



Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - VRF Lite

External connectivity from data centers is a prime requirement. Virtual eXtensible Local Area Network (VXLAN) Border Gateway Protocol (BGP) Ethernet VPN (EVPN) based data center fabrics provide east-west connectivity by distributing IP-MAC reachability information among various devices within the fabric. While the EVPN Multi-Site feature provides inter site connectivity, the VRF Lite feature is used for connecting the fabric to an external Layer 3 domain. Tenants, typically represented by virtual routing and forwarding instances (VRFs) can procure external connectivity via special nodes called borders. In this way, tenant workloads in one data center fabric can have Layer 3 connectivity to hosts within the same VRF in other fabrics. This chapter describes LAN Fabric provisioning of the Nexus 9000-based border devices through the Cisco® Data Center Network Manager (DCNM) for the VRF Lite use case. This use case covers VRF extension from border devices connected to edge routers that in turn provide connectivity to the external fabric.

- [Prerequisites](#) , on page 1
- [Sample Scenario](#) , on page 2
- [VRF Lite Configuration](#) , on page 4
- [Deploying VRF Instances on Border Leafs](#) , on page 11
- [Undeploying VRF Instances on the Border Leafs](#) , on page 16
- [Additional References](#) , on page 21
- [Appendix](#) , on page 21

Prerequisites

- The VRF Lite feature requires Cisco Nexus 9000 Series NX-OS Release 7.0(3)I6(2) or later.
- Familiarity with VXLAN BGP EVPN data center fabric architecture and top-down based LAN fabric provisioning through the DCNM.
- Fully configured VXLAN BGP EVPN fabrics including underlay and overlay configurations on the various leaf and spine devices, external fabric configuration through DCNM, and relevant external fabric device configuration (edge routers, for example).
 - A VXLAN BGP EVPN fabric (and its connectivity to an external Layer 3 domain for north-south traffic flow) can be configured manually or using DCNM. This document explains the process to connect the fabric to an edge router (outside the fabric, towards the external fabric) through DCNM. So, you should know how to configure and deploy VXLAN BGP EVPN and external fabrics through

DCNM. For more details, see the **Control** chapter in the *Cisco DCNM LAN Fabric User Guide, Release 11.0(1)*.

- Ensure that the role of the designated border leaf switches is *Border*. To verify, right-click the switch and click **Set role**. You can see that (*current*) is added to the current role of the switch. If the current role is not *Border* or *Border Gateway*, you should remove the device from the fabric and discover it again through DCNM using the POAP bootstrap option and re-provision the configurations for the device.



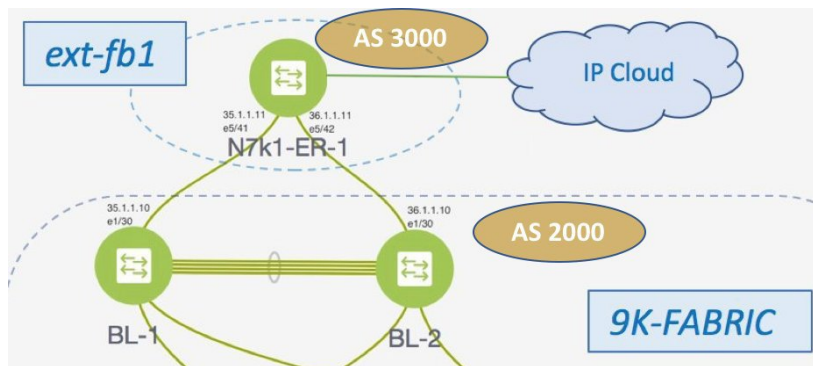
Note For an explanation on the VRF Lite feature, see the [Cisco Programmable Fabric with VXLAN BGP EVPN Configuration Guide](#) document.

Sample Scenario

The VRF Lite feature is explained through an example scenario. Consider a VXLAN BGP EVPN fabric, *9K-FABRIC*, whose border devices *BL-1* and *BL-2* are connected through an edge router in the fabric *ext-fb1*, to a shared IP core. This document will show you how to enable Layer 3 north-south traffic between the VXLAN fabric border devices and the edge router.



Note In this scenario, DCNM allows provisioning for fabric switches and the border devices. The edge router connected to the border devices needs to be manually configured (the edge router *N7k1-ER1* [or *ER1*] in the *ext-fb1* fabric is connected to *BL1* and *BL2* in *9K-FABRIC*).



Network configurations for the fabric is provisioned through DCNM. For external Layer 3 reachability from hosts connected to leaf switches within the fabric, border devices need to be provisioned with the appropriate VRF configuration. Multiple border devices in the fabric ensure redundancy in the case of failures as well as effective load distribution.

N7k1-ER1 (or *ER1*) is directly attached to the 2 border leaves. From the VXLAN fabric's point of view, the edge router belongs to an external fabric, *ext-fb1*, with a different AS number. For representation purposes, the *ext-fb1* fabric is created as an external fabric through DCNM, and *ER1* is associated with it in DCNM.



Note External fabric creation is a prerequisite for this use case. To create an external fabric *ext-fb1* in DCNM, follow these steps.

1. Click Control > Fabric Builder.

The Fabric Builder page comes up.

2. Click Create Fabric.

The Add Fabric screen comes up.

3. Enter the fabric name (*ext-fb1*) and select *External_Fabric* in the Fabric Template drop-down box.

4. Enter the BGP AS number and click Save.

5. The *ext-fb1* fabric is created as an external fabric.

VRF Lite—This requires setting up the border leaf configuration for enabling the VRF Lite feature by establishing eBGP peering from the border leaf to appropriate external devices like the edge routers. In this context, border leaves are special devices that allow clear control and data plane segregation from the fabric domain to the external Layer 3 domain (while allowing for policy enforcement points for any inter-fabric traffic).

The steps involved to enable VRF Lite feature are:

- 1. Connecting the VXLAN BGP EVPN fabric with the edge router**—Top-Down deployment for the VRF Lite feature configures route maps and an eBGP session in the default VRF through an interface (parent interface) connected to the edge router. This is a one-time setup for each edge router connected to a border leaf.
- 2. VRF extensions**—For each VRF that is to be extended, a unique sub interface towards the edge router and an eBGP session through this sub interface is configured on the border leaf. This is a per-VRF configuration. The corresponding configurations have to be manually enabled on the edge router too.

The end-to-end configurations can be split into these 2 steps:

1. VRF Lite configurations on the border leaves (*BL-1, BL-2*)

- 1. VRF Lite function on *BL-1* and *BL-2* in *9K-FABRIC* that are directly connected to *ER-1*.**
- 2. Configurations on edge router *ER-1* -** These configurations are not in the scope of DCNM provisioning and this document. It is mentioned here for completeness and sample configurations are provided in the *Appendix* section.

2. Deploying VRF instances on the border leaves (*BL-1, BL-2*)

For this example, multiple VRFs will be configured on the border leaves in *9K-FABRIC*.

After successful VRF Lite deployment at the border leaves and on the edge router, traffic will flow between them.



Note In the DCNM topology view, the lines connecting devices managed by DCNM (for example, *BL-1* to *N7k1-SPINE-1*) symbolize a physical cable connection. They do not indicate that the connection is functional and traffic flows between them.

To start off with, let us consider VRF Lite provisioning on border leafs *BL-1* and *BL-2* through DCNM Top-Down LAN Fabric Provisioning.

VRF Lite Configuration

VRF Lite Configuration (on *BL-1* towards *ER-1* in *9K-FABRIC*)

1. Click **Control** > **Network & VRFs**. The LAN Fabric Provisioning page appears.
2. Click **Continue**. The Select a Fabric page comes up.
3. Select *9K-FABRIC* from the drop-down box since you are configuring border leaf *BL-1* in the fabric *9K-FABRIC*.

In the same page, click **Fabric Extension Setup** since the purpose of this task is to allow *9K-FABRIC* to communicate to the edge router in the external fabric.

The Fabric Extension screen comes up.

The screenshot shows the 'Fabric Extension' configuration page. At the top, there is a title 'Fabric Extension' and a close button. Below it is the 'Inter-Fabric Connections' section, which includes a 'Selected 0 / Total 0' indicator and a 'Show Quick Filter' dropdown. The main area is a table with the following columns: Type, Source Fabric, Source Device, Source Interface, Destination Fabric, Destination Device, Destination Interface, Configuration, and Status. The table is currently empty, displaying 'No data available' at the bottom.

The **Inter-Fabric Connections** section lists previously created external connections from the border leafs in *9K-FABRIC*. This section is empty as this is the first time you are adding an external connection. Each row represents a physical or logical connection between a border leaf in *9K-FABRIC* and the edge router in the *ext-fb1* fabric. For each connection, the source fabric, source device, source interface, destination fabric, destination device, and destination interface are listed along with the type of external connectivity.

To extend the fabric through VRF-Lite, you should first create an extension.

Extension from *BL-1* to *ER-1*

1. Click on the + icon (at the top left part of the screen) to add a new external connection. The **Add Inter-Fabric Connection** screen appears.

Add Inter-Fabric Connections

1 Fabric Interconnect →
 2 Define Variables →
 3 Preview & Deploy

•
•
•

* Extension Type	<input type="text" value="VRF_LITE"/>	
* Base Template	<input type="text" value="ext_base_setup"/>	
* Extension Template	<input type="text" value="ext_fabric_setup"/>	
* Source Fabric	<input type="text" value="9K-FABRIC"/>	
* Destination Fabric	<input type="text"/>	
* Source Device	<input type="text"/>	<small>ⓘ VRF_LITE: Set switch role - Border; MULTISITE: Set switch role - "Border Gateway"</small>
* Source Interface	<input type="text"/>	
* Destination Device	<input type="text"/>	
* Destination Interface	<input type="text"/>	

Previous
Next
Save & Deploy
Cancel

By default, *VRF_LITE* is populated in the **Extension Type** field.

Base Template—By default, the *ext_base_setup* base template is populated. This template represents a one-time configuration pushed to the border leaf *BL-1*.

Extension Template—*ext_fabric_setup*, as the name indicates, represents the template that outputs the configuration required to setup the connection between the border leaf and the edge router. As opposed to the configuration represented by the Base Template that is applied only once per border leaf, the Extension Template generated configuration is executed once for every connection between a border leaf and the edge router.



Note These templates are auto-populated with corresponding pre-packaged default templates based on your selection. You can add, edit or delete user-defined templates. For more details, see the *Template Library* section in the *Control* chapter .

Source Fabric—This field is pre-populated with *9K-FABRIC* since the VRF Lite connection is between *BL-1* in *9K-FABRIC* and *ER-1* in the *ext-fb1* fabric.

Destination Fabric—Choose *ext-fb1*.

Source Device and **Source Interface**—Choose *BL-1* as the source device and an Ethernet interface that needs to be connected to *ER-1*.

Destination Device and **Destination Interface**—Choose *ER-1* as the destination device and the Ethernet interface that connects to the border leaf *BL-1*.

Note that based on the selection of the source device and source interface, the destination information will be auto-populated based on CDP information if available. There is extra validation performed to ensure that the destination external device is indeed part of the destination fabric.



Note If the destination device is a non Cisco Nexus device (such as a Cisco ASR or Catalyst device), then you can manually type in the switch name and the interface name (for example, *Catalyst9400* and *Ethernet1/1*) in these fields.

After filling up the Fabric Interconnect section, the screen looks like this.

Add Inter-Fabric Connections

1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

- * Extension Type: VRF_LITE
- * Base Template: ext_base_setup
- * Extension Template: ext_fabric_setup
- * Source Fabric: 9K-FABRIC
- * Destination Fabric: ext-fb1
- * Source Device: BL-1
- * Source Interface: Ethernet1/30
- * Destination Device: N7k1-ER-1
- * Destination Interface: Ethernet5/41

Buttons: Previous, Next, Save & Deploy, Cancel

2. Click **Next** to go to the **Define Variables** section. The fields are:

IF_NAME—In this field, the interface name is auto-populated from the previous step.

Interface IP_MASK—Fill up this field with the IP address and mask of the *BL-1* interface that connects to *ER-1*.


NEIGHBOR_IP—Fill up this field with the IP address of the *ER-1* interface that connects to *BL-1*.

NEIGHBOR_ASN—In this field, the AS number of *ER-1*'s fabric (*ext-fb1*) will be auto-populated.

Extension Type—In this field, **VRF_Lite** will be auto-populated.

A sample screenshot of the fully filled up screen:

Add Inter-Fabric Connections



Network Profile

General	
* IF_NAME	Ethernet1/30 ?
* IP_MASK	35.1.1.10/24 ?
* NEIGHBOR_IP	35.1.1.11 ?
* NEIGHBOR_ASN	3000 ?
* Extension Type	VRF_LITE ?

Previous Next Save & Deploy Cancel

- Click **Next** to go to the **Preview and Deploy** section. The two sections of the screen are shown in the 2 images:

Switch:

Generated Configuration:

```
ip prefix-list default-route seq 5 permit 0.0.0.0/0 le 1
ip prefix-list host-route seq 5 permit 0.0.0.0/0 eq 32
route-map EXTCON-RMAP-FILTER deny 10
  match ip address prefix-list default-route
route-map EXTCON-RMAP-FILTER deny 20
  match ip address prefix-list host-route
route-map EXTCON-RMAP-FILTER permit 1000

ipv6 prefix-list default-route-v6 seq 5 permit 0::/0
ipv6 prefix-list host-route-v6 seq 5 permit 0::/0 eq 128
route-map EXTCON-RMAP-FILTER-V6 deny 10
  match ipv6 address prefix-list default-route-v6
route-map EXTCON-RMAP-FILTER-V6 deny 20
  match ip address prefix-list host-route-v6
route-map EXTCON-RMAP-FILTER-V6 permit 1000

interface Ethernet1/30
```

Generated Config

Previous Next Save & Deploy Cancel

Add Inter-Fabric Connections



1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

Switch:

Generated Configuration:

```
interface Ethernet1/30
  no switchport
  ip address 35.1.1.10/24
  no shutdown

router bgp 2000
  address-family ipv4 unicast
  redistribute direct route-map RMAP-REDIST-DIRECT
  neighbor 35.1.1.11 remote-as 3000
  update-source Ethernet1/30
  address-family ipv4 unicast
  next-hop-self
```

Previous Next Save & Deploy Cancel

**Note**


If a VXLAN BGP EVPN fabric border leaf is connected to more than one edge router, the prefix-list and route map configurations are pushed only for the first fabric extension instance. Similarly when deleting fabric extension instances on a border leaf, the global configurations (prefix-list and route-maps) are removed from the border leaf only after the last fabric extension instance is deleted.

In this screen, you can preview the configurations that will be deployed to *BL-1*. Note that no configuration will be pushed to the external device (edge router) itself.

A one-time configuration of route maps along with the parent interface connection is displayed. Also, you can see that BGP peering information in the default routing table is configured for *BL-1*. The corresponding BGP configurations should be manually enabled on *ER-1*.

4. Click **Save and Deploy** to complete the task. This results in the configuration getting pushed to *BL-1*. The external connection will appear in the Fabric Extension screen, under **Inter-Fabric Connections**.


At this stage, an extension is enabled from *BL-1* to *ER-1*, as indicated by the arrow in the image.



Next, you need to enable an extension from *BL-2* to *ER-1* too.

VRF Lite Configuration (on *BL-2* towards *ER-1* in *9K-FABRIC*)

As described in the previous section, enable an extension from *BL-2* to *ER-1*. After configurations are pushed to *BL-2*, an extension will be enabled from *BL-2* to *ER-1*, as shown in the screen shot.



A preview of the configurations on *BL-2* is given in these 2 screen shots.

Add Inter-Fabric Connections

Switch:

Generated Configuration:

```

ip prefix-list default-route seq 5 permit 0.0.0.0/0 le 1
ip prefix-list host-route seq 5 permit 0.0.0.0/0 eq 32
route-map EXTCON-RMAP-FILTER deny 10
  match ip address prefix-list default-route
route-map EXTCON-RMAP-FILTER deny 20
  match ip address prefix-list host-route
route-map EXTCON-RMAP-FILTER permit 1000

ipv6 prefix-list default-route-v6 seq 5 permit 0::/0
ipv6 prefix-list host-route-v6 seq 5 permit 0::/0 eq 128
route-map EXTCON-RMAP-FILTER-V6 deny 10
  match ipv6 address prefix-list default-route-v6
route-map EXTCON-RMAP-FILTER-V6 deny 20
  match ip address prefix-list host-route-v6
route-map EXTCON-RMAP-FILTER-V6 permit 1000

```

Generated Config

4

Add Inter-Fabric Connections ✕

1 Fabric Interconnect → 2 Define Variables → 3 Preview & Deploy

Switch:

Generated Configuration:

```
interface Ethernet1/30
  no switchport
  ip address 36.1.1.10/24
  no shutdown

router bgp 2000
  address-family ipv4 unicast
  redistribute direct route-map RMAP-REDIST-DIRECT
  neighbor 36.1.1.11 remote-as 3000
  update-source Ethernet1/30
  address-family ipv4 unicast
  next-hop-self
```

Previous Next Save & Deploy Cancel

Edge Router Configurations

Apart from the DCNM provisioning on the border leaves in the two fabrics, you should also enable appropriate configurations on *ER-1* for connectivity between the edge router and the border leaves. Sample *ER-1* configuration is provided in the *Appendix* section for your reference.

What to do next—As noted earlier, the end-to-end VRF-Lite configurations through DCNM Top-Down provisioning includes these 2 steps:

1. VRF Lite configurations on the border leaves (*BL-1*, *BL-2*)
2. Deploying VRF Instances on the border leaves (*BL-1*, *BL-2*)

At this stage, the first step explanation is complete. The next section explains how VRF extension configuration is pushed to the border leaves.

Deploying VRF Instances on Border Leaves

Before you begin—In this scenario, we will deploy three VRF instances, *MyVRF-50016*, *MyVRF-50018*, and *MyVRF-50019* on the border leaves *BL-1* and *BL-2* in *9K-FABRIC*. You should ensure that you have already deployed the corresponding network(s) on the fabric's leaf switches.

After deploying one network on the leaf switches, you will have to deploy the associated VRF on the border leafs so that the network(s) can be extended from/to the *9K-FABRIC*. To know how to create a fabric, and networks and VRFs, see the *Control* chapter in the *Cisco DCNM LAN Fabric User Guide, Release 11.0(1)*.

In the Select a Fabric page, ensure that you select *9K-FABRIC* in the drop-down box and click **Continue** (at the top right part of the screen). After clicking **Continue**, the **Networks** page comes up.

Click on **VRF View**. The **VRFs** page comes up.

We will deploy 3 new VRF instances *MyVRF-50016*, *MyVRF-50018*, and *MyVRF-50019* on the border leafs. To do that, select the checkboxes (in the extreme left column).

Fabric Selected: 9K-FABRIC

Selected 3 / Total 138

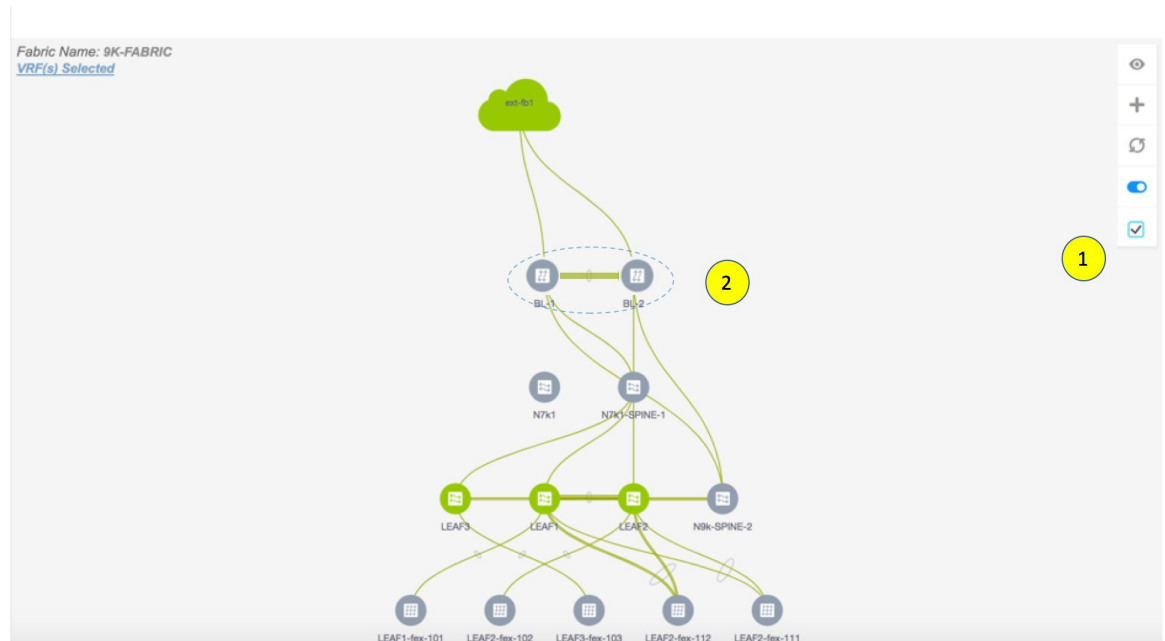
<input type="checkbox"/>	VRF Name	VRF ID	Status
<input type="checkbox"/>	MyVRF_50000	50000	DEPLOYED
<input checked="" type="checkbox"/>	MyVRF_50016	50016	DEPLOYED
<input checked="" type="checkbox"/>	MyVRF_50018	50018	NA
<input checked="" type="checkbox"/>	MyVRF_50019	50019	NA
<input type="checkbox"/>	MyVRF_50500	50500	DEPLOYED

Click the **Continue** button at the top right part of the screen. The VRF Deployment page (Topology View) comes up. You can deploy VRFs on multiple switches simultaneously, but with the same role. So, deploy the selected VRFs on the border leafs.



Note

In the image, you can see that the VRF instances are deployed on the leaf switches (green color indicates deployed status). Note that the color code, and hence the deployment state on switches is contextual and specific to the selection. In this scenario, the deployed state only depicts that the 3 selected VRFs are deployed on leaf switches LEAF3, LEAF1 and LEAF2. It does not display information about other VRF deployment instances, if any.



Select the multi-select check box from the panel of options available (*Step 1* in the image).

Then, click your mouse (or track pad) and drag the cursor across *BL-1* and *BL-2* (*Step 2* in the image).

Immediately, the **Switches Deploy** screen (for VRFs) appears. A tab is displayed for each VRF.

Click the checkbox next to the **Switch** column. Both the border leaf check boxes will be selected automatically. Alternatively, you can select check boxes next to the switches.

Switches Deploy

Fabric Name: 9K-FABRIC

MyVRF_50016 MyVRF_50018 MyVRF_50019

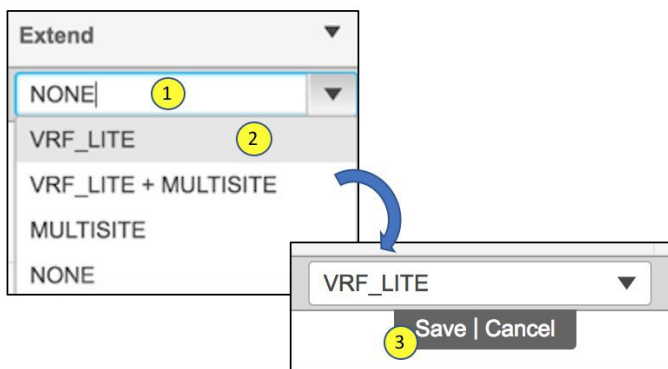
Deploy Options:

Select the row and click on the cell to edit and save changes

<input type="checkbox"/>	Switch	VLAN	Extend	Status
<input checked="" type="checkbox"/>	BL-1	2001	NONE	NA
<input checked="" type="checkbox"/>	BL-2	2001	NONE	NA

Click on **NONE** in the **Extend** column, select *VRF_LITE* and click on the **Save** button below it.

Repeat this action for the second row too. A sample screenshot:



This creates a VRF Lite extension for this VRF, as seen in the **Extension Details** section that appears at the bottom part of the screen.

In the **Extension Details** section, select the **Source Switch** checkbox (or ensure that you select the check box in each row). This is how the screen looks when you select both the switches in the **Extension Details** section.

The corresponding dot1Q tag for the VRF is auto-populated in the **DOT1Q_ID** field.

Switches Deploy ✕

Fabric Name: 9K-FABRIC

MyVRF_50016 | MyVRF_50018 | MyVRF_50019

Deploy Options:

① Select the row and click on the cell to edit and save changes

<input type="checkbox"/>	Switch	VLAN	Extend	Status
<input checked="" type="checkbox"/>	BL-1	2001	VRF_LITE	PENDING
<input checked="" type="checkbox"/>	BL-2	2001	VRF_LITE	PENDING

Extension Details

<input type="checkbox"/>	Source Switch	Type	IF_NAME	DOT1Q_ID	IP_MASK
<input checked="" type="checkbox"/>	BL-1	VRF_LITE	Ethernet1/30	3	35.1.1.10/24
<input checked="" type="checkbox"/>	BL-2	VRF_LITE	Ethernet1/30	3	36.1.1.10/24

Now, select the *MyVRF_50018* and *MyVRF_50019* and similarly update relevant parameters.

Click the **Save** button at the bottom right part of the Switches Deploy screen to save all VRFs' configurations on the selected switches. The VRF Deployment screen (Topology view) appears.

BL-1 and *BL-2* icons will be displayed in blue color, indicating that a deployment is pending. If you want to check your configurations, click on the Preview (eye) icon.

Dashboard Fabric Selection Network Selection Network Deployment Deploy Detailed View

Preview Configuration

Select a Switch: BL-1 Select a VRF: MyVRF_50016

Generated Configuration:

```

maximum-paths ibgp 2

network 0::0

neighbor 35.1.1.11 remote-as 3000
address-family ipv4 unicast
send-community both
route-map EXTCON-RMAP-FILTER out

neighbor 35:1:1:1:2 remote-as 3000
address-family ipv6 unicast
send-community both
route-map EXTCON-RMAP-FILTER-V6 out

interface Ethernet1/30.3
encapsulation dot1q 3
vrf member MyVRF_50016
ip address 35.1.1.10/24

ipv6 address 35:1:1:1:1/64
no shutdown

configure terminal

```

Generated Config

Topology View: A network diagram showing a cloud connected to BL-1 and BL-2. Below them are N7K1 and N7K-SPINE-1, and N7K-SPINE-2. At the bottom are several LEAF switches (LEAF3, LEAF2, LEAF1) and N7K-SPINE-2. A 'Generated Config' label points to the configuration window.

Preview Configuration

Select a Switch: BL-2 Select a VRF: MyVRF_50019

Generated Configuration:

```

configure profile 9K-FABRIC-Default_VRF_Extension-50019
vlan 2003
  vn-segment 50019
  interface vlan 2003
    vrf member MyVRF_50019
    ip forward
    ipv6 forward
    no ip redirects
    no ipv6 redirects
    mtu 9216
    no shut

interface nve 1
  member vni 50019 associate-vrf

vrf context MyVRF_50019
vni 50019
rd auto
address-family ipv4 unicast
route-target both auto
route-target both auto evpn

ip route 0/0 36.1.1.11
address-family ipv6 unicast
route-target both auto
route-target both auto evpn

ipv6 route 0::/0 36:1:1:1:2

```

Generated Co

You can select a switch and a VRF to view corresponding configurations. Configuration details of *MyVRF_5016* that is pushed to *BL-1* are included in the Appendix section.

After you verify that the configurations that are generated from the profiles are correct for the selected switches, click the **Deploy** button (on the top right part of the Topology View screen) to deploy the *MyVRF_50016*, *MyVRF_50018*, and *MyVRF_50019* VRF configurations on *BL-1* and *BL-2*.

DCNM shows the deployment status in the topology by highlighting the switch icons with different colors, yellow for In Progress, green for Deployed, and red for Out of sync status.

When the switch icons turn green, it indicates that the *MyVRF_50016*, *MyVRF_50018*, and *MyVRF_50019* VRF configurations have been deployed on the border leafs of the *9K-FABRIC*. You can also click the **Detailed View** option to see the status.

After configurations in *9K-FABRIC* are complete, you should enable configurations in *Fabric2* too.

Resources

The Resources page (Resource Allocation section) gives information of all the resources allocated or deployed on each device per fabric. This includes the network VLANs, VRF VLANs, and the sub interface dot1q identifiers employed for the VRF Lite extension. Once a VRF is undeployed, the associated resources in the Resource Allocation section will be unallocated and updated immediately.

To access the Resource Allocation page, click **Control > Management > Resources**.

As we can see in the screenshot below, after deploying VRF instances *MyVRF_50016*, *MyVRF_50018* and *MyVRF_50019* on the border leafs, the associated VLAN-VRF mapping is displayed in the Resource Allocation screen.

Scope Type	Scope	Allocated Resource	Allocated To	Resource Type	Is Allocated?	Allocated On
Device	FDO20401LB4	2004	MyVRF_50000	TOP_DOWN_VRF_VLAN	Yes	7/27/2018, 2:07...
Device	FDO20401LB4	2005	MyVRF_500017	TOP_DOWN_VRF_VLAN	Yes	7/27/2018, 2:37...
Device	FDO20401LB4	2000	MyVRF_50003	TOP_DOWN_VRF_VLAN	Yes	7/27/2018, 7:07...
Device	FDO20401LB4	2001	MyVRF_50016	TOP_DOWN_VRF_VLAN	Yes	7/30/2018, 3:49...
Device	FDO20401LB4	2002	MyVRF_50018	TOP_DOWN_VRF_VLAN	Yes	7/30/2018, 3:49...
Device	FDO20401LB4	2003	MyVRF_50019	TOP_DOWN_VRF_VLAN	Yes	7/30/2018, 3:50...
Device	SAL18432P6G	2	Loopback2	LOOPBACK_ID	Yes	7/23/2018, 11:2...
DeviceInterface	FDO20401LB4	2	MyVRF_50003	TOP_DOWN_L3_DOT1Q	Yes	7/30/2018, 3:22...
DeviceInterface	FDO20401LB4	3	MyVRF_50016	TOP_DOWN_L3_DOT1Q	Yes	7/30/2018, 3:54...
DeviceInterface	FDO20401LB4	4	MyVRF_50018	TOP_DOWN_L3_DOT1Q	Yes	7/30/2018, 3:53...
DeviceInterface	FDO20401LB4	5	MyVRF_50019	TOP_DOWN_L3_DOT1Q	Yes	7/30/2018, 3:53...

The VRF instances *MyVRF_50016*, *MyVRF_50018*, and *MyVRF_50019* are deployed on *BL-1*, with their corresponding VLANs *2001*, *2002*, and *2003*.

Also, the corresponding dot1Q IDs *3*, *4*, and *5* are displayed

Undeploying VRF Instances on the Border Leafs

VRFs can be deployed/undeployed on the border leafs. The following steps will demonstrate undeployment of VRFs on the border leafs.

For *9K-FABRIC*, navigate to the **Networks** page and click **VRF View**. The VRFs page will be displayed.

Select *MyVRF-50018* and *MyVRF-50019* and click **Continue**.

Fabric Selected: 9K-FABRIC

VRFs Selected 2 / Total 138

Show All

<input type="checkbox"/>	VRF Name	VRF ID	Status
<input type="checkbox"/>	MyVRF_50000	50000	DEPLOYED
<input type="checkbox"/>	MyVRF_50016	50016	DEPLOYED
<input checked="" type="checkbox"/>	MyVRF_50018	50018	DEPLOYED
<input checked="" type="checkbox"/>	MyVRF_50019	50019	DEPLOYED
<input type="checkbox"/>	MyVRF_50500	50500	DEPLOYED
<input type="checkbox"/>	VRF 50011	50011	UNDEPLOYED

The Topology View page is displayed. Follow similar steps as described in the Deploying VRFs section on the border leafs.

Select *BL-1* and *BL-2* switches in the topology page. The **Switches Deploy** screen will be displayed.

A tab is displayed for each VRF. *MyVRF_50018* is currently selected in the below screenshot.

Switches Deploy

Fabric Name: 9K-FABRIC

MyVRF_50018 MyVRF_50019

Deploy Options:

① Select the row and click on the cell to edit and save changes

<input type="checkbox"/>	Switch	VLAN	Extend	Status
<input checked="" type="checkbox"/>	BL-1	2002	VRF_LITE	DEPLOYED
<input checked="" type="checkbox"/>	BL-2	2002	VRF_LITE	DEPLOYED

Extension Details

<input type="checkbox"/>	Source Switch	Type	IF_NAME	DOT1Q_ID	IP_MASK
<input checked="" type="checkbox"/>	BL-1	VRF_LITE	Ethernet1/30	4	35.1.1.10/24
<input checked="" type="checkbox"/>	BL-2	VRF_LITE	Ethernet1/30	4	36.1.1.10/24

Double click the checkbox next to the **Switch** column or uncheck the check box next to *BL-1* and *BL-2*. Both of the check boxes will be de-selected and the **Extension Details** section will disappear at the bottom