

Managing a Brownfield VXLAN BGP EVPN Fabric

- Overview, on page 1
- Prerequisites, on page 1
- Guidelines and Limitations, on page 2
- Fabric Topology Overview, on page 3
- NDFC Brownfield Deployment Tasks, on page 4
- Verifying the Existing VXLAN BGP EVPN Fabric, on page 4
- Creating a VXLAN EVPN Fabric Using the **Easy_Fabric** Template, on page 7
- Adding Switches and Transitioning VXLAN Fabric Management to NDFC, on page 26
- Configuration Profiles Support for Brownfield Migration, on page 30
- Manually Adding PIM-BIDIR Configuration for Leaf or Spine Post Brownfield Migration , on page 30
- Migrating a VXLAN EVPN Multi-Site Fabric with Border Gateway Switches , on page 31

Overview

This use case shows how to migrate an existing VXLAN BGP EVPN fabric to Cisco NDFC. The transition involves migrating existing network configurations to Nexus Dashboard Fabric Controller.

Typically, your fabric would be created and managed through manual CLI configuration or custom automation scripts. Now, you can start managing the fabric through Nexus Dashboard Fabric Controller. After the migration, the fabric underlay and overlay networks will be managed by NDFC.

For information about MSD fabric migration, see *Migrating an MSD Fabric with Border Gateway Switches*.

Prerequisites

- NDFC-supported NX-OS software versions. For details, refer Cisco Nexus Dashboard Fabric Controller Release Notes.
- Underlay routing protocol is OSPF or IS-IS.
- The following fabric-wide loopback interface IDs must not overlap:
  - Routing loopback interface for IGP/BGP.
  - VTEP loopback ID
  - Underlay rendezvous point loopback ID if ASM is used for multicast replication.

- BGP configuration uses the ‘router-id’, which is the IP address of the routing loopback interface.
- If the iBGP peer template is configured, then it must be configured on the leaf switches and route reflectors. The template name that needs to be used between leaf and route reflector should be identical.
- The BGP route reflector and multicast rendezvous point (if applicable) functions are implemented on spine switches. Leaf switches do not support the functions.
- Familiarity with VXLAN BGP EVPN fabric concepts and functioning of the fabric from the Nexus Dashboard Fabric Controller perspective.
- Fabric switch nodes are operationally stable and functional and all fabric links are up.
- vPC switches and the peer links are up before the migration. Ensure that no configuration updates are in progress or changes pending.
- Create an inventory list of the switches in the fabric with their IP addresses and credentials. Nexus Dashboard Fabric Controller uses this information to connect to the switches.
- Shut down any other controller software you are using presently so that no further configuration changes are made to the VXLAN fabric. Alternatively, disconnect the network interfaces from the controller software (if any) so that no changes are allowed on the switches.
- The switch overlay configurations must have the mandatory configurations defined in the shipping NDFC Universal Overlay profiles. Additional network or VRF overlay related configurations found on the switches are preserved in the freeform configuration associated with the network or VRF NDFC entries.
- All the overlay network and VRF profile parameters such as VLAN name and route map name should be consistent across all devices in the fabric for the brownfield migration to be successful.

Guidelines and Limitations

- Brownfield import must be completed for the entire fabric by adding all the switches to the NDFC fabric.
- Configuring `terminal color` or variations of this configuration (such as `terminal color persist`) on switches can cause a brownfield import to fail.
- On the **Create Fabric** window, the **Advanced** > **Overlay Mode** fabric setting decides how the overlays will be migrated. If the default config-profile is set, then the VRF and Network overlay configuration profiles will be deployed to switches as part of the migration process. In addition, there will be diffs to remove some of the redundant overlay CLI configurations. These are non network impacting.
- From the **Overlay Mode** drop-down list, if CLI is set, then VRF and Network overlay configurations stay on the switch as-is with no or little changes to address any consistency differences.
- The brownfield import in NDFC supports the simplified NX-OS VXLAN EVPN configuration CLIs. For more information, see Cisco Nexus 9000 Series NX-OS VXLAN Configuration Guide, Release 10.2(x).
- The following features are unsupported.
  - Super Spine roles
  - ToR
  - eBGP underlay

    - Layer 3 port channel

- Take a backup of the switch configurations and save them before migration.
- No configuration changes (unless instructed to do so in this document) must be made to the switches until the migration is completed. Else, significant network issues can occur.
- Migration to Cisco Nexus Dashboard Fabric Controller is only supported for Cisco Nexus 9000 switches.
- The Border Spine and Border Gateway Spine roles are supported for the brownfield migration.
- First, note the guidelines for updating the settings. Then update each VXLAN fabric settings as explained below:
    - Some values (BGP AS Number, OSPF, etc) are considered as reference points to your existing fabric, and the values you enter must match the existing fabric values.
    - For some fields (such as IP address range, VXLAN ID range), the values that are auto-populated or entered in the settings are only used for future allocation. The existing fabric values are honored during migration.
    - Some fields relate to new functions that may not exist in your existing fabric (such as advertise-pip). Enable or disable it as per your need.
    - At a later point in time, after the fabric transition is complete, you can update settings if needed.

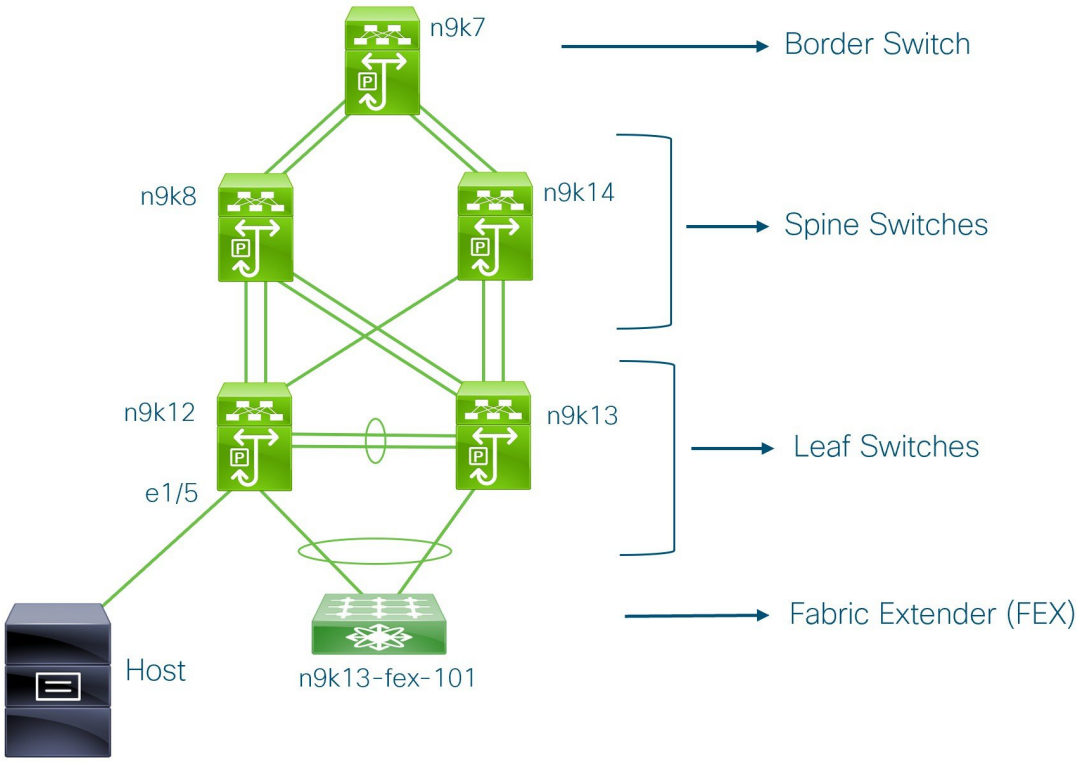
Fabric Topology Overview

This example use case uses the following hardware and software components:

- Five Cisco Nexus 9000 Series Switches
- One Fabric Extender or FEX
- One host

For information about the supported software images, see *Compatibility Matrix for Cisco NDFC*.

Before we start the transition of the existing fabric, let us see its topology.



You can see that there is a border switch, two spine switches, two leaf switches, and a Fabric Extender or FEX.

A host is connected to the n9k12 leaf switch through the interface Ethernet 1/5.

NDFC Brownfield Deployment Tasks

The following tasks are involved in a Brownfield migration:

1. Verifying the Existing VXLAN BGP EVPN Fabric, on page 4
2. Creating a VXLAN EVPN Fabric Using the Easy_Fabric Template
3. Adding Switches and Transitioning VXLAN Fabric Management to NDFC, on page 26

Verifying the Existing VXLAN BGP EVPN Fabric

Let us check the network connectivity of the **n9k12** switch from the console terminal.

