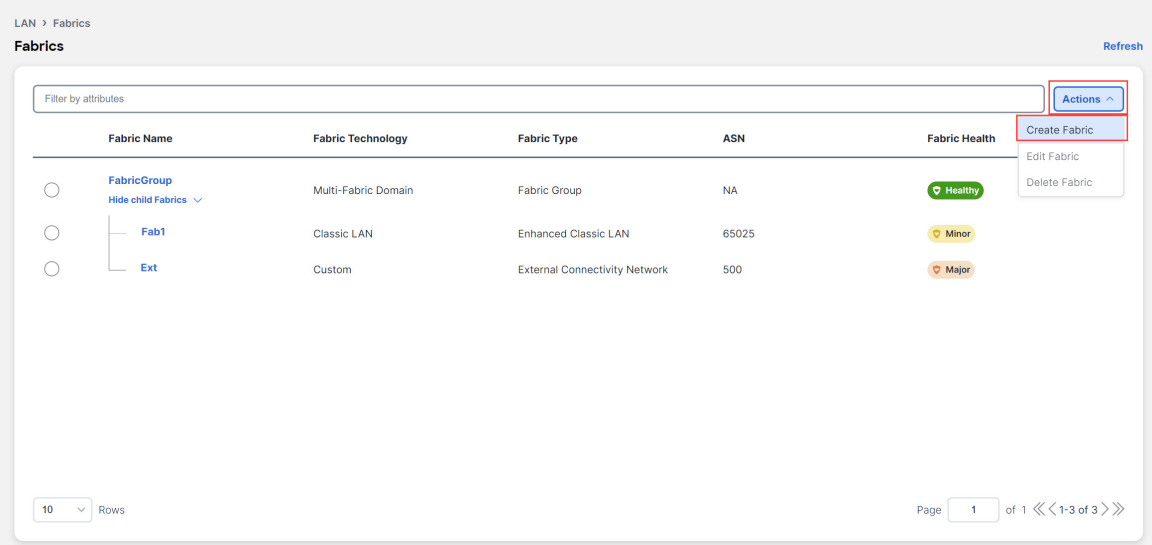
If you would like to have a group visualization for the Topological view, you can create a `Fabric Group` fabric type with the Access-Aggregrate and Core fabrics as child members of this group. You will be using the Fabric Group fabric template for these configurations. For more information on that fabric template, see *Enhanced Classic LAN*.

Procedure

Step 1 In NDFC, navigate to **LAN** > **Fabrics**, if you are not there already.

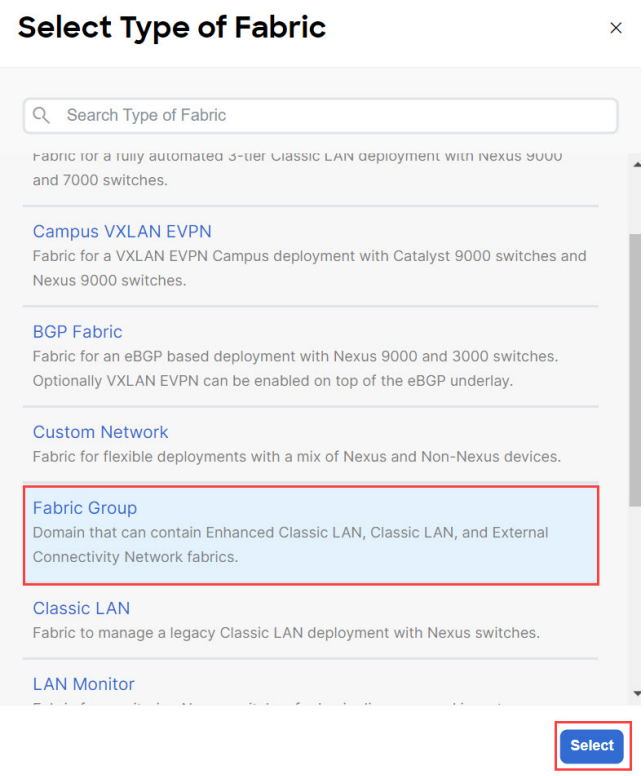
A page showing all of the configured fabrics appears.

Step 2 Click **Actions** > **Create Fabric**.



Step 3 In the **Create Fabric** screen, enter a name for the new fabric (for example, `Classic-Group`), then click **Choose Template**.

Step 4 Choose the `Fabric Group` template, then click **Select**.



Step 5 Click **Save** in the **Create Fabric** page.

You are returned to the **LAN Fabrics** page.

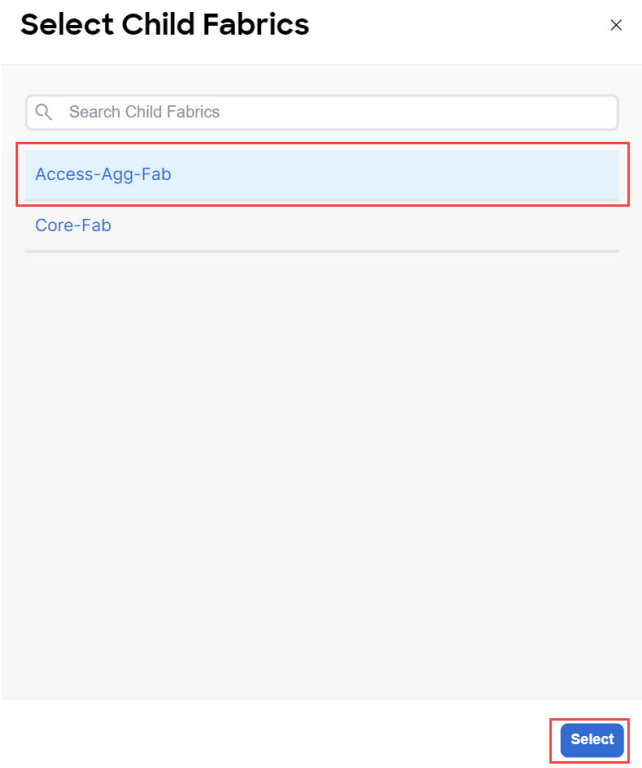
Step 6 Double-click the fabric group that you just created.

The **Fabric Overview** page appears for the fabric group.

Step 7 Click **Actions** > **Add Child Fabric**.

The **Select Child Fabrics** page appears.

Step 8 Choose the Enhanced Classic LAN fabric that you created using the procedures provided in Configure the Enhanced Classic LAN Fabric, on page 15, then click **Select**.



You are returned to the **Fabric Overview** page.

Step 9 Click **Actions** > **Add Child Fabric** again.

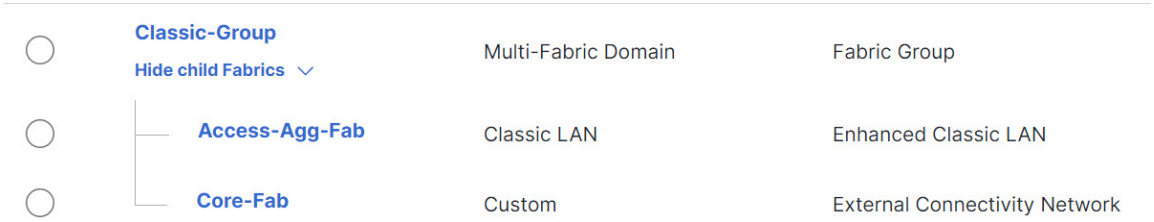
The **Select Child Fabrics** page appears.

Step 10 Choose the External Connectivity Network fabric that you created using the procedures provided in Configure the External Connectivity Network Fabric, on page 31, then click **Select**.

You are returned to the **Fabric Overview** page.

Step 11 Click the **X** in the upper right corner of the page.

The **LAN Fabrics** page appears again, with the newly created fabric group added to the list of configured fabrics.

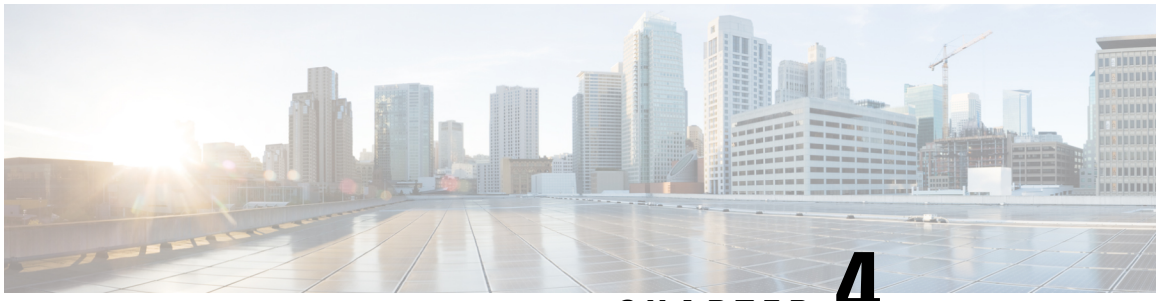


Right-click operations are now available from the Topology page per switch.

Step 12 If necessary, navigate to **LAN** > **Switches**, select one or more switches, then click **Actions** > **Deploy**.

What to do next

Begin the Day 1 configurations using the procedures provided in Day 1 Configurations, on page 43.



CHAPTER 4

Day 1 Configurations

- Create a Layer 2 Network, on page 43
- Create a Layer 3 Network, on page 48
- Configure the VRF Lite Extension, on page 54

Create a Layer 2 Network

Procedure

Step 1 In NDFC, navigate to **LAN** > **Fabrics**, if you are not there already.

A page showing all of the configured fabrics appears.

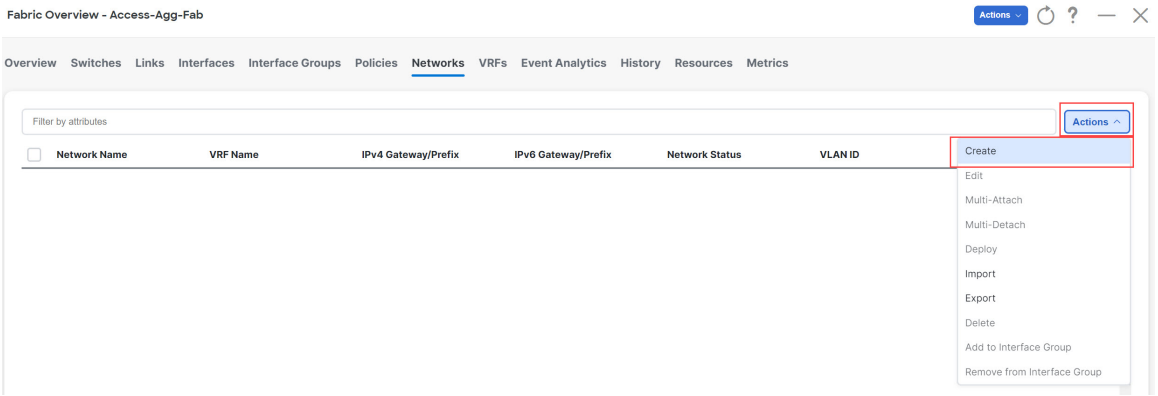
Step 2 Double-click the Enhanced Classic LAN fabric that you created using the procedures provided in Configure the Enhanced Classic LAN Fabric, on page 15.

The **Fabric Overview** page for that fabric appears.

Step 3 Click the **Networks** tab.

Step 4 Create the Layer 2 network.

a) In the **Networks** tab, click **Actions** > **Create**.

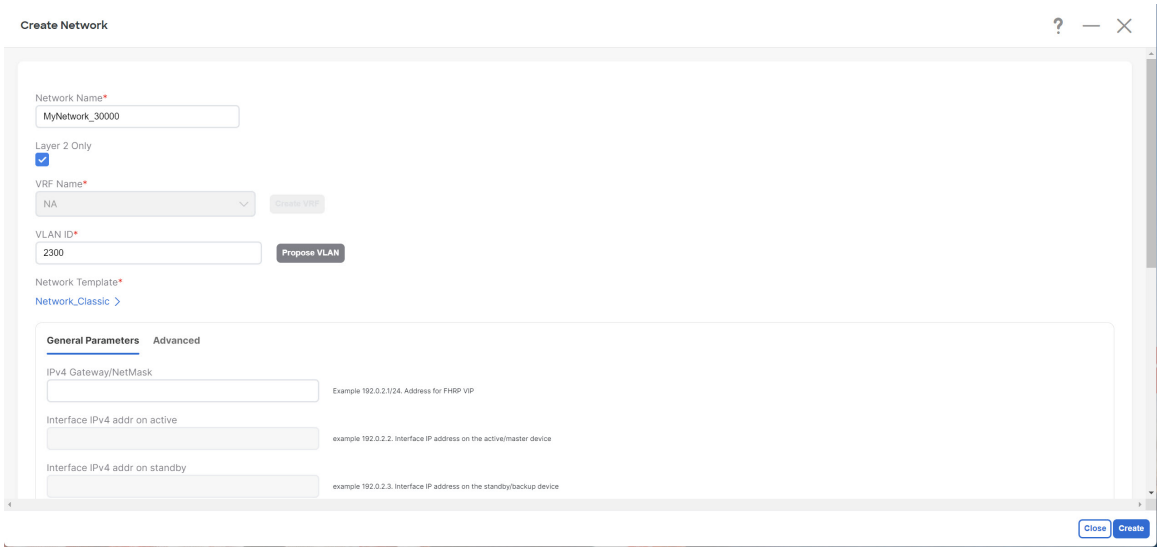


The **Create Network** window appears.

b) In the **Create Network** screen, enter the necessary information.

- A default name for the Layer 2 network is automatically generated in the **Network Name** field, but you can change it, if necessary.
- In the **Layer 2 Only** field, check the box to enable this option. This specifies that this network is Layer 2 only.
- In the **VLAN ID** field, enter a value to use for the associated VLAN, or click **Propose VLAN** to have NDFC propose a VLAN ID for your Layer 2 network, based on the available resources (the range is customizable in **Fabric Settings**).
- In the **Network Template** field, leave the default `Network_Classic` option selected.

  This is the correct template to use for the Layer 2 network.
- The gateway for a Layer 2 network resides outside of the fabric; therefore, the IP addresses in the **General Parameters** page are left empty.



c) Fill in the remaining fields as necessary for your Layer 2 network, if necessary.

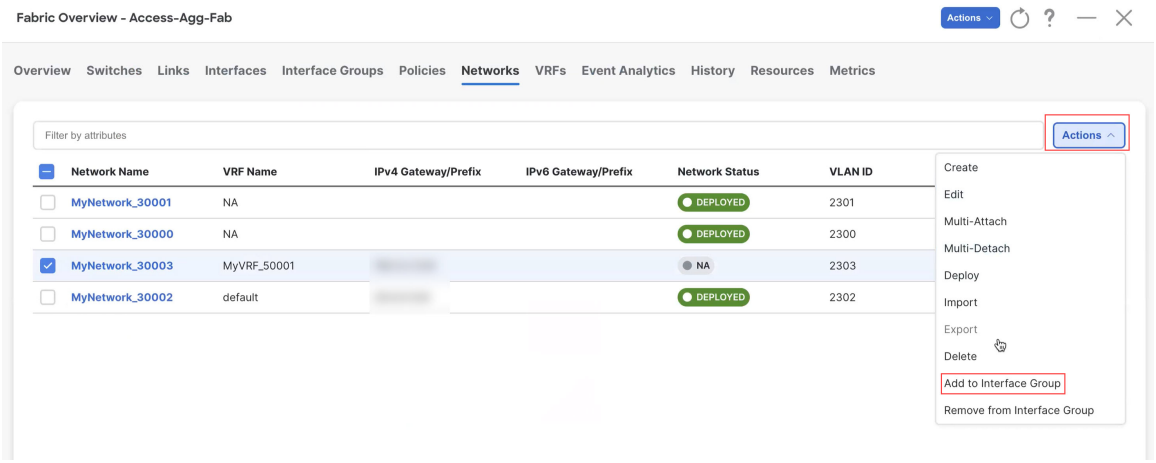
d) Click **Create**.

You are returned to the **Networks** tab for the Enhanced Classic LAN fabric.

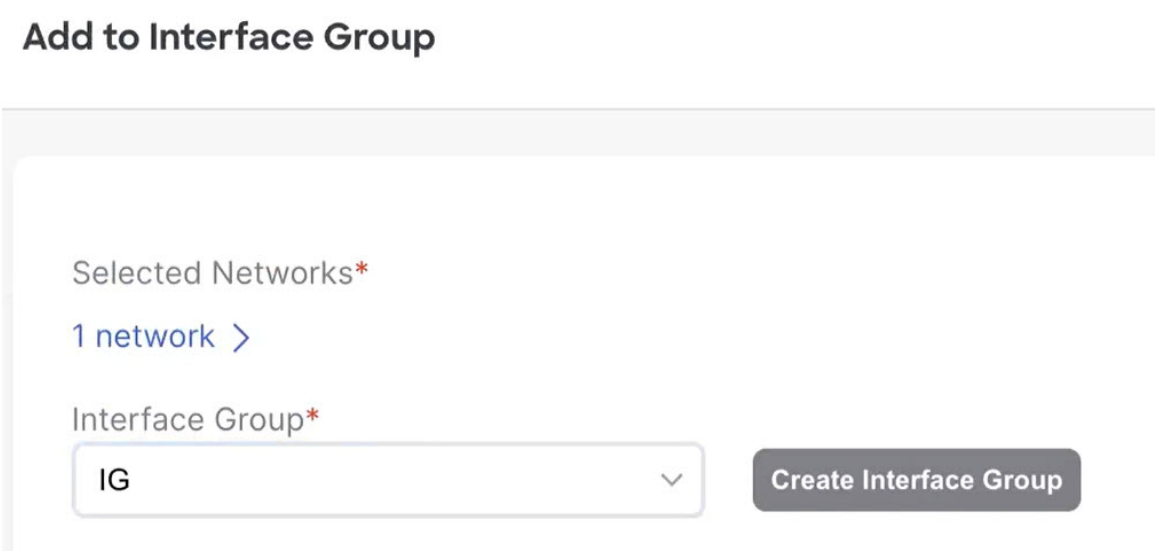
Step 5 Add the network to an interface group, if necessary.

Interface groups are useful if you are trying to deploy a network to a group of interfaces.

a) In the **Networks** tab for the Enhanced Classic LAN fabric, click on the Layer 2 network that you just created, then click **Actions** > **Add to Interface Group**.



b) Select the interface group that you want to add, or click **Create Interface Group** to create a new interface group to add.



For more detailed procedures on adding a network to an interface group, see Interface Groups in the *Cisco NDFC-Fabric Controller Configuration Guide* .

Step 6 Attach the network.

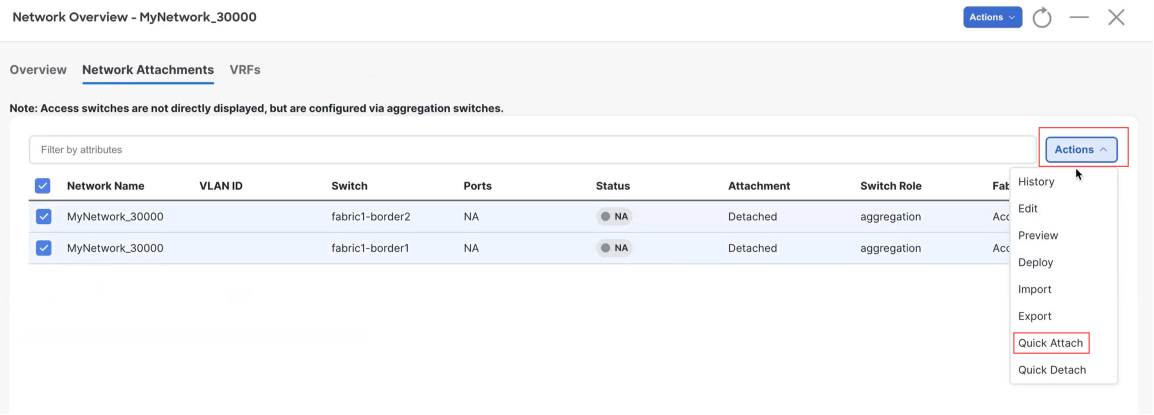
Once you've created the Layer 2 network, you can attach it to host-facing ports on the Access switch, which will then allow the VLAN on these trunk or access ports and also on the vPC/port channel/standalone ports between the Access and Aggregation switches.

Determine if you want to perform a **quick attach** or a **multi-attach**.

- If you want to perform a **quick attach**, where you will attach this network to the selected switches, follow these steps:

a) In the **Networks** tab for the Enhanced Classic LAN fabric, double-click on the Layer 2 network that you just created.
b) Click the **Network Attachments** tab.

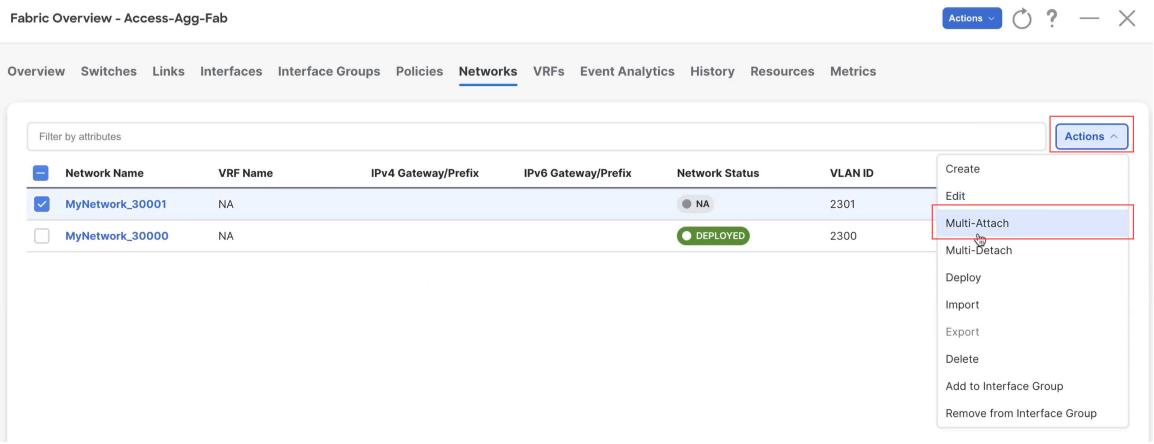
c) Locate the switches with **aggregation** shown in the **Switch Role** column and click the boxes next to those switches.



d) Click **Actions** > **Quick Attach**, then go to Step 7, on page 47.

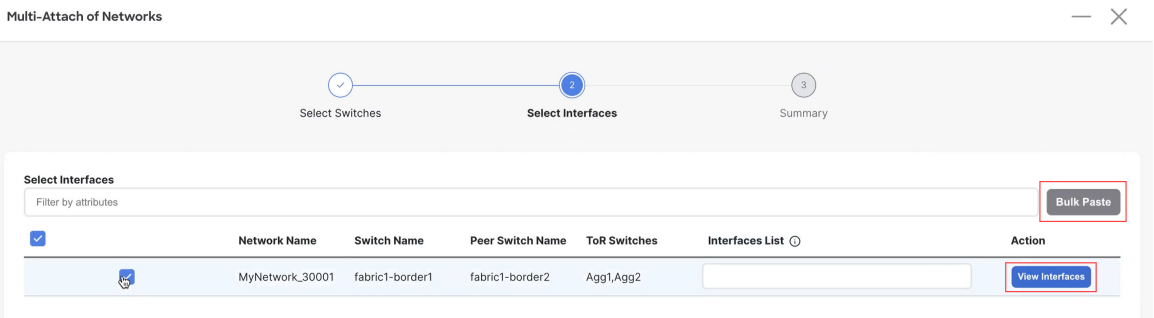
• If you want to perform a **multi-attach**:

a) In the **Networks** tab for the Enhanced Classic LAN fabric, click on the Layer 2 network that you just created, then click **Actions** > **Multi-Attach**.



b) Click the box next to each switch that you want to attach to the network, then click **Next**.

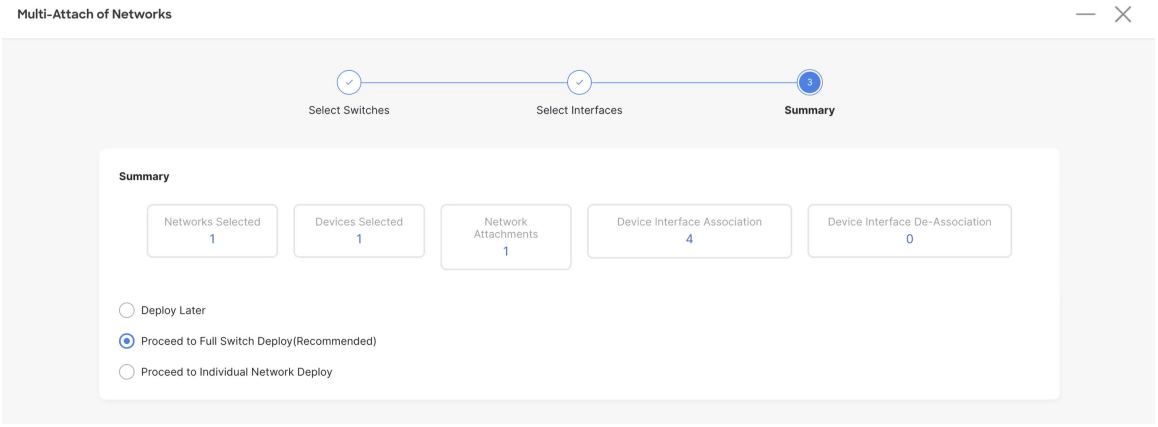
c) Click the box next to each interface that you want to select, then select the interfaces using either **Bulk Paste** or **View Interfaces**.



• If you select **Bulk Paste**, enter the interfaces to be pasted in the **Interfaces List** area, then click **Save**.

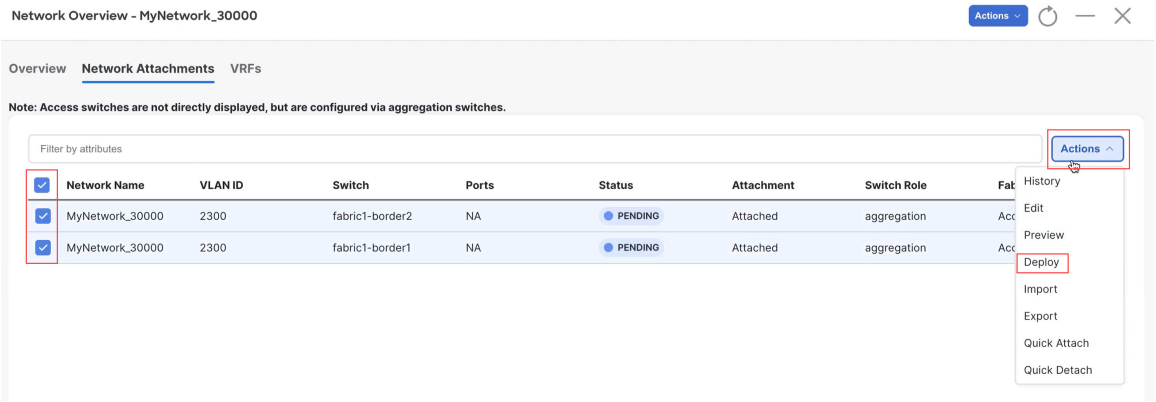
- If you select **View Interfaces**, click the boxes next to the specific ports that you want to attach in the following screen, then click **Save**.

d) In the Select Interfaces window, click **Next**.
e) Make the necessary deployment selection in the next window, then click **Save**.

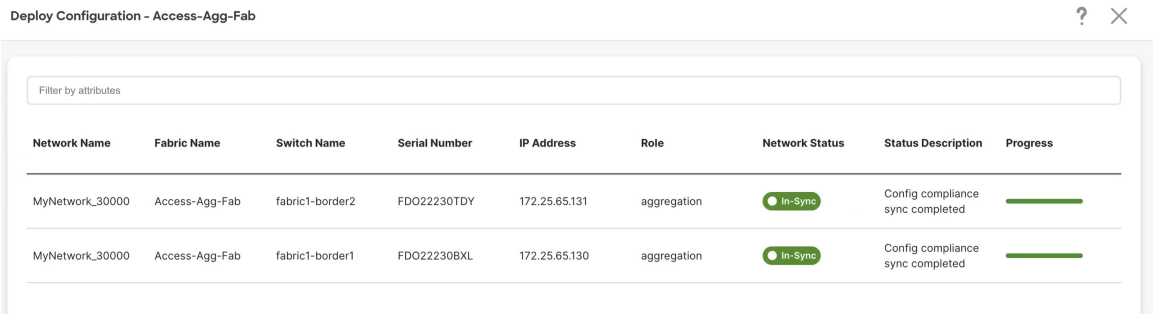


Step 7 Deploy the network.

In the **Network Attachments** window, select the networks that you just attached, then click **Actions** > **Deploy**.



Step 8 When the recalculation process is completed, click **Deploy** and verify that the status shown in the **Config Status** column shows as **In-Sync**.



What to do next

Create a Layer 3 network using the procedures provided in Create a Layer 3 Network, on page 48.

Create a Layer 3 Network

A Layer 3 network can be in either a default or a custom VRF.

Before you begin

Procedure

Step 1 Navigate to the **Networks** tab for the Enhanced Classic LAN fabric, if you are not there already.

a) In NDFC, navigate to **LAN** > **Fabrics**.

A page showing all of the configured fabrics appears.

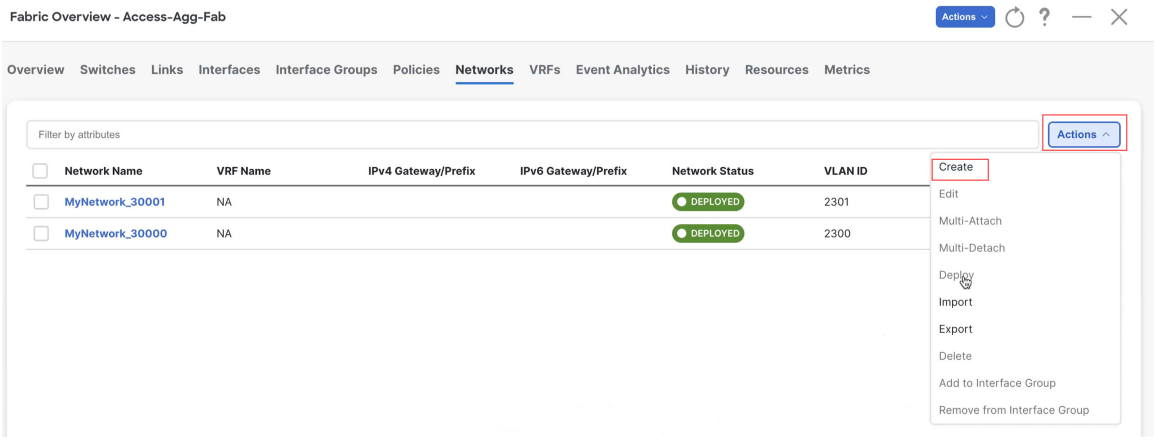
b) Double-click the Enhanced Classic LAN fabric that you created using the procedures provided in Configure the Enhanced Classic LAN Fabric, on page 15.

The **Fabric Overview** page for that fabric appears.

c) Click the **Networks** tab.

Step 2 Create the Layer 3 network.

a) In the **Networks** tab, click **Actions** > **Create**.



The **Create Network** window appears.

b) In the **Create Network** screen, click the **General Parameters** tab and enter the necessary information.

- Change the name for the Layer 3 network in the **Network Name** field, if necessary.

  The entry in the **Network Name** field is auto-populated but can be changed, if necessary.

- In the **Layer 2 Only** field, verify that there is no check in the box so that this option is not enabled.

  You do not want the **Layer 2 Only** option enabled for a Layer 3 network.

c) Determine if you want to use the default VRF or if you want to create a VRF for this Layer 3 network.

- If you want to use the default VRF for this Layer 3 network, in the **VRF Name** field, choose `default` from the list of VRFs available. Go to 2.d, on page 50.
- Follow these steps if you want to create a custom VRF for this Layer 3 network:

  1. In the **VRF Name** field, click **Create VRF**.

     The **Create VRF** page appears.

     Change the name for the custom VRF for the Layer 3 network in the **VRF Name** field, if necessary. The entry in the **VRF Name** field is auto-populated but can be changed, if necessary.

  2. Click the **General Parameters** tab and enter the necessary information.

     (Optional) Enable the **Enable Auto Peering over SVI Between VPC Aggregations** option if desired, which enables per VRF iBGP or OSPF peering between Aggregation switches. Note that the protocol used is based on the VRF-Lite routing protocol configured in the Fabric setting.

  3. Click the **Advanced** tab and enter the necessary information.

     The settings in the **Advanced** area include options for BGP authentication, route maps, and static 0/0 configurations, where you could configure a default (0/0) route towards the core switch.

  4. Click **Save** to save the custom VRF for the Layer 3 network.

     You are returned to the **Create Network** page.

     Once this Layer 3 network is attached in the next step, if the **Enable Auto Peering over SVI Between VPC Aggregations** option is enabled, NDFC will create the configuration for an iBGP peering session or an

OSPF neighborship between the Aggregation switches for this new VRF, including the VLAN ID and the IP address for the SVI.

d) In the **Create Network** screen, enter the necessary information in the remaining fields in the **General Parameters** tab.

- In the **VLAN ID** field, enter a value to use for the associated VLAN, or click **Propose VLAN** to have NDFC propose a VLAN ID for your Layer 3 network, based on the available resources (the range is customizable in **Fabric Settings**).
- Define the gateway IP address in either the **IPv4 Gateway/Netmask** or the **IPv6 Gateway/Netmask** fields.

  You must define the IPv4 or IPv6 gateway IP address for the Layer 3 network. For these procedures, the gateway is the Aggregation switch within the fabric.

  Note
  For the next two fields in the **General Parameters** tab, note that you will be defining the active and standby devices in a later step in these procedures.

- Define the interface address on the active/master device in either the **Interface IPv4 addr on active** or the **Interface IPv6 addr on active** fields.
- Define the interface address on the standby/backup device in either the **Interface IPv4 addr on standby** or the **Interface IPv6 addr on standby** fields.

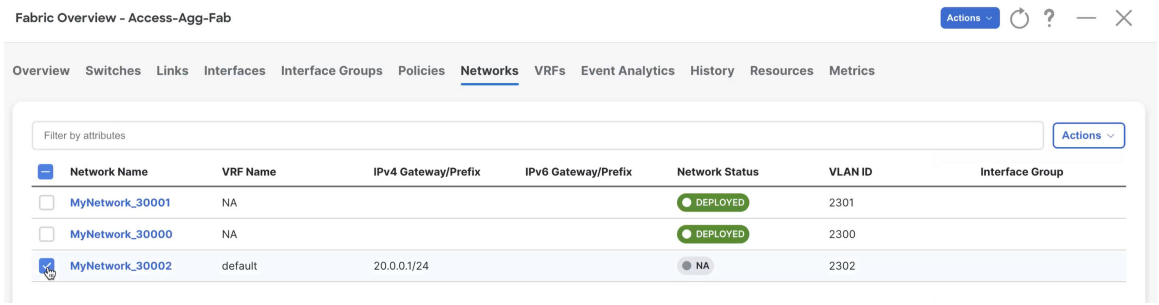
e) In the **Create Network** screen, click the **Advanced** tab and enter the necessary information.

Customize the First Hop Redundancy Protocol settings (either HSRP or VRRP, based on the fabric settings) in this page. You can determine which of the Aggregation switches will be used for the active and standby First Hop Redundancy Protocol when you attach the network.

f) Fill in the remaining fields as necessary for your Layer 3 network.

g) Click **Create**.

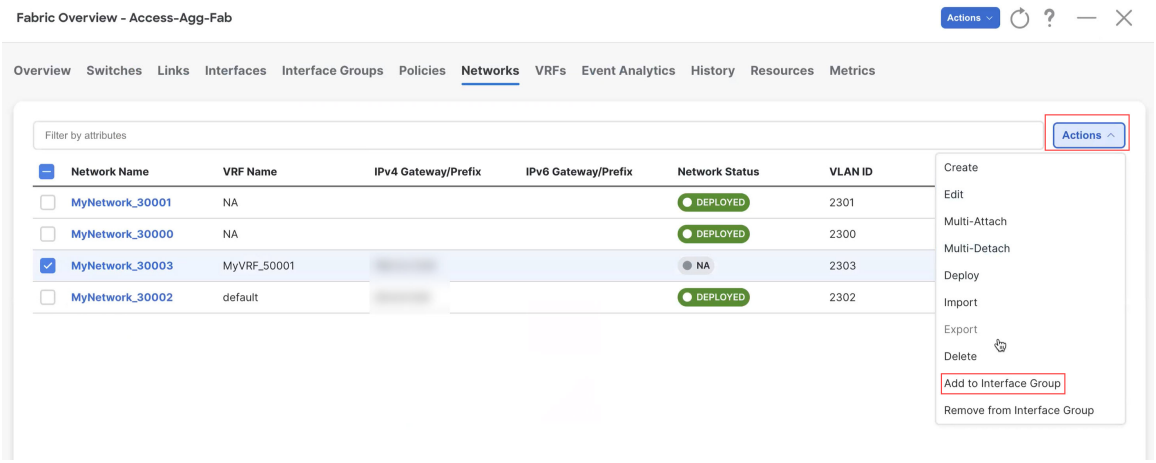
The **Networks** page displays the newly created Layer 3 network and accompanying VRF.



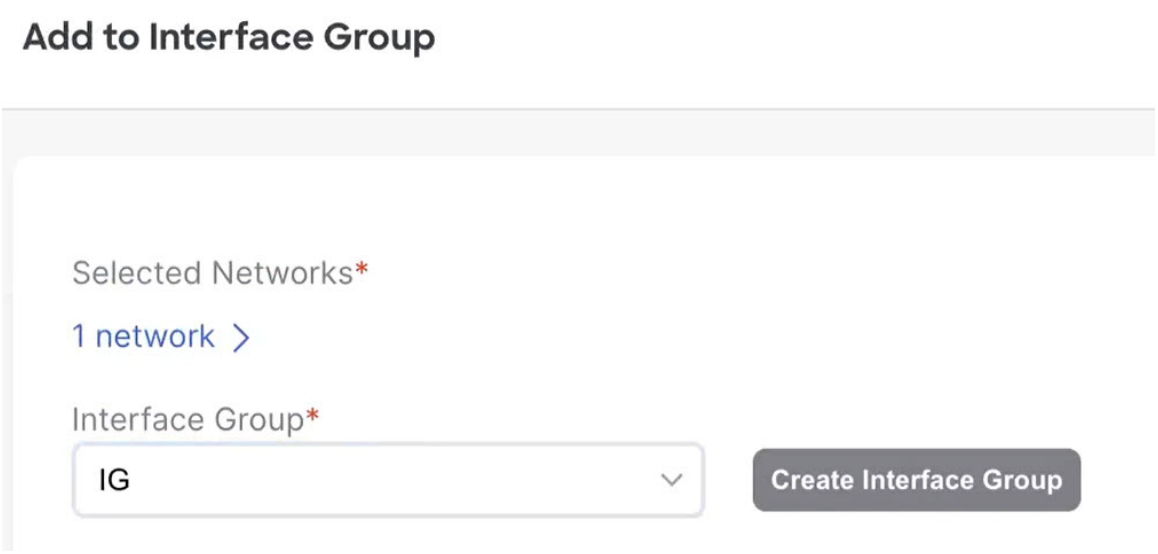
Step 3 Add the network to an interface group, if necessary.

Interface groups are useful if you are trying to deploy a network to a group of interfaces.

a) In the **Networks** tab for the Enhanced Classic LAN fabric, click on the Layer 3 network that you just created, then click **Actions** > **Add to Interface Group**.



b) Select the interface group that you want to add, or click **Create Interface Group** to create a new interface group to add.



For more detailed procedures on adding a network to an interface group, see Interface Groups in the *Cisco NDFC-Fabric Controller Configuration Guide* .

Step 4 Attach the network.

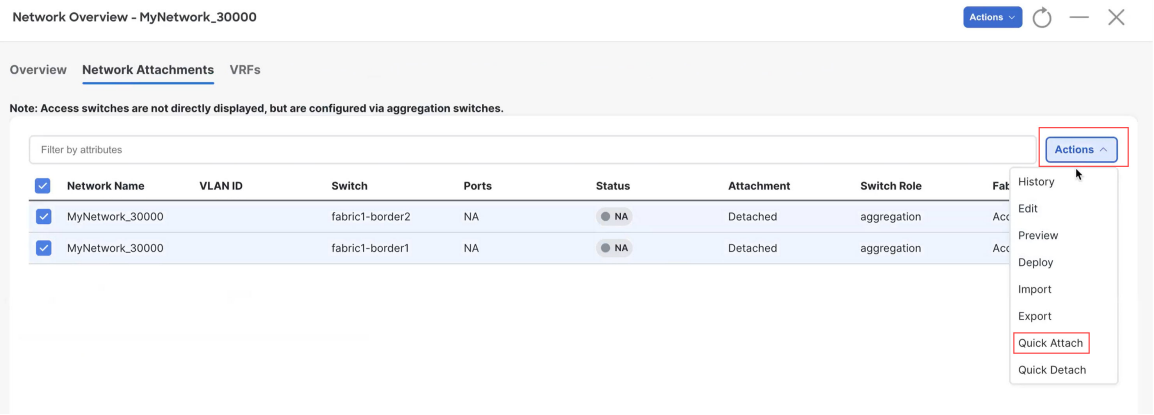
Once you've created the Layer 3 network, you can attach it to host-facing ports on the Access switch, which will then allow the VLAN on these trunk or access ports and also on the vPC/port channel/standalone ports between the Access and Aggregation switches.

Determine if you want to perform a **quick attach** or a **multi-attach**.

- If you want to perform a **quick attach**, where you will attach this network to the selected switches, follow these steps:

a) In the **Networks** tab for the Enhanced Classic LAN fabric, double-click on the Layer 3 network that you just created.
b) Click the **Network Attachments** tab.

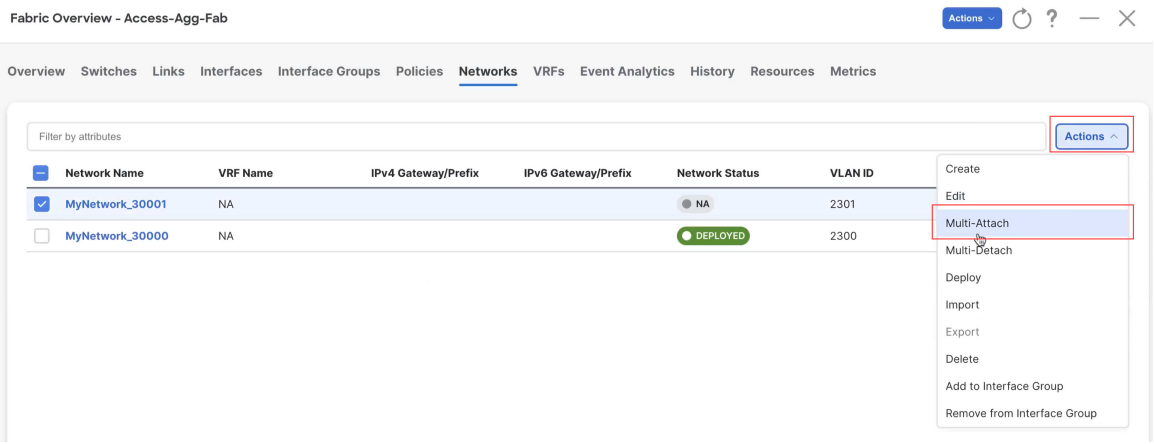
c) Locate the switches with **aggregation** shown in the **Switch Role** column and click the boxes next to those switches.



d) Click **Actions** > **Quick Attach**, then go to Step 5, on page 53.

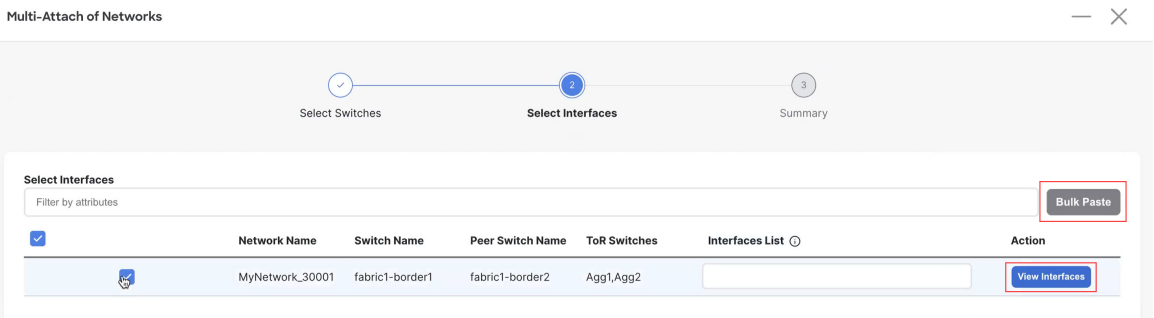
• If you want to perform a **multi-attach**:

a) In the **Networks** tab for the Enhanced Classic LAN fabric, click on the Layer 3 network that you just created, then click **Actions** > **Multi-Attach**.



b) Click the box next to each switch that you want to attach to the network, then click **Next**.

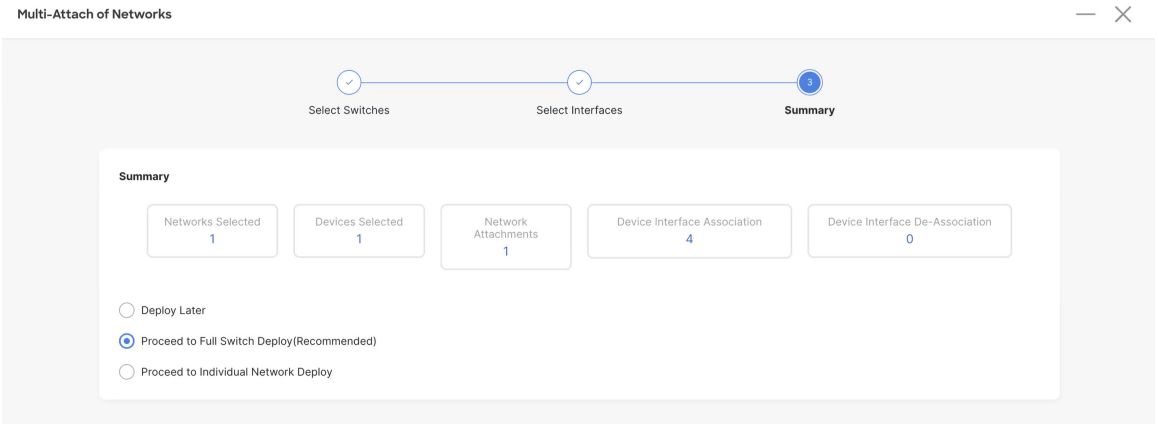
c) Click the box next to each interface that you want to select, then select the interfaces using either **Bulk Paste** or **View Interfaces**.



• If you select **Bulk Paste**, enter the interfaces to be pasted in the **Interfaces List** area, then click **Save**.

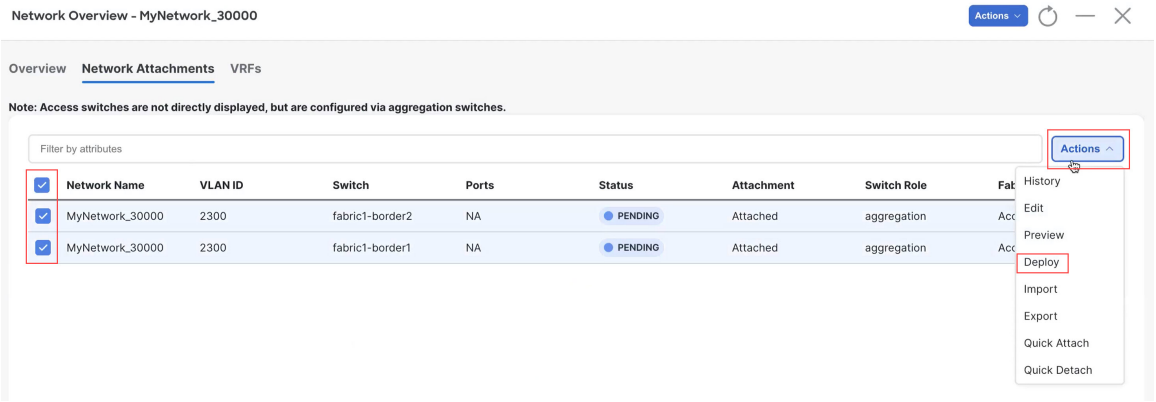
- If you select **View Interfaces**, click the boxes next to the specific ports that you want to attach in the following screen, then click **Save**.

d) In the Select Interfaces window, click **Next**.
e) Make the necessary deployment selection in the next window, then click **Save**.

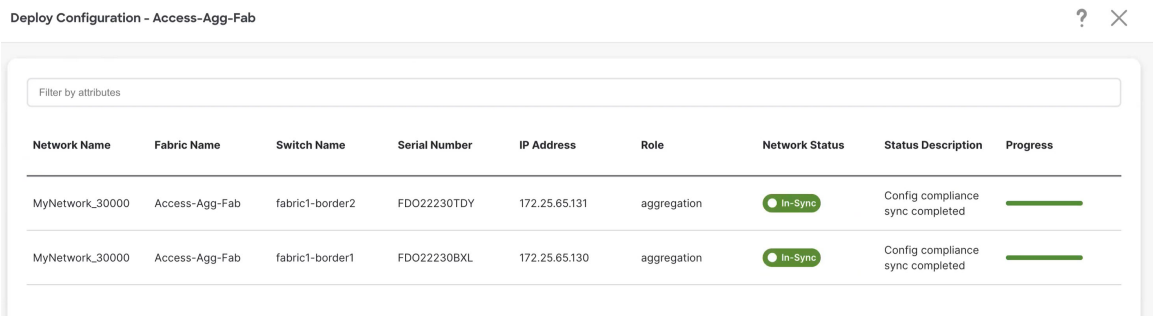


Step 5 Deploy the network.

In the **Network Attachments** window, select the networks that you just attached, then click **Actions** > **Deploy**.



Step 6 When the recalculation process is completed, click **Deploy** and verify that the status shown in the **Config Status** column shows as **In-Sync**.



What to do next

Configure the VRF Lite extension using the procedures provided in Configure the VRF Lite Extension, on page 54.

Configure the VRF Lite Extension

In these procedures, you will configure the VRF Lite extension between the Aggregation and Core switches. NDFC supports either automatic or manual configuration modes for VRF Lite between the Aggregation and Core switches. For this use case, we will configure VRF Lite using the automatic configuration.

You can go through the configurations in this section whether you have a three-tier topology or a two-tier (Collapsed Core) topology. However, in order to be able to configure the VRF Lite extension, you must have created a custom VRF for the Layer 3 network in Create a Layer 3 Network, on page 48, rather than using a default VRF.

Procedure

Step 1 Verify that the necessary default configurations in the **Resources** page are set correctly for VRF Lite.

a) In NDFC, navigate to **LAN** > **Fabrics**, if you are not there already.

A page showing all of the configured fabrics appears.

b) Double-click the Enhanced Classic LAN fabric that you created using the procedures provided in Configure the Enhanced Classic LAN Fabric, on page 15.

The **Fabric Overview** page for that fabric appears.

c) Click **Actions** > **Edit Fabric**.

The **Edit Fabric** page appears.

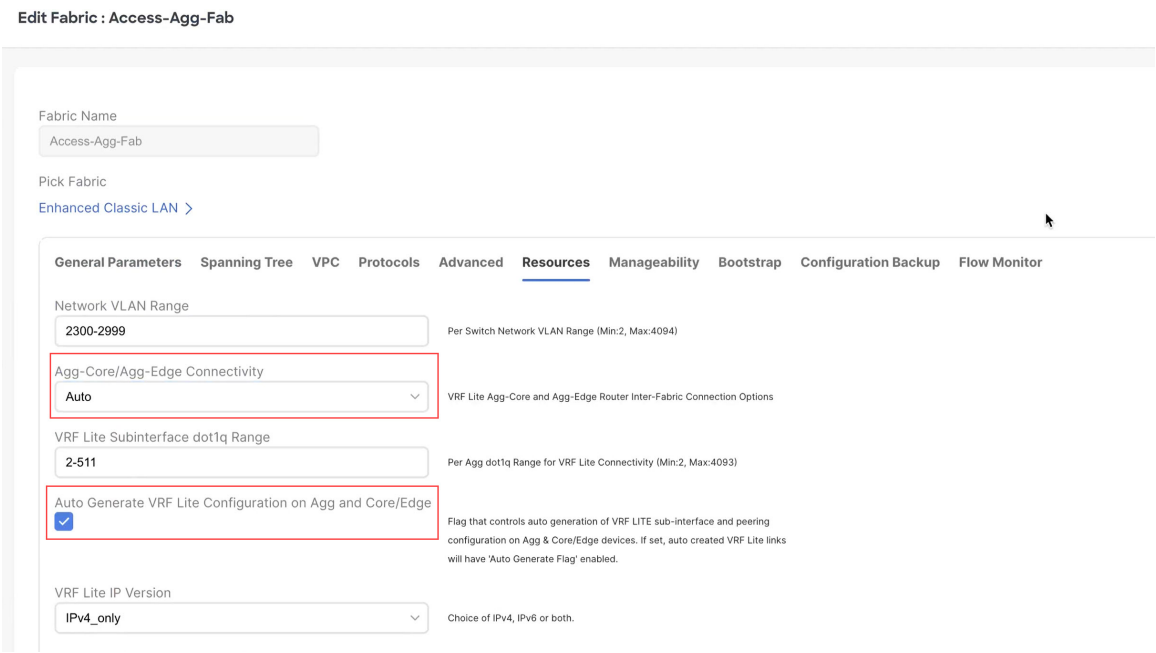
d) Click the **Resources** tab.

e) Verify that the necessary default configurations in the **Resources** page are set correctly for VRF Lite.

- In the **Agg-Core/Agg-Edge Connectivity** field, verify that the **Auto** option is set.

The Aggregation-Core peering protocol that will be used is based on the setting that you entered in the **Routing Protocol** field in Configure the Enhanced Classic LAN Fabric, on page 15.

- Check the box in the **Auto Generate VRF Lite Configuration on Agg and Core/Edge** field to enable this option.
  - If the **Routing Protocol** field is set to **ebgp** and you are using the Cisco Nexus 7000 or 9000 Series switches, or the Cisco Catalyst 9000 series switches for the Core layer, enabling the **Auto Generate VRF Lite Configuration on Agg and Core/Edge** option automatically generates the VRF Lite configuration on the Aggregation and Core switches.
  - If the **Routing Protocol** field is set to **ospf** and you are using the Cisco Nexus 7000 or 9000 Series switches for the Core layer, enabling the **Auto Generate VRF Lite Configuration on Agg and Core/Edge** option automatically generates the VRF Lite configuration on the Aggregation and Core switches.
  - In other cases, such as using Cisco ASR 9000 Series Aggregation Service Routers as the Core Router or Edge Router, then the VRF Lite intent and configurations will not be automatically generated on the Core Router or Edge Router. Instead, for each VRF, you must manually create a policy using the necessary policy.



Step 2 Verify that the links between the Aggregation and Core switches have the correct templates attached and the proper settings applied.

a) Navigate back to the **Fabric Overview** page for the Enhanced Classic LAN fabric.

The **Overview** page for that fabric appears.

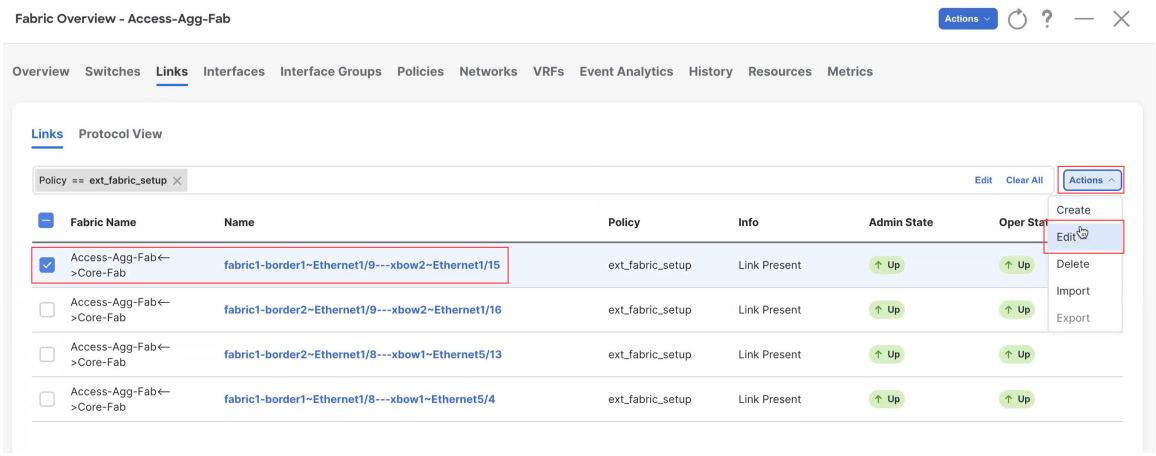
b) Click the **Links** tab.

c) Locate the links from the Aggregation switches to the Core switches in this page.

For example, you would have two Aggregation switches and one Core switch for this use case, so you would therefore locate these two links in this page:

- The link between the first Aggregation switch and the Core switch
- The link between the second Aggregation switch and the Core switch

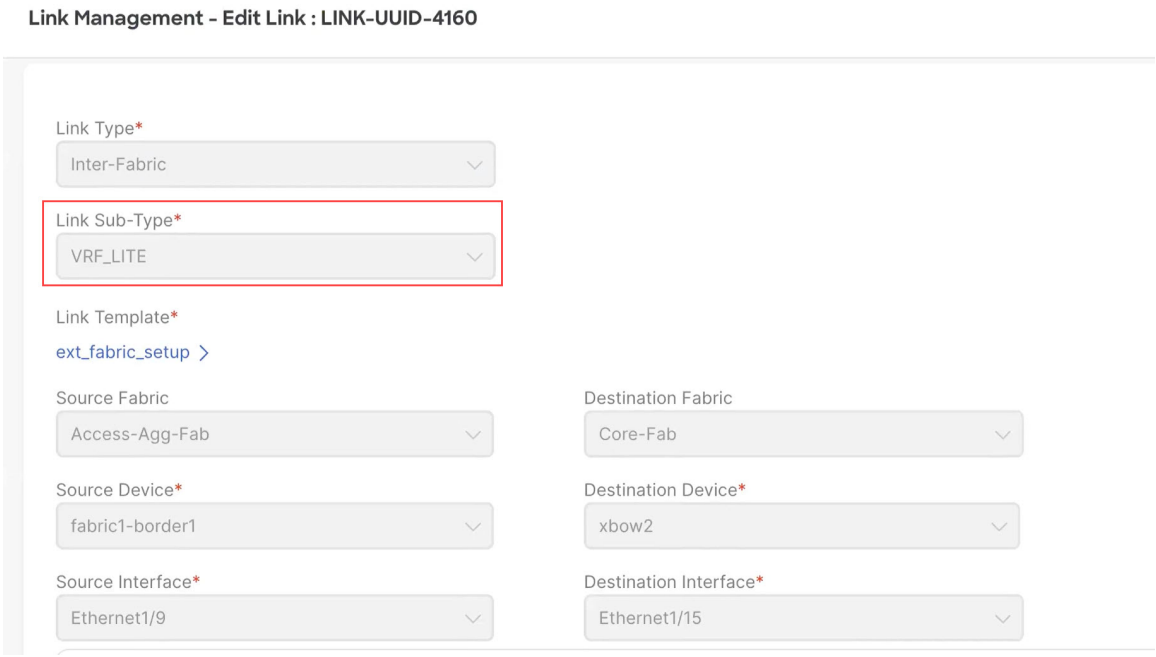
d) Click on the box next to the link between the first Aggregation switch and the Core switch, then click **Actions** > **Edit**.



The **Link Management - Edit Link** page appears for this link.

e) Verify that `VRF_LITE` is automatically selected in the **Link Sub-Type** field.

f) Verify that all of the remaining parameters are automatically populated correctly, such as the source and destination fabrics, devices, and interfaces.



g) (Optional) Repeat these steps for the remaining links from the Aggregation switches to the Core switches in this page, if you want additional verification.

For example, since we have two Aggregation switches in this use case, you would click the link between the second Aggregation switch and the Core switch in this page, then repeat these steps to verify that the links between the second Aggregation switch and the Core switch have the correct templates attached and the proper settings applied.

Step 3 In the **Fabric Overview** page for the Enhanced Classic LAN fabric, click the **VRFs** tab.

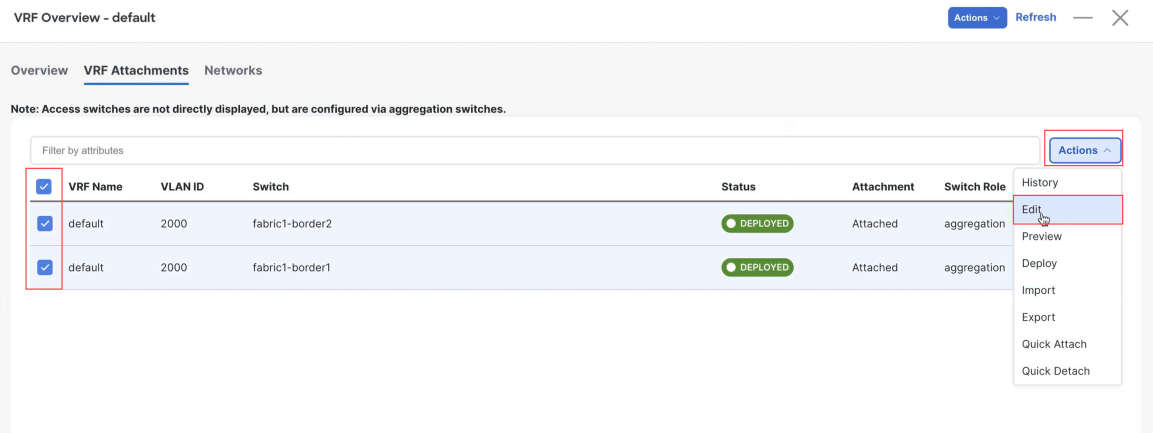
Step 4 Double-click the custom VRF that you created for the Layer 3 network.

The **VRF Overview** page appears for this VRF.

Step 5 Click the **VRF Attachments** tab.

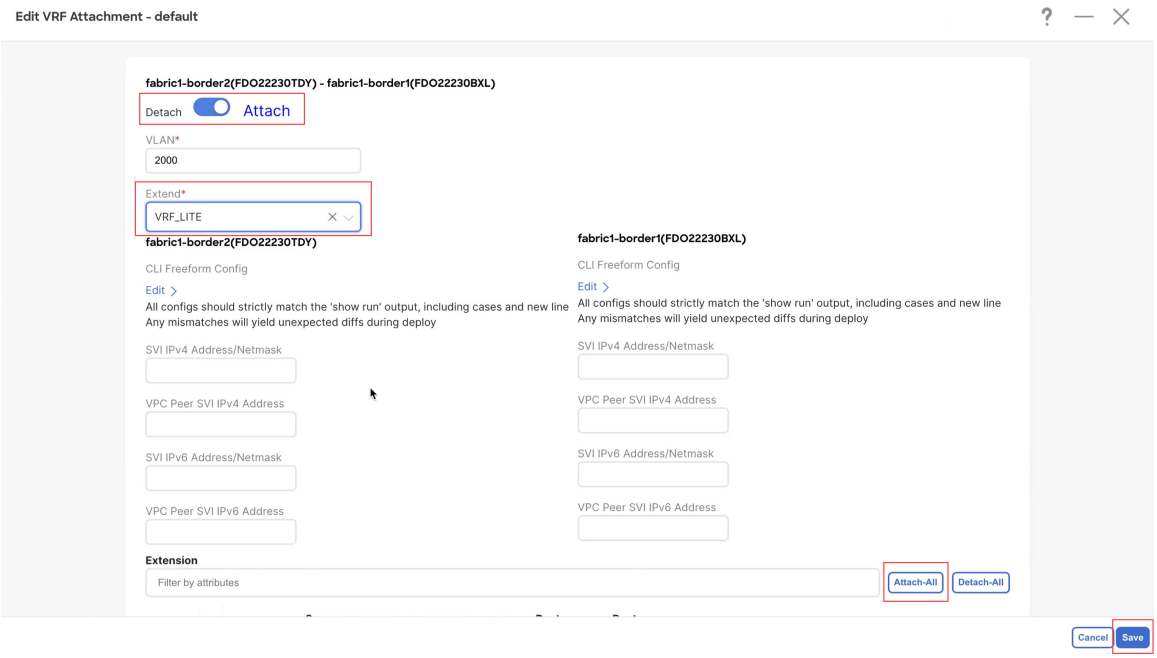
The VRF attachments for this VRF are displayed.

Step 6 Click the boxes next to the VRF attachments that you want to extend using VRF Lite and click **Actions** > **Edit**.



Step 7 In the **Edit VRF Attachment** page, make the necessary configurations to extend the VRF attachments using VRF Lite to attach the Aggregation switches to the Core switche.

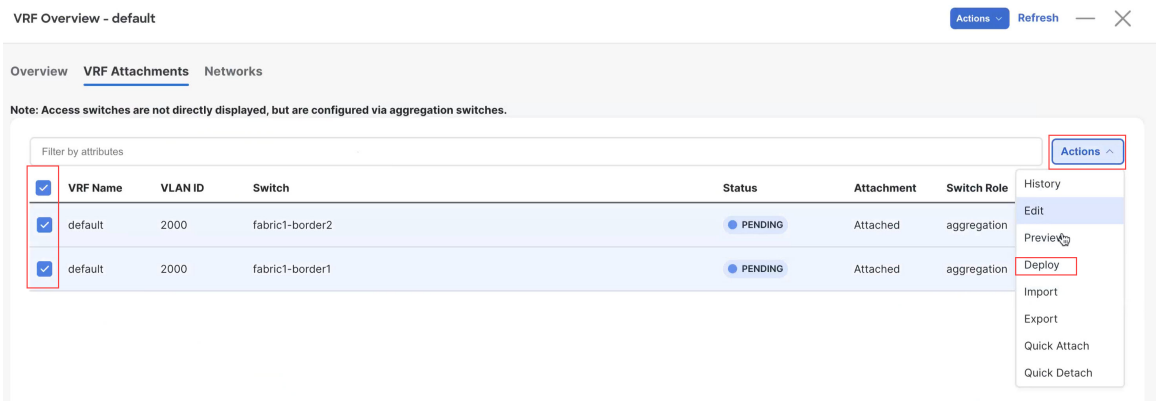
- Flip the switch to **Attach** at the top of the page.
- In the **Extend** field, choose the **VRF_LITE** option.
- Click **Attach-All**.



Click **Save** once you have completed the necessary configurations in the **Edit VRF Attachment** page. You are returned to the **VRF Attachments** page.

Step 8 Deploy the configuration using either of the two methods described below.

- In the **VRF Attachments** page, select the two VRF attachments and click **Actions** > **Deploy**.

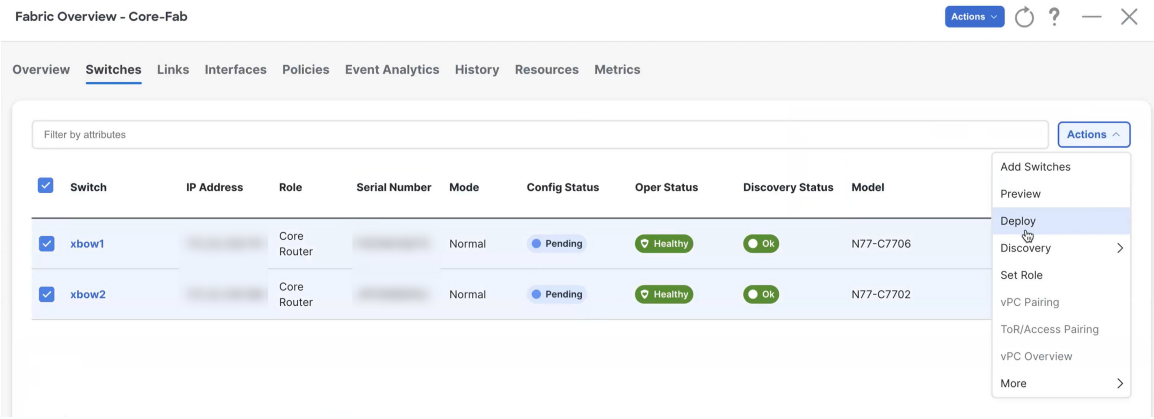


- In the **Fabric Overview** page for the Enhanced Classic LAN fabric, click the **Switches** tab and click the boxes next to the two Aggregation switches, then click **Actions** > **Deploy**.

Step 9 If you created a External Connectivity Network fabric for the core tier using the procedures in Configure the External Connectivity Network Fabric, on page 31, navigate back to **LAN** > **Fabrics** and double-click on the External Connectivity Network fabric.

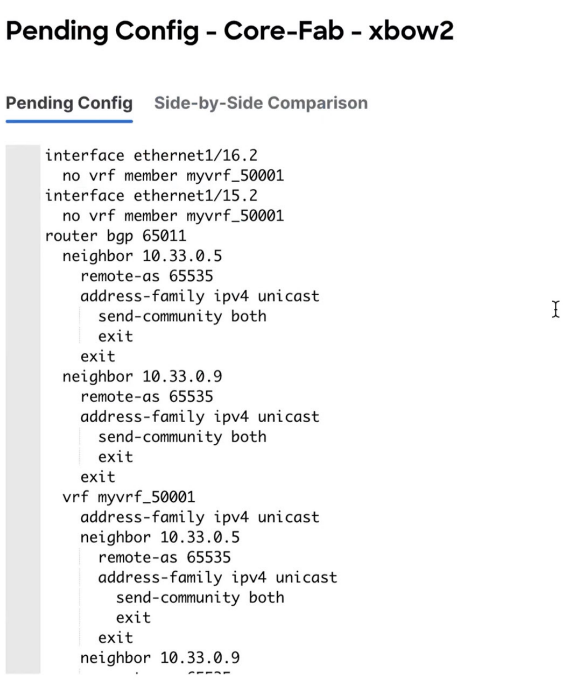
You must perform the same operations on the External Connectivity Network fabric to enable the pending configurations for VRF Lite to also be pushed to the Core switches.

Step 10 In the **Fabric Overview** page for the External Connectivity Network fabric, click the **Switches** tab and click the box next to the Core switch, then click **Actions** > **Deploy**.

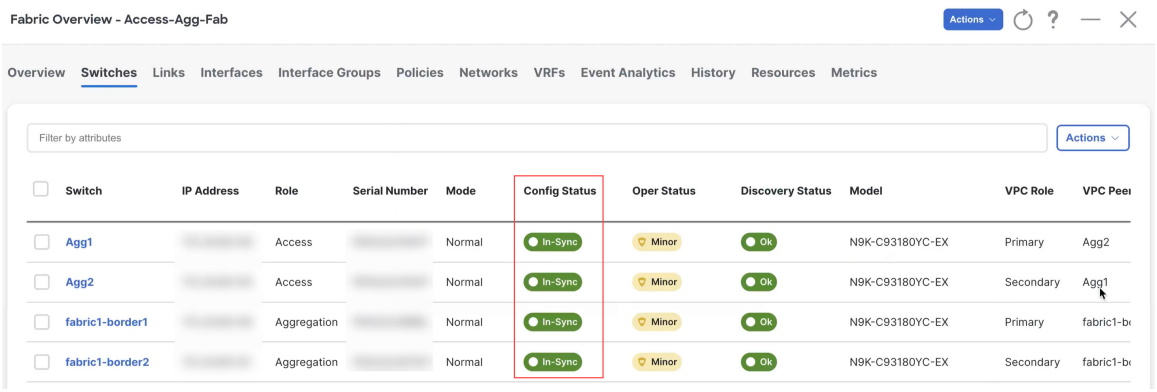


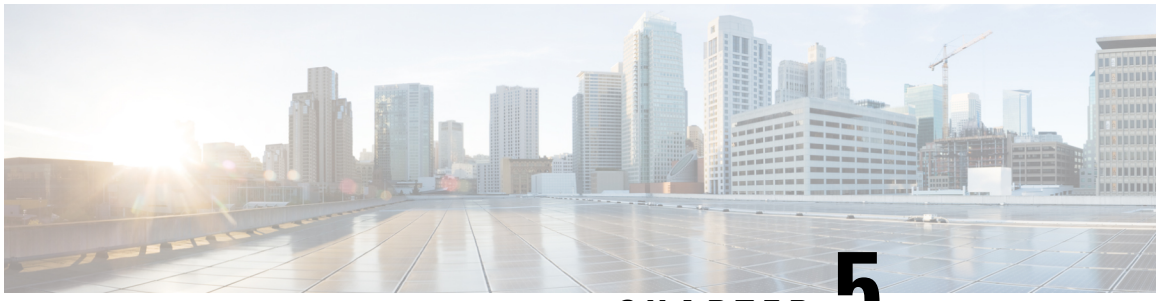
Step 11 Preview the configuration updates as the deployment process progresses.

You can click on the blue link in the **Pending Config** column to get additional information on the changes that are being configured for the Core switch.



Step 12 When the recalculation process is completed, click **Deploy All** and verify that the status shown in the **Config Status** column shows as **In-Sync**.


Fabric Overview - Access-Agg-Fab
Actions
Overview
Switches
Links
Interfaces
Interface Groups
Policies
Networks
VRFs
Event Analytics
History
Resources
Metrics
Filter by attributes
Actions
Switch
IP Address
Role
Serial Number
Mode
Config Status
Oper Status
Discovery Status
Model
VPC Role
VPC Peer
Agg1
Access
Normal
In-Sync
Minor
Ok
N9K-C93180YC-EX
Primary
Agg2
Agg2
Access
Normal
In-Sync
Minor
Ok
N9K-C93180YC-EX
Secondary
Agg1
fabric1-border1
Aggregation
Normal
In-Sync
Minor
Ok
N9K-C93180YC-EX
Primary
fabric1-b
fabric1-border2
Aggregation
Normal
In-Sync
Minor
Ok
N9K-C93180YC-EX
Secondary
fabric1-b



CHAPTER 5

Day 2 Configurations

- Locating Information for Day 2 Configurations, on page 61

Locating Information for Day 2 Configurations

All of the maintenance and operational features shown below are supported for classic LAN networks in the same way that they're supported for standard VXLAN fabrics. The features shown below are agnostic to fabric type.

You can find additional information on each of the following features in the appropriate *Cisco NDFC-Fabric Controller Configuration Guide*, depending on your release. As a reference point, the following table describes where you might find the relevant information for each feature in the release 12.1.1e version of the *Cisco NDFC-Fabric Controller Configuration Guide*.

Feature	Location in Release 12.1.3 Document
Image Management: Upgrades, Downgrades, EPLDs, RPMs, SMUs	*Image Management*
Inventory View	*About Fabric Overview for LAN Operational Mode Setups*
Event Analytics	*Event Analytics*
Deployment History, Audit Logs	History, Viewing Audit History
Backup and Restore	*Backup and Restore*
Performance Metrics	Metrics
Programmable Reports	*Programmable Reports*
Virtual infrastructure (Virtual Machine Manager, Kubernetes, OpenStack) Visibility	*Virtual Infrastructure Manager*