Procedure

Step 1 Verify the Network Virtual Interface or NVE in the fabric.

```
n9k12# show nve vni summary
Codes: CP - Control Plane        DP - Data Plane
       UC - Unconfigured

Total CP VNIs: 84    [Up: 84, Down: 0]
Total DP VNIs: 0    [Up: 0, Down: 0]
```

There are 84 VNIs in the control plane and they are up. Before the Brownfield migration, make sure that all the VNIs are up.

Step 2 Check consistency and failures of vPC.

```
n9k12# show vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 2
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary
Number of vPCs configured         : 40
Peer Gateway                      : Enabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled, timer is off.(timeout = 300s)
Delay-restore status              : Timer is off.(timeout = 60s)
Delay-restore SVI status          : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router    : Disabled
.
.
.
```

Step 3 Check the EVPN neighbors of the **n9k-12** switch.

```
n9k12# show bgp l2vpn evpn summary
BGP summary information for VRF default, address family L2VPN EVPN
BGP router identifier 192.168.0.4, local AS number 65000
BGP table version is 637, L2VPN EVPN config peers 2, capable peers 2
243 network entries and 318 paths using 57348 bytes of memory
BGP attribute entries [234/37440], BGP AS path entries [0/0]
BGP community entries [0/0], BGP clusterlist entries [2/8]

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.0.0     4 65000     250      91      637    0    0 01:26:59 75
192.168.0.1     4 65000     221      63      637    0    0 00:57:22 75
```

You can see that there are two neighbors corresponding to the spine switches.

Note that the ASN is 65000.

Step 4 Verify the VRF information.

```
n9k12# show run vrf internet

!Command: show running-config vrf Internet
```

```
!Running configuration last done at: Fri Aug  9 01:38:02 2019
!Time: Fri Aug  9 02:48:03 2019

version 7.0(3)I7(6) Bios:version 07.59

interface Vlan347
  vrf member Internet

interface Vlan349
  vrf member Internet

interface Vlan3962
  vrf member Internet

interface Ethernet1/25
  vrf member Internet

interface Ethernet1/26
  vrf member Internet
vrf context Internet
  description Internet
  vni 16777210
  ip route 204.90.141.0/24 204.90.140.129 name LC-Networks
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
router ospf 300
  vrf Internet
    router-id 204.90.140.3
    redistribute direct route-map allow
    redistribute static route-map static-to-ospf
router bgp 65000
  vrf Internet
    address-family ipv4 unicast
      advertise l2vpn evpn
```

The VRF **Internet** is configured on this switch.

The host connected to the **n9k-12** switch is part of the VRF **Internet**.

You can see the VLANs associated with this VRF.

Specifically, the host is part of **Vlan349**.

Step 5 Verify the layer 3 interface information.

```
n9k12# show run interface vlan349

!Command: show running-config interface Vlan349
!Running configuration last done at: Fri Aug  9 01:38:02 2019
!Time: Fri Aug  9 02:49:27 2019

version 7.0(3)I7(6) Bios:version 07.59

interface Vlan349
  no shutdown
  vrf member Internet
  no ip redirects
  ip address 204.90.140.134/29
  no ipv6 redirects
  fabric forwarding mode anycast-gateway
```

Note that the IP address is **204.90.140.134**. This IP address is configured as the anycast gateway IP.

Step 6 Verify the physical interface information. This switch is connected to the Host through the interface Ethernet 1/5.

```
n9k12# show run interface ethernet1/5

!Command: show running-config interface Ethernet1/5
!Running configuration last done at: Fri Aug  9 01:38:02 2019
!Time: Fri Aug  9 02:50:05 2019

version 7.0(3)I7(6) Bios:version 07.59

interface Ethernet1/5
  description to host
  switchport mode trunk
  switchport trunk native vlan 349
  switchport trunk allowed vlan 349,800,815
  spanning-tree bpduguard enable
  mtu 9050
```

You can see that this interface is connected to the host and is configured with VLAN 349.

Step 7 Verify the connectivity from the host to the anycast gateway IP address.

```
host# ping 204.90.140.134 count unlimited interval 1
PING 204.90.140.134 (204.90.140.134): 56 data bytes
64 bytes from 204.90.140.134: icmp_seq=0 ttl=254 time=1.078 ms
64 bytes from 204.90.140.134: icmp_seq=1 ttl=254 time=1.129 ms
64 bytes from 204.90.140.134: icmp_seq=2 ttl=254 time=1.151 ms
64 bytes from 204.90.140.134: icmp_seq=3 ttl=254 time=1.162 ms
64 bytes from 204.90.140.134: icmp_seq=4 ttl=254 time=1.84 ms
64 bytes from 204.90.140.134: icmp_seq=5 ttl=254 time=1.258 ms
64 bytes from 204.90.140.134: icmp_seq=6 ttl=254 time=1.273 ms
64 bytes from 204.90.140.134: icmp_seq=7 ttl=254 time=1.143 ms
```

We let the ping command run in the background while we transition the existing brownfield fabric into Nexus Dashboard Fabric Controller.

Creating a VXLAN EVPN Fabric Using the Easy_Fabric Template

This topic describes how to create a new VXLAN EVPN fabric using the **Easy_Fabric** template and contains descriptions for the IPv4 underlay. For information about the IPv6 underlay, see IPv6 Underlay Support for Easy Fabric.

1. Navigate to the **LAN Fabrics** page:

   LAN > **Fabrics**

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

   The **Create Fabric** window appears.

3. Enter a unique name for the fabric in the **Fabric Name** field, then click **Choose Fabric**.

   A list of all available fabric templates are listed.

4. From the available list of fabric templates, choose the **Easy_Fabric** template, then click **Select**.

5. Enter the necessary field values to create a fabric.

The tabs and their fields in the screen are explained in the following sections. The overlay and underlay network parameters are included in these tabs.

Note If you're creating a standalone fabric as a potential member fabric of an MSD fabric (used for provisioning overlay networks for fabrics that are connected through EVPN Multi-Site technology), see Multi-Site Domain for VXLAN BGP EVPN Fabrics before creating the member fabric.

- General Parameters, on page 8
- Replication, on page 9
- VPC, on page 11
- Protocols, on page 12
- Advanced, on page 15
- Resources, on page 19
- Manageability, on page 21
- Bootstrap, on page 22
- Configuration Backup, on page 24
- Flow Monitor, on page 24

6. When you have completed the necessary configurations, click **Save**.
   - Click on the fabric to display a summary in the slide-in pane.
   - Click on the Launch icon to display the Fabric Overview.

General Parameters

The **General Parameters** tab is displayed by default. The fields in this tab are described in the following table.

Field	Description
BGP ASN	Enter the BGP AS number the fabric is associated with. This must be same as existing fabric.
Enable IPv6 Underlay	Enable the IPv6 underlay feature. For information, see IPv6 Underlay Support for Easy Fabric.
Enable IPv6 Link-Local Address	Enables the IPv6 Link-Local address.
Fabric Interface Numbering	Specifies whether you want to use point-to-point (**p2p**) or unnumbered networks.
Underlay Subnet IP Mask	Specifies the subnet mask for the fabric interface IP addresses.
Underlay Subnet IPv6 Mask	Specifies the subnet mask for the fabric interface IPv6 addresses.

Field	Description
Underlay Routing Protocol	The IGP used in the fabric, OSPF, or IS-IS.
Route-Reflectors (RRs)	The number of spine switches that are used as route reflectors for transporting BGP traffic. Choose 2 or 4 from the drop-down box. The default value is 2. To deploy spine devices as RRs, Nexus Dashboard Fabric Controller sorts the spine devices based on their serial numbers, and designates two or four spine devices as RRs. If you add more spine devices, existing RR configuration won't change. *Increasing the count* – You can increase the route reflectors from two to four at any point in time. Configurations are automatically generated on the other two spine devices designated as RRs. *Decreasing the count* – When you reduce four route reflectors to two, remove the unneeded route reflector devices from the fabric. Follow these steps to reduce the count from 4 to 2. **1.** Change the value in the drop-down box to 2. **2.** Identify the spine switches designated as route reflectors. An instance of the **rr_state** policy is applied on the spine switch if it's a route reflector. To find out if the policy is applied on the switch, right-click the switch, and choose **View/edit policies**. In the View/Edit Policies screen, search **rr_state** in the **Template** field. It is displayed on the screen. **3.** Delete the unneeded spine devices from the fabric (right-click the spine switch icon and choose **Discovery > Remove from fabric**). If you delete existing RR devices, the next available spine switch is selected as the replacement RR. **4.** Click **Deploy Config** in the fabric topology window. You can preselect RRs and RPs before performing the first **Save & Deploy** operation. For more information, see *Preselecting Switches as Route-Reflectors and Rendezvous-Points*.
Anycast Gateway MAC	Specifies the anycast gateway MAC address.
Enable Performance Monitoring	Check the check box to enable performance monitoring. Ensure that you do not clear interface counters from the Command Line Interface of the switches. Clearing interface counters can cause the Performance Monitor to display incorrect data for traffic utilization. If you must clear the counters and the switch has both `clear counters` and `clear counters snmp` commands (not all switches have the `clear counters snmp` command), ensure that you run both the main and the SNMP commands simultaneously. For example, you must run the `clear counters interface ethernet slot/port` command followed by the `clear counters interface ethernet slot/port snmp` command. This can lead to a one time spike.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Replication

The fields in the **Replication** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
Replication Mode	The mode of replication that is used in the fabric for BUM (Broadcast, Unknown Unicast, Multicast) traffic. The choices are Ingress Replication or Multicast. When you choose Ingress replication, the multicast related fields get disabled. You can change the fabric setting from one mode to the other, if no overlay profile exists for the fabric.
Multicast Group Subnet	IP address prefix used for multicast communication. A unique IP address is allocated from this group for each overlay network. The replication mode change isn't allowed if a policy template instance is created for the current mode. For example, if a multicast related policy is created and deployed, you can't change the mode to Ingress.
Enable Tenant Routed Multicast (TRM)	Check the check box to enable Tenant Routed Multicast (TRM) that allows overlay multicast traffic to be supported over EVPN/MVPN in the VXLAN BGP EVPN fabric.
Default MDT Address for TRM VRFs	The multicast address for Tenant Routed Multicast traffic is populated. By default, this address is from the IP prefix specified in the **Multicast Group Subnet** field. When you update either field, ensure that the TRM address is chosen from the IP prefix specified in **Multicast Group Subnet**. For more information, see Overview of Tenant Routed Multicast.
Rendezvous-Points	Enter the number of spine switches acting as rendezvous points.
RP mode	Choose from the two supported multicast modes of replication, ASM (for Any-Source Multicast [ASM]) or BiDir (for Bidirectional PIM [BIDIR-PIM]). When you choose ASM, the BiDir related fields aren't enabled. When you choose BiDir, the BiDir related fields are enabled. **Note** BIDIR-PIM is supported on Cisco's Cloud Scale Family platforms 9300-EX and 9300-FX/FX2, and software release 9.2(1) onwards. When you create a new VRF for the fabric overlay, this address is populated in the **Underlay Multicast Address** field, in the **Advanced** tab.
Underlay RP Loopback ID	The loopback ID used for the rendezvous point (RP), for multicast protocol peering purposes in the fabric underlay.
Underlay Primary RP Loopback ID	Enabled if you choose BIDIR-PIM as the multicast mode of replication. The primary loopback ID used for the phantom RP, for multicast protocol peering purposes in the fabric underlay.
Underlay Backup RP Loopback ID	Enabled if you choose BIDIR-PIM as the multicast mode of replication. The secondary loopback ID used for the phantom RP, for multicast protocol peering purposes in the fabric underlay.
Underlay Second Backup RP Loopback Id	Used for the second fallback Bidir-PIM Phantom RP.
Underlay Third Backup RP Loopback Id	Used for the third fallback Bidir-PIM Phantom RP.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

VPC

The fields in the **VPC** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
vPC Peer Link VLAN	VLAN used for the vPC peer link SVI.
Make vPC Peer Link VLAN as Native VLAN	Enables vPC peer link VLAN as Native VLAN.
vPC Peer Keep Alive option	Choose the management or loopback option. If you want to use IP addresses assigned to the management port and the management VRF, choose management. If you use IP addresses assigned to loopback interfaces (and a non-management VRF), choose loopback. If you use IPv6 addresses, you must use loopback IDs.
vPC Auto Recovery Time	Specifies the vPC auto recovery time-out period in seconds.
vPC Delay Restore Time	Specifies the vPC delay restore period in seconds.
vPC Peer Link Port Channel ID	Specifies the Port Channel ID for a vPC Peer Link. By default, the value in this field is 500.
vPC IPv6 ND Synchronize	Enables IPv6 Neighbor Discovery synchronization between vPC switches. The check box is enabled by default. Uncheck the check box to disable the function.
vPC advertise-pip	Select the check box to enable the Advertise PIP feature. You can enable the advertise PIP feature on a specific vPC as well. .
Enable the same vPC Domain Id for all vPC Pairs	Enable the same vPC Domain ID for all vPC pairs. When you select this field, the **vPC Domain Id** field is editable.
vPC Domain Id	Specifies the vPC domain ID to be used on all vPC pairs.
vPC Domain Id Range	Specifies the vPC Domain Id range to use for new pairings.
Enable QoS for Fabric vPC-Peering	Enable QoS on spines for guaranteed delivery of vPC Fabric Peering communication. . **Note** QoS for vPC fabric peering and queuing policies options in fabric settings are mutually exclusive.
QoS Policy Name	Specifies QoS policy name that should be same on all fabric vPC peering spines. The default name is **spine_qos_for_fabric_vpc_peering**.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Protocols

The fields in the **Protocols** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
Underlay Routing Loopback Id	The loopback interface ID is populated as 0 since loopback0 is usually used for fabric underlay IGP peering purposes.
Underlay VTEP Loopback Id	The loopback interface ID is populated as 1 since loopback1 is used for the VTEP peering purposes.
Underlay Anycast Loopback Id	The loopback interface ID is greyed out and used for vPC Peering in VXLANv6 Fabrics only.
Underlay Routing Protocol Tag	The tag defining the type of network.
OSPF Area ID	The OSPF area ID, if OSPF is used as the IGP within the fabric. **Note** The OSPF or IS-IS authentication fields are enabled based on your selection in the **Underlay Routing Protocol** field in the **General** tab.
Enable OSPF Authentication	Select the check box to enable OSPF authentication. Deselect the check box to disable it. If you enable this field, the OSPF Authentication Key ID and OSPF Authentication Key fields get enabled.
OSPF Authentication Key ID	The Key ID is populated.
OSPF Authentication Key	The OSPF authentication key must be the 3DES key from the switch. **Note** Plain text passwords are not supported. Log in to the switch, retrieve the encrypted key and enter it in this field. Refer, *Retrieving the Authentication Key* section for details.
IS-IS Level	Select the IS-IS level from this drop-down list.
Enable IS-IS Network Point-to-Point	Enables network point-to-point on fabric interfaces which are numbered.
Enable IS-IS Authentication	Select the check box to enable IS-IS authentication. Deselect the check box to disable it. If you enable this field, the IS-IS authentication fields are enabled.
IS-IS Authentication Keychain Name	Enter the Keychain name, such as CiscoisisAuth.
IS-IS Authentication Key ID	The Key ID is populated.
IS-IS Authentication Key	Enter the Cisco Type 7 encrypted key. **Note** Plain text passwords are not supported. Log in to the switch, retrieve the encrypted key and enter it in this field. Refer the Retrieving the Authentication Key section for details.

Field	Description
Set IS-IS Overload Bit	When enabled, set the overload bit for an elapsed time after a reload.
IS-IS Overload Bit Elapsed Time	Allows you to clear the overload bit after an elapsed time in seconds.
Enable BGP Authentication	Select the check box to enable BGP authentication. Deselect the check box to disable it. If you enable this field, the BGP Authentication Key Encryption Type and BGP Authentication Key fields are enabled. **Note** If you enable BGP authentication using this field, leave the iBGP Peer-Template Config field blank to avoid duplicate configuration.
BGP Authentication Key Encryption Type	Choose the 3 for 3DES encryption type, or 7 for Cisco encryption type.
BGP Authentication Key	Enter the encrypted key based on the encryption type. **Note** Plain text passwords are not supported. Log in to the switch, retrieve the encrypted key and enter it in the BGP Authentication Key field. Refer the Retrieving the Authentication Key section for details.
Enable PIM Hello Authentication	Select this check box to enable PIM hello authentication on all the intra-fabric interfaces of the switches in a fabric. This check box is editable only for the Multicast replication mode. Note this check box is valid only for the IPv4 underlay.
PIM Hello Authentication Key	Specifies the PIM hello authentication key. For more information, see Retrieving PIM Hello Authentication Key. To retrieve the PIM Hello Authentication Key, perform the following steps: **1.** SSH into the switch. **2.** On an unused switch interface, enable the following: `switch(config)# interface e1/32` `switch(config-if)# ip pim hello-authentication ah-md5 pimHelloPassword` In this example, **pimHelloPassword** is the cleartext password that has been used. **3.** Enter the **show run interface** command to retrieve the PIM hello authentication key. `switch(config-if)# show run interface e1/32 | grep pim` `ip pim sparse-mode` `ip pim hello-authentication ah-md5 3 d34e6c5abc7fecf1caa3b588b09078e0` In this example, **d34e6c5abc7fecf1caa3b588b09078e0** is the PIM hello authentication key that should be specified in the fabric settings.

Field	Description
Enable BFD	Check the check box to enable **feature bfd** on all switches in the fabric. This feature is valid only on IPv4 underlay and the scope is within a fabric. BFD within a fabric is supported natively. The BFD feature is disabled by default in the Fabric Settings. If enabled, BFD is enabled for the underlay protocols with the default settings. Any custom required BFD configurations must be deployed via the per switch freeform or per interface freeform policies. The following config is pushed after you select the **Enable BFD** check box: `feature bfd` For information about BFD feature compatibility, refer your respective platform documentation and for information about the supported software images, see *Compatibility Matrix for Cisco Nexus Dashboard Fabric Controller*.
Enable BFD for iBGP	Check the check box to enable BFD for the iBGP neighbor. This option is disabled by default.
Enable BFD for OSPF	Check the check box to enable BFD for the OSPF underlay instance. This option is disabled by default, and it is grayed out if the link state protocol is ISIS.
Enable BFD for ISIS	Check the check box to enable BFD for the ISIS underlay instance. This option is disabled by default, and it is grayed out if the link state protocol is OSPF.
Enable BFD for PIM	Check the check box to enable BFD for PIM. This option is disabled by default, and it is be grayed out if the replication mode is Ingress. Following are examples of the BFD global policies: `router ospf <ospf tag>` `  bfd` `router isis <isis tag>` `  address-family ipv4 unicast` `  bfd` `ip pim bfd` `router bgp <bgp asn>` `  neighbor <neighbor ip>` `  bfd`
Enable BFD Authentication	Check the check box to enable BFD authentication. If you enable this field, the **BFD Authentication Key ID** and **BFD Authentication Key** fields are editable. **Note** BFD Authentication is not supported when the **Fabric Interface Numbering** field under the **General** tab is set to **unnumbered**. The BFD authentication fields will be grayed out automatically. BFD authentication is valid for only for P2P interfaces.
BFD Authentication Key ID	Specifies the BFD authentication key ID for the interface authentication. The default value is 100.
BFD Authentication Key	Specifies the BFD authentication key. For information about how to retrieve the BFD authentication parameters. .

Field	Description
iBGP Peer-Template Config	Add iBGP peer template configurations on the leaf switches to establish an iBGP session between the leaf switch and route reflector. If you use BGP templates, add the authentication configuration within the template and uncheck the Enable BGP Authentication check box to avoid duplicate configuration. In the sample configuration, the 3DES password is displayed after password 3. `router bgp 65000` `  password 3 sd8478fswerdfw3434fsw4f4w34sdsd8478fswerdfw3434fsw4f4w` The following fields can be used to specify different configurations: • **iBGP Peer-Template Config** – Specifies the config used for RR and spines with border role. • **Leaf/Border/Border Gateway iBGP Peer-Template Config** – Specifies the config used for leaf, border, or border gateway. If this field is empty, the peer template defined in **iBGP Peer-Template Config** is used on all BGP enabled devices (RRs, leafs, border, or border gateway roles). In a brownfield migration, if the spine and leaf use different peer template names, both **iBGP Peer-Template Config** and **Leaf/Border/Border Gateway iBGP Peer-Template Config** fields need to be set according to the switch config. If spine and leaf use the same peer template name and content (except for the "route-reflector-client" CLI), only **iBGP Peer-Template Config** field in fabric setting needs to be set. If the fabric settings on iBGP peer templates do not match the existing switch configuration, an error message is generated and the migration will not proceed.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Advanced

The fields in the **Advanced** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
VRF Template	Specifies the VRF template for creating VRFs.
Network Template	Specifies the network template for creating networks.
VRF Extension Template	Specifies the VRF extension template for enabling VRF extension to other fabrics.
Network Extension Template	Specifies the network extension template for extending networks to other fabrics.
Overlay Mode	VRF/Network configuration using config-profile or CLI, default is config-profile. For more information, see Overlay Mode.
Site ID	The ID for this fabric if you are moving this fabric within an MSD. The site ID is mandatory for a member fabric to be a part of an MSD. Each member fabric of an MSD has a unique site ID for identification.

Field	Description
Intra Fabric Interface MTU	Specifies the MTU for the intra fabric interface. This value should be an even number.
Layer 2 Host Interface MTU	Specifies the MTU for the layer 2 host interface. This value should be an even number.
Unshut Host Interfaces by Default	Check this check box to unshut the host interfaces by default.
Power Supply Mode	Choose the appropriate power supply mode.
CoPP Profile	Choose the appropriate Control Plane Policing (CoPP) profile policy for the fabric. By default, the strict option is populated.
VTEP HoldDown Time	Specifies the NVE source interface hold down time.
Brownfield Overlay Network Name Format	Enter the format to be used to build the overlay network name during a brownfield import or migration. The network name should not contain any white spaces or special characters except underscore (_) and hyphen (-). The network name must not be changed once the brownfield migration has been initiated. See the *Creating Networks for the Standalone Fabric* section for the naming convention of the network name. The syntax is **[<string> | $$VLAN_ID$$] $$VNI$$ [<string>| $$VLAN_ID$$]** and the default value is **Auto_Net_VNI$$VNI$$_VLAN$$VLAN_ID$$**. When you create networks, the name is generated according to the syntax you specify. The following list describes the variables in the syntax: • **$$VNI$$**: Specifies the network VNI ID found in the switch configuration. This is a mandatory keyword required to create unique network names. • **$$VLAN_ID$$**: Specifies the VLAN ID associated with the network. VLAN ID is specific to switches, hence Nexus Dashboard Fabric Controller picks the VLAN ID from one of the switches, where the network is found, randomly and use it in the name. We recommend not to use this unless the VLAN ID is consistent across the fabric for the VNI. • **<string>**: This variable is optional and you can enter any number of alphanumeric characters that meet the network name guidelines. An example overlay network name: Site_VNI12345_VLAN1234 **Note** Ignore this field for greenfield deployments. The **Brownfield Overlay Network Name Format** applies for the following brownfield imports: • CLI-based overlays • Configuration profile-based overlay
Enable CDP for Bootstrapped Switch	Enables CDP on management (mgmt0) interface for bootstrapped switch. By default, for bootstrapped switches, CDP is disabled on the mgmt0 interface.

Field	Description
Enable VXLAN OAM	Enables the VXLAM OAM functionality for devices in the fabric. This is enabled by default. Uncheck the check box to disable VXLAN OAM function. If you want to enable the VXLAN OAM function on specific switches and disable on other switches in the fabric, you can use freeform configurations to enable OAM and disable OAM in the fabric settings. **Note** The VXLAN OAM feature in Cisco Nexus Dashboard Fabric Controller is only supported on a single fabric or site.
Enable Tenant DHCP	Check the check box to enable feature dhcp and associated configurations globally on all switches in the fabric. This is a pre-requisite for support of DHCP for overlay networks that are part of the tenant VRFs. **Note** Ensure that **Enable Tenant DHCP** is enabled before enabling DHCP-related parameters in the overlay profiles.
Enable NX-API	Specifies enabling of NX-API on HTTPS. This check box is checked by default.
Enable NX-API on HTTP Port	Specifies enabling of NX-API on HTTP. Enable this check box and the **Enable NX-API** check box to use HTTP. This check box is checked by default. If you uncheck this check box, the applications that use NX-API and supported by Cisco Nexus Dashboard Fabric Controller, such as Endpoint Locator (EPL), Layer 4-Layer 7 services (L4-L7 services), VXLAN OAM, and so on, start using the HTTPS instead of HTTP. **Note** If you check the **Enable NX-API** check box and the **Enable NX-API on HTTP** check box, applications use HTTP.
Enable Policy-Based Routing (PBR)	Check this check box to enable routing of packets based on the specified policy. Starting with Cisco NX-OS Release 7.0(3)I7(1) and later releases, this feature works on Cisco Nexus 9000 Series switches with Nexus 9000 Cloud Scale (Tahoe) ASICs. This feature is used along with the Layer 4-Layer 7 service workflow. For information on Layer 4-Layer 7 service, refer the *Layer 4-Layer 7 Service* chapter.
Enable Strict Config Compliance	Enable the Strict Config Compliance feature by selecting this check box. It enables bi-directional compliance checks to flag additional configs in the running config that are not in the intent/expected config. By default, this feature is disabled.
Enable AAA IP Authorization	Enables AAA IP authorization, when IP Authorization is enabled in the remote authentication server. This is required to support Nexus Dashboard Fabric Controller in scenarios where customers have strict control of which IP addresses can have access to the switches.
Enable NDFC as Trap Host	Select this check box to enable Nexus Dashboard Fabric Controller as an SNMP trap destination. Typically, for a native HA Nexus Dashboard Fabric Controller deployment, the eth1 VIP IP address will be configured as SNMP trap destination on the switches. By default, this check box is enabled.
Anycast Border Gateway advertise-pip	Enables to advertise Anycast Border Gateway PIP as VTEP. Effective on MSD fabric 'Recalculate Config'.

Field	Description
Greenfield Cleanup Option	Enable the switch cleanup option for switches imported into Nexus Dashboard Fabric Controller with Preserve-Config=No, without a switch reload. This option is typically recommended only for the fabric environments with Cisco Nexus 9000v Switches to improve on the switch clean up time. The recommended option for Greenfield deployment is to employ Bootstrap or switch cleanup with a reboot. In other words, this option should be unchecked.
Enable Precision Time Protocol (PTP)	Enables PTP across a fabric. When you check this check box, PTP is enabled globally and on core-facing interfaces. Additionally, the **PTP Source Loopback Id** and **PTP Domain Id** fields are editable. For more information, see Precision Time Protocol for Easy Fabric.
PTP Source Loopback Id	Specifies the loopback interface ID Loopback that is used as the Source IP Address for all PTP packets. The valid values range from 0 to 1023. The PTP loopback ID cannot be the same as RP, Phantom RP, NVE, or MPLS loopback ID. Otherwise, an error will be generated. The PTP loopback ID can be the same as BGP loopback or user-defined loopback which is created from Nexus Dashboard Fabric Controller. If the PTP loopback ID is not found during **Deploy Config**, the following error is generated: Loopback interface to use for PTP source IP is not found. Create PTP loopback interface on all the devices to enable PTP feature.
PTP Domain Id	Specifies the PTP domain ID on a single network. The valid values range from 0 to 127.
Enable MPLS Handoff	Check the check box to enable the MPLS Handoff feature. For more information, see the MPLS SR and LDP Handoff chapter in External/WAN Layer 3 Connectivity for VXLAN BGP EVPN Fabrics.
Underlay MPLS Loopback Id	Specifies the underlay MPLS loopback ID. The default value is 101.
Enable TCAM Allocation	TCAM commands are automatically generated for VXLAN and vPC Fabric Peering when enabled.
Enable Default Queuing Policies	Check this check box to apply QoS policies on all the switches in this fabric. To remove the QoS policies that you applied on all the switches, uncheck this check box, update all the configurations to remove the references to the policies, and save and deploy. Pre-defined QoS configurations are included that can be used for various Cisco Nexus 9000 Series Switches. When you check this check box, the appropriate QoS configurations are pushed to the switches in the fabric. The system queuing is updated when configurations are deployed to the switches. You can perform the interface marking with defined queuing policies, if required, by adding the required configuration to the per interface freeform block. Review the actual queuing policies by opening the policy file in the template editor. From Cisco Nexus Dashboard Fabric Controller Web UI, choose **Operations > Templates**. Search for the queuing policies by the policy file name, for example, **queuing_policy_default_8q_cloudscale**. Choose the file. From the **Actions** drop-down list, select **Edit template content** to edit the policy. See the *Cisco Nexus 9000 Series NX-OS Quality of Service Configuration Guide* for platform specific details.
N9K Cloud Scale Platform Queuing Policy	Choose the queuing policy from the drop-down list to be applied to all Cisco Nexus 9200 Series Switches and the Cisco Nexus 9000 Series Switches that ends with EX, FX, and FX2 in the fabric. The valid values are **queuing_policy_default_4q_cloudscale** and **queuing_policy_default_8q_cloudscale**. Use the **queuing_policy_default_4q_cloudscale** policy for FEXes. You can change from the **queuing_policy_default_4q_cloudscale** policy to the **queuing_policy_default_8q_cloudscale** policy only when FEXes are offline.

Field	Description
N9K R-Series Platform Queuing Policy	Choose the queuing policy from the drop-down list to be applied to all Cisco Nexus switches that ends with R in the fabric. The valid value is **queuing_policy_default_r_series**.
Other N9K Platform Queuing Policy	Choose the queuing policy from the drop-down list to be applied to all other switches in the fabric other than the switches mentioned in the above two options. The valid value is **queuing_policy_default_other**.
Enable MACsec	Enables MACsec for the fabric. For more information, see Enabling MACsec. *Freeform CLIs* - Fabric level freeform CLIs can be added while creating or editing a fabric. They are applicable to switches across the fabric. You must add the configurations as displayed in the running configuration, without indentation. Switch level freeform configurations should be added via the switch freeform on NDFC. For more information, see Enabling Freeform Configurations on Fabric Switches.
Leaf Freeform Config	Add CLIs that should be added to switches that have the *Leaf*, *Border*, and *Border Gateway* roles.
Spine Freeform Config	Add CLIs that should be added to switches with a *Spine*, *Border Spine*, *Border Gateway Spine*, and *Super Spine* roles.
Intra-fabric Links Additional Config	Add CLIs that should be added to the intra-fabric links.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Resources

The fields in the **Resources** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
Manual Underlay IP Address Allocation	*Do not* check this check box if you are transitioning your VXLAN fabric management to Nexus Dashboard Fabric Controller. • By default, Nexus Dashboard Fabric Controller allocates the underlay IP address resources (for loopbacks, fabric interfaces, etc) dynamically from the defined pools. If you check the check box, the allocation scheme switches to static, and some of the dynamic IP address range fields are disabled. • For static allocation, the underlay IP address resources must be populated into the Resource Manager (RM) using REST APIs. • The Underlay RP Loopback IP Range field stays enabled if BIDIR-PIM function is chosen for multicast replication. • Changing from static to dynamic allocation keeps the current IP resource usage intact. Only future IP address allocation requests are taken from dynamic pools.
Underlay Routing Loopback IP Range	Specifies loopback IP addresses for the protocol peering.

Field	Description
Underlay VTEP Loopback IP Range	Specifies loopback IP addresses for VTEPs.
Underlay RP Loopback IP Range	Specifies the anycast or phantom RP IP address range.
Underlay Subnet IP Range	IP addresses for underlay P2P routing traffic between interfaces.
Underlay MPLS Loopback IP Range	Specifies the underlay MPLS loopback IP address range. For eBGP between Border of Easy A and Easy B, Underlay routing loopback and Underlay MPLS loopback IP range must be a unique range. It should not overlap with IP ranges of the other fabrics, else VPNv4 peering will not come up.
Underlay Routing Loopback IPv6 Range	Specifies Loopback0 IPv6 Address Range
Underlay VTEP Loopback IPv6 Range	Specifies Loopback1 and Anycast Loopback IPv6 Address Range.
Underlay Subnet IPv6 Range	Specifies IPv6 Address range to assign Numbered and Peer Link SVI IPs.
BGP Router ID Range for IPv6 Underlay	Specifies BGP router ID range for IPv6 underlay.
Layer 2 VXLAN VNI Range	Specifies the overlay VXLAN VNI range for the fabric (min:1, max:16777214).
Layer 3 VXLAN VNI Range	Specifies the overlay VRF VNI range for the fabric (min:1, max:16777214).
Network VLAN Range	VLAN range for the per switch overlay network (min:2, max:4094).
VRF VLAN Range	VLAN range for the per switch overlay Layer 3 VRF (min:2, max:4094).
Subinterface Dot1q Range	Specifies the subinterface range when L3 sub interfaces are used.
VRF Lite Deployment	Specify the VRF Lite method for extending inter fabric connections. The VRF Lite Subnet IP Range field specifies resources reserved for IP address used for VRF Lite when VRF Lite IFCs are auto-created. If you select Back2Back&ToExternal, then VRF Lite IFCs are auto-created.
Auto Deploy for Peer	This check box is applicable for VRF Lite deployment. When you select this checkbox, auto-created VRF Lite IFCs will have the **Auto Generate Configuration for Peer** field in the **VRF Lite** tab set. To access VRF Lite IFC configuration, navigate to the **Links** tab, select the particular link, and then choose **Actions** > **Edit**. You can check or uncheck the check box when the **VRF Lite Deployment** field is not set to **Manual**. This configuration only affects the new auto-created IFCs and does not affect the existing IFCs. You can edit an auto-created IFC and check or uncheck the **Auto Generate Configuration for Peer** field. This setting takes priority always.

Field	Description
Auto Deploy Default VRF	When you select this check box, the **Auto Generate Configuration on default VRF** field is automatically enabled for auto-created VRF Lite IFCs. You can check or uncheck this check box when the **VRF Lite Deployment** field is not set to **Manual**. The **Auto Generate Configuration on default VRF** field when set, automatically configures the physical interface for the border device, and establishes an EBGP connection between the border device and the edge device or another border device in a different VXLAN EVPN fabric.
Auto Deploy Default VRF for Peer	When you select this check box, the **Auto Generate Configuration for NX-OS Peer on default VRF** field is automatically enabled for auto-created VRF Lite IFCs. You can check or uncheck this check box when the **VRF Lite Deployment** field is not set to **Manual**. The **Auto Generate Configuration for NX-OS Peer on default VRF** field when set, automatically configures the physical interface and the EBGP commands for the peer NX-OS switch. **Note** To access the **Auto Generate Configuration on default VRF** and **Auto Generate Configuration for NX-OS Peer on default VRF** fields for an IFC link, navigate to the **Links** tab, select the particular link and choose **Actions** > **Edit**.
Redistribute BGP Route-map Name	Defines the route map for redistributing the BGP routes in default VRF.
VRF Lite Subnet IP Range and **VRF Lite Subnet Mask**	These fields are populated with the DCI subnet details. Update the fields as needed. The values shown in your screen are automatically generated. If you want to update the IP address ranges, VXLAN Layer 2/Layer 3 network ID ranges or the VRF/Network VLAN ranges, ensure the following: **Note** When you update a range of values, ensure that it does not overlap with other ranges. You should only update one range of values at a time. If you want to update more than one range of values, do it in separate instances. For example, if you want to update L2 and L3 ranges, you should do the following. **1.** Update the L2 range and click **Save**. **2.** Click the **Edit Fabric** option again, update the L3 range and click **Save**.
Service Network VLAN Range	Specifies a VLAN range in the Service Network VLAN Range field. This is a per switch overlay service network VLAN range. The minimum allowed value is 2 and the maximum allowed value is 3967.
Route Map Sequence Number Range	Specifies the route map sequence number range. The minimum allowed value is 1 and the maximum allowed value is 65534.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Manageability

The fields in the **Manageability** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
DNS Server IPs	Specifies the comma separated list of IP addresses (v4/v6) of the DNS servers.
DNS Server VRFs	Specifies one VRF for all DNS servers or a comma separated list of VRFs, one per DNS server.
NTP Server IPs	Specifies comma separated list of IP addresses (v4/v6) of the NTP server.
NTP Server VRFs	Specifies one VRF for all NTP servers or a comma separated list of VRFs, one per NTP server.
Syslog Server IPs	Specifies the comma separated list of IP addresses (v4/v6) IP address of the syslog servers, if used.
Syslog Server Severity	Specifies the comma separated list of syslog severity values, one per syslog server. The minimum value is 0 and the maximum value is 7. To specify a higher severity, enter a higher number.
Syslog Server VRFs	Specifies one VRF for all syslog servers or a comma separated list of VRFs, one per syslog server.
AAA Freeform Config	Specifies the AAA freeform configurations. If AAA configurations are specified in the fabric settings, **switch_freeform** PTI with source as **UNDERLAY_AAA** and description as **AAA Configurations** will be created.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Bootstrap

The fields in the **Bootstrap** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
Enable Bootstrap	Select this check box to enable the bootstrap feature. Bootstrap allows easy day-0 import and bring-up of new devices into an existing fabric. Bootstrap leverages the NX-OS POAP functionality. After you enable bootstrap, you can enable the DHCP server for automatic IP address assignment using one of the following methods: • External DHCP Server: Enter information about the external DHCP server in the **Switch Mgmt Default Gateway** and **Switch Mgmt IP Subnet Prefix** fields. • Local DHCP Server: Enable the **Local DHCP Server** check box and enter details for the remaining mandatory fields.
Enable Local DHCP Server	Select this check box to initiate enabling of automatic IP address assignment through the local DHCP server. When you select this check box, the **DHCP Scope Start Address** and **DHCP Scope End Address** fields become editable. If you do not select this check box, Nexus Dashboard Fabric Controller uses the remote or external DHCP server for automatic IP address assignment.

Field	Description
DHCP Version	Select DHCPv4 or DHCPv6 from this drop-down list. When you select DHCPv4, the **Switch Mgmt IPv6 Subnet Prefix** field is disabled. If you select DHCPv6, the **Switch Mgmt IP Subnet Prefix** is disabled. **Note** Cisco Nexus 9000 and 3000 Series Switches support IPv6 POAP only when switches are either Layer-2 adjacent (eth1 or out-of-band subnet must be a /64) or they are L3 adjacent residing in some IPv6 /64 subnet. Subnet prefixes other than /64 are not supported.
DHCP Scope Start Address and **DHCP Scope End Address**	Specifies the first and last IP addresses of the IP address range to be used for the switch out of band POAP.
Switch Mgmt Default Gateway	Specifies the default gateway for the management VRF on the switch.
Switch Mgmt IP Subnet Prefix	Specifies the prefix for the Mgmt0 interface on the switch. The prefix should be between 8 and 30. *DHCP scope and management default gateway IP address specification* - If you specify the management default gateway IP address 10.0.1.1 and subnet mask 24, ensure that the DHCP scope is within the specified subnet, between 10.0.1.2 and 10.0.1.254.
Switch Mgmt IPv6 Subnet Prefix	Specifies the IPv6 prefix for the Mgmt0 interface on the switch. The prefix should be between 112 and 126. This field is editable if you enable IPv6 for DHCP.
Enable AAA Config	Select this check box to include AAA configurations from the Manageability tab as part of the device start-up config post bootstrap.
DHCPv4/DHCPv6 Multi Subnet Scope	Specifies the field to enter one subnet scope per line. This field is editable after you check the **Enable Local DHCP Server** check box. The format of the scope should be defined as: **DHCP Scope Start Address, DHCP Scope End Address, Switch Management Default Gateway, Switch Management Subnet Prefix** For example: 10.6.0.2, 10.6.0.9, 10.6.0.1, 24
Bootstrap Freeform Config	(Optional) Enter additional commands as needed. For example, if you require some additional configurations to be pushed to the device and be available post device bootstrap, they can be captured in this field, to save the desired intent. After the devices boot up, they will contain the configuration defined in the **Bootstrap Freeform Config** field. Copy-paste the running-config to a **freeform config** field with correct indentation, as seen in the running configuration on the NX-OS switches. The freeform config must match the running config. For more information, see Enabling Freeform Configurations on Fabric Switches.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Configuration Backup

The fields in the **Configuration Backup** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
Hourly Fabric Backup	Select the check box to enable an hourly backup of fabric configurations and the intent. The hourly backups are triggered during the first 10 minutes of the hour.
Scheduled Fabric Backup	Check the check box to enable a daily backup. This backup tracks changes in running configurations on the fabric devices that are not tracked by configuration compliance.
Scheduled Time	Specify the scheduled backup time in a 24-hour format. This field is enabled if you check the **Scheduled Fabric Backup** check box. Select both the check boxes to enable both back up processes. The backup process is initiated after you click **Save**. The scheduled backups are triggered exactly at the time you specify with a delay of up to two minutes. The scheduled backups are triggered regardless of the configuration deployment status. The number of fabric backups that will be retained on NDFC is decided by the **Settings** > **Server Settings** > **LAN Fabric** > **Maximum Backups per Fabric**. The number of archived files that can be retained is set in the **# Number of archived files per device to be retained:** field in the **Server Properties** window. **Note** To trigger an immediate backup, do the following: **1.** Choose **LAN > Topology**. **2.** Click within the specific fabric box. The fabric topology screen comes up. **3.** From the **Actions** pane at the left part of the screen, click **Re-Sync Fabric**. You can also initiate the fabric backup in the fabric topology window. Click **Backup Now** in the **Actions** pane.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Flow Monitor

The fields in the **Flow Monitor** tab are described in the following table. Most of the fields are automatically generated based on Cisco-recommended best practice configurations, but you can update the fields if needed.

Field	Description
Enable Netflow	Check this check box to enable Netflow on VTEPs for this Fabric. By default, Netflow is disabled. On Enable, NetFlow configuration will be applied to all VTEPS that support netflow. **Note** When Netflow is enabled on the fabric, you can choose not to have netflow on a particular switch by having a dummy no_netflow PTI. If netflow is not enabled at the fabric level, an error message is generated when you enable netflow at the interface, network, or vrf level. For information about Netflow support for Cisco NDFC, refer to Netflow Support.

In the **Netflow Exporter** area, click **Actions > Add** to add one or more Netflow exporters. This exporter is the receiver of the netflow data. The fields on this screen are:

- **Exporter Name** – Specifies the name of the exporter.
- **IP** – Specifies the IP address of the exporter.
- **VRF** – Specifies the VRF over which the exporter is routed.
- **Source Interface** – Enter the source interface name.
- **UDP Port** – Specifies the UDP port over which the netflow data is exported.

Click **Save** to configure the exporter. Click **Cancel** to discard. You can also choose an existing exporter and select **Actions > Edit** or **Actions > Delete** to perform relevant actions.

In the **Netflow Record** area, click **Actions > Add** to add one or more Netflow records. The fields on this screen are:

- **Record Name** – Specifies the name of the record.
- **Record Template** – Specifies the template for the record. Enter one of the record templates names. In Release 12.0.2, the following two record templates are available for use. You can create custom netflow record templates. Custom record templates saved in the template library are available for use here.
  - **netflow_ipv4_record** – to use the IPv4 record template.
  - **netflow_l2_record** – to use the Layer 2 record template.
- **Is Layer2 Record** – Check this check box if the record is for Layer2 netflow.

Click **Save** to configure the report. Click **Cancel** to discard. You can also choose an existing record and select **Actions > Edit** or **Actions > Delete** to perform relevant actions.

In the **Netflow Monitor** area, click **Actions > Add** to add one or more Netflow monitors. The fields on this screen are:

- **Monitor Name** – Specifies the name of the monitor.
- **Record Name** – Specifies the name of the record for the monitor.
- **Exporter1 Name** – Specifies the name of the exporter for the netflow monitor.
- **Exporter2 Name** – (optional) Specifies the name of the secondary exporter for the netflow monitor.

The record name and exporters referred to in each netflow monitor must be defined in "**Netflow Record**" and "**Netflow Exporter**".

Click **Save** to configure the monitor. Click **Cancel** to discard. You can also choose an existing monitor and select **Actions > Edit** or **Actions > Delete** to perform relevant actions.

What's next: Complete the configurations in another tab if necessary, or click **Save** when you have completed the necessary configurations for this fabric.

Adding Switches and Transitioning VXLAN Fabric Management to NDFC

Let us discover and add switches to the newly created fabric.

Procedure

Step 1 Double click on the newly created fabric to view the **Fabric Overview** screen.

Click on **Switches** tab.

Step 2 From the **Actions** drop-down list, select **Add Switches**.

The **Add Switches** window appears.

Similarly, you can add switches on **Topology** window. On Topology window, choose a fabric, right-click on a fabric and click **Add Switches**.

Step 3 On the **Add Switches - Fabric** screen, enter the **Seed Switch Details**.

Enter the IP address of the switch in the **Seed IP** field. Enter the username and password of the switches that you want to discover.

By default, the value in the **Max Hops** field is **2**. The switch with the specified IP address and the switches that are 2 hops from it will be populated after the discovery is complete.

Make sure to check the **Preserve Config** check box. This ensures that the current configuration of the switches will be retained.

Step 4 Click **Discover Switches**.

The switch with the specified IP address and switches up to two hops away (depending on the setting of Max Hops) from it are populated in the Scan Details section.

Step 5 Check the check box next to the switches that have to be imported into the fabric and click **Import into fabric**.

It is best practice to discover multiple switches at the same time in a single attempt. The switches must be cabled and connected to the NDFC server and the switch status must be manageable.

If switches are imported in multiple attempts, then please ensure that all the switches are added to the fabric before proceeding with the Brownfield import process.

Step 6 Click **Import into fabric**.

The switch discovery process is initiated. The **Progress** column displays progress for all the selected switches. It displays **done** for each switch after completion.

Note
You should not close the screen and try to import switches again until all selected switches are imported or an error message comes up.

If an error message comes up, close the screen. The fabric topology screen comes up. The error messages are displayed at the top-right part of the screen. Resolve the errors and initiate the import process again by clicking **Add Switches** in the **Actions** panel.

Step 7 After a successful import, the progress bar shows **Done** for all the switches. Click **Close**.

After closing the window, the fabric topology window comes up again. The switch is in Migration Mode now, and the Migration mode label is displayed on the switch icons.

At this point, you must not try to add Greenfield or *new* switches. Support is not available for adding new switches during the migration process. It might lead to undesirable consequences for your network. However, you can add a new switch after the migration process is complete.

Step 8 After all the network elements are discovered, they are displayed in the **Topology** window in a connected topology. Each switch is assigned the **Leaf** role by default.

The switch discovery process might fail for a few switches, and the Discovery Error message is displayed. However, such switches are still displayed in the fabric topology. You should remove such switches from the fabric (Right-click the switch icon and click **Discovery > Remove** from fabric), and import them again.

You should not proceed to the next step until all switches in the existing fabric are discovered in NDFC.

If you choose the Hierarchical layout for display (in the Actions panel), the topology automatically gets aligned as per role assignment, with the leaf switches at the bottom, the spine switches connected on top of them, and the border switches at the top.

Note
The supported roles for switches with the Cisco NX-OS Release 7.0(3)I4(8b) and 7.0(4)I4(x) images are Border Leaf, Border Spine, Leaf, and Spine

Step 9 Select the switch, click **Actions > Set Role**. On the Select Role screen, select **Border** and click **Select**.

Similarly, set the **Spine** role for the **n9k-14** and **n9k-8** spine switches.

Note
You need to manually create a vPC pairing when the L3 keep alive is configured on the switches. Otherwise, the vPC configuration is automatically picked up from the switches.

vPC Pairing - The vPC pairing must be done for switches where the Layer 3 vPC peer-keep alive is used. The vPC configuration is automatically picked up from the switches when the vPC peer keep alive is established through the management option. This pairing reflects in the GUI only after the migration is complete.

a. Right-click the switch icon and click vPC Pairing to set a vPC switch pair.

The Select vPC peer screen comes up. It lists potential vPC peer switches.

b. Select the appropriate switch and click OK. The fabric topology comes up again. The vPC pair is formed.

Note
Check if you have added all switches from the current fabric. If you have missed adding switches, add them now. Once you are certain that you have imported all existing switches, move to the next step, the Save and Deploy option.

Step 10 From the Fabric Overview **Actions** drop-down list, choose **Recalculate and Deploy**.

When you click **Recalculate and Deploy**, NDFC obtains switch configurations and populates the state of every switch from the current running configuration to the current expected configuration, which is the intended state maintained in NDFC.

If there are configuration mismatches, **Pending Config** column shows the number of lines of difference. Click on the Pending Config column to view the **Pending Config** and **Side-by-Side Comparison** of the running configuration. Click **Deploy** to apply the configurations.

After the migration of underlay and overlay networks, the **Deploy Configuration**screen comes up.

Note

- The brownfield migration requires best practices to be followed on the existing fabric such as maintain consistency of the overlay configurations.
- The Brownfield migration may take some time to complete since it involves collecting the running configuration from switches, build the NDFC configuration intent based on these, consistency checks etc.
- Any errors or inconsistencies that are found during the migration is reported in fabric errors. The switches continue to remain in the Migration mode. You should fix these errors and complete the migration again by clicking **Deploy** until no errors are reported.

Step 11 After the configurations are generated, review them by clicking the links in the **Preview Config** column.

We strongly recommend that you preview the configuration before proceeding to deploy it on the switches. Click the Preview Configuration column entry. The **Preview Config** screen comes up. It lists the pending configurations on the switch.

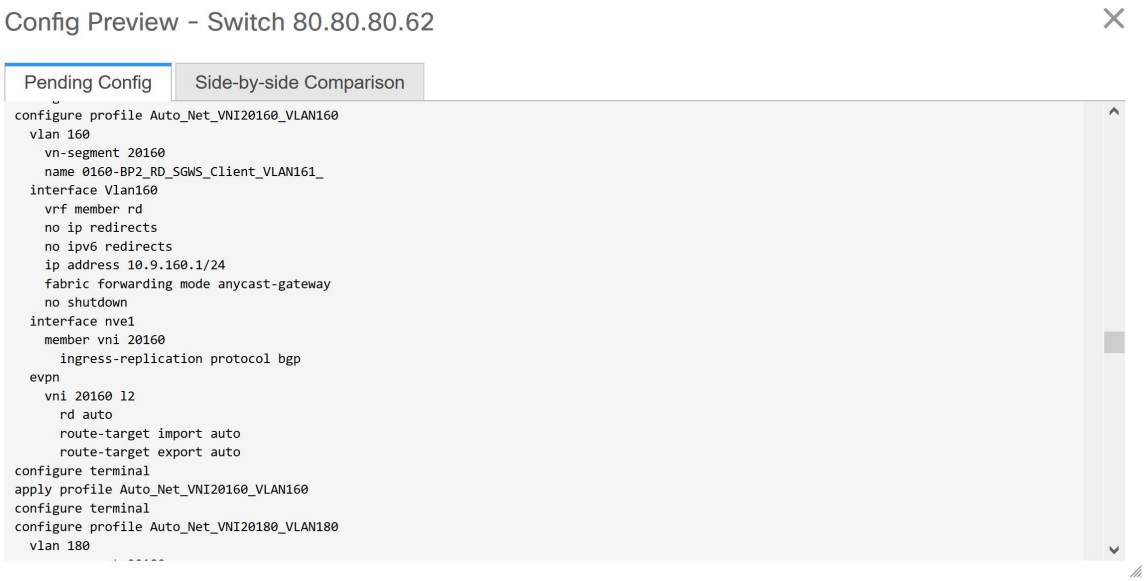
The Side-by-side Comparison tab displays the running configuration and expected configuration side-by-side.

The **Pending Config** tab displays the set of configurations that need to be deployed on a switch in order to go from the current running configuration to the current expected or intended configuration.

The **Pending Config** tab may show many configuration lines that will be deployed to the switches. Typically, on a successful brownfield import, these lines correspond to the configuration profiles pushed to the switches for a overlay network configuration. Note that the existing network and VRF-related overlay configurations are not removed from the switches.

Note

The configuration profiles are NDFC required constructs for managing the VXLAN configurations on the switches. During the Brownfield import process, they capture the same information as the original VXLAN configurations already present on the switches. In the following image, the configuration profile with **vlan 160** is applied.



As part of the import process, after the configuration profiles are applied, the original CLI based configuration references will be removed from the switches. These are the 'no' CLIs that will be seen towards the end of the diffs. The VXLAN configurations on the switches will be persisted in the configuration profiles. In the following image, you can see that the configurations will be removed, specifically, **no vlan 160**.

The removal of CLI based configuration is allowed if the **Overlay Mode** is set to **config-profile**, and not CLI.



The **Side-by-side Comparison** tab displays the Running Config and Expected Config on the switch.

Step 12 Close the **Config Preview Switch** window after reviewing the configurations.

Step 13 Click **Deploy Config** to deploy the pending configuration onto the switches.

If the **Status** column displays **FAILED**, investigate the reason for failure to address the issue.

The progress bar shows **100%** for each switch. After correct provisioning and successful configuration compliance, close the screen.

In the fabric topology screen that comes up, all imported switch instances are displayed in green color, indicating successful configuration. Also, the **Migration Mode** label is not displayed on any switch icon.

NDFC has successfully imported a VXLAN-EVPN fabric.

Post-transitioning of VXLAN fabric management to NDFC - This completes the transitioning process of VXLAN fabric management to NDFC. Now, you can add new switches and provision overlay networks for your fabric. For details, refer the respective section in the Fabrics topic in the configuration guide.

For more information, see Fabric Overview.

Configuration Profiles Support for Brownfield Migration

Cisco NDFC supports the Brownfield import of fabrics with VXLAN overlay provisioned with configuration profiles. This import process recreates the overlay configuration intent based on the configuration profiles. The underlay migration is performed with the usual Brownfield migration.

This feature can be used to recover your existing Easy fabric when a NDFC backup is not available to be restored. In this case, you must install the latest NDFC release, create a fabric, and then import the switches into the fabric.

Note that this feature is not recommended for the NDFC upgrade. For more information, see *Cisco NDFC Installation and Upgrade Guide*.

The following are the guidelines for the support of configuration profiles:

- The Brownfield migration of configuration profiles is supported for the **Data Center VXLAN EVPN** template.
- The configuration profiles on the switches must be a subset of the default overlay **Universal** profiles. If extra configuration lines are present that are not part of the **Universal** profiles, unwanted profile refreshes will be seen. In this case, after you recaluclate and deploy configuration, review the diffs using the **Side-by-side Comparison** feature and deploy the changes.
- Brownfield migration with switches having a combination of VXLAN overlay configuration profiles and regular CLIs is not supported. If this condition is detected, an error is generated, and migration is aborted. All the overlays must be with either configuration profiles or regular CLIs only.

Manually Adding PIM-BIDIR Configuration for Leaf or Spine Post Brownfield Migration

After brownfield migration, if you add new spine or leaf switches, you should manually configure the PIM-BIDIR feature.

The following procedure shows how to manually configure the PIM-BIDIR feature for a new Leaf or Spine:

Procedure

Step 1 Check the **base_pim_bidir_11_1** policies that are created for an RP added through the brownfield migration. Check the RP IP and Multicast Group used in each **ip pim rp-address** *RP_IP* **group-list** *MULTICAST_GROUP* **bidir** command.

Step 2 Add respective **base_pim_bidir_11_1** policies from the **View/Edit Policies** window for the new Leaf or Spine, push the config for each **base_pim_bidir_11_1** policy.

Migrating a VXLAN EVPN Multi-Site Fabric with Border Gateway Switches

When you migrate an existing VXLAN EVPN Multi-Site fabric with a border gateway switch into NDFC, make sure to note the following guidelines:

- Uncheck all **Auto** IFC creation related fabric settings. Review the settings and ensure they are unchecked as follows:
  - **Data Center VXLAN EVPN** fabric

    Uncheck **Auto Deploy Both** check box under **Resources** tab.
  - **VXLAN EVPN Multi-Site** fabric

    Uncheck **Multi-Site Underlay IFC Auto Deployment Flag** check box under **DCI** tab.
- Underlay Multisite peering: The eBGP peering and corresponding routed interfaces for underlay extensions between sites are captured in **switch_freeform** and **routed_inerfaces**, and optionally in the **interface_freeform** configs. This configuration includes all the global configs for multisite. Loopbacks for EVPN multisite are also captured via the appropriate interface templates.
- Overlay Multisite peering: The eBGP peering is captured as part of **switch_freeform** as the only relevant config is under **router bgp**.
- Overlays containing Networks or VRFs: The corresponding intent is captured with the profiles on the Border Gateways with **extension_type = MULTISITE**.

1. Create all the required fabrics including the Data Center VXLAN EVPN and Multi-Site Interconnect Network fabrics with the required fabric settings. Disable the Auto VRF-Lite options as mentioned above. For more information, refer to *Creating VXLAN EVPN Fabric* and *External Fabric* sections.
2. Import all the switches into all the required fabrics and set roles accordingly.
3. Click **Recalculate and Deploy** in each of the fabrics and ensure that the Brownfield Migration process reaches the 'Deployment' phase. Now, do not click **Deploy Configuration**.
4. Create the **VXLAN EVPN Multi-Site** fabric with the required fabric settings and disable the **Auto MultiSite IFC** options as shown in Guidelines. For more information, see *Creating a VXLAN EVPN Multi-Site Fabric*.

5. Move all the member fabrics into the VXLAN EVPN Multi-Site fabric. Do not proceed further till this step is completed successfully. For more information, see *Moving the Member1 Fabric Under VXLAN EVPN Multi-Site-Parent-Fabric*.



Note The Overlay Networks and VRFs definitions in each of the Easy Fabrics must be symmetric for them to get added successfully to the VXLAN EVPN Multi-Site fabric. Errors will be reported if any mismatches are found. These must be fixed by updating the overlay information in the fabric(s) and added to the VXLAN EVPN Multi-Site fabric.

6. Create all the Multisite Underlay IFCs such that they match the IP address and settings of the deployed configuration.



Note Additional interface configurations must be added to the Source/Destination interface freeform fields in the **Advanced** section as needed.

For more information, see *Configuring Multi-Site Overlay IFCs*.

7. Create all the Multisite Overlay IFCs such that they match the IP address and settings of the deployed configuration. You will need to add the IFC links. For more information, see *Configuring Multi-Site Overlay IFCs*.
8. If there are VRF-Lite IFCs also, create them as well.



Note If the Brownfield Migration is for the case where Configuration Profiles already exist on the switches, the VRF-Lite IFCs will be created automatically in Step #3.

9. If Tenant Routed Multicast (TRM) is enabled in the VXLAN EVPN Multi-Site fabric, edit all the TRM related VRFs and Network entries in VXLAN EVPN Multi-Site and enable the TRM parameters.

   This step needs to be performed if TRM is enabled in the fabric. If TRM is not enabled, you still need to edit each Network entry and save it.
10. Now click **Recalculate and Deploy** in the VXLAN EVPN Multi-Site fabric, but, do not click **Deploy Configuration**.
11. Navigate to each member fabric, click **Recalculate and Deploy**, and then click **Deploy Configuration**.

This completes the Brownfield Migration. You can now manage all the networks or VRFs for BGWs by using the regular NDFC Overlay workflows.

When you migrate an existing VXLAN EVPN Multi-Site fabric with border gateway switches (BGW) that has a Layer-3 port-channel for Underlay IFCs, make sure to do the following steps:



Note Ensure that the child fabrics are added into VXLAN EVPN Multi-Site before migrating the VXLAN EVPN Multi-Site fabric.

1. Click on appropriate VXLAN EVPN Multi-Site child fabric and navigate to **Fabrics** > **Interfaces** to view the BGW. Choose an appropriate Layer-3 port channel to use for underlay IFC.
2. On **Policy** column, choose **int_port_channel_trunk_host_11_1** from drop-down list. Enter the associated port-channel interface members and then click **Save**.
3. Navigate to the **Tabular view** of the VXLAN EVPN Multi-Site fabric. Edit layer-3 port link, choose the multisite underlay IFC link template, enter source and destination IP addresses. These IP addresses are the same as existing configuration values on the switches
4. Do the steps from step 7 to 11 from above procedure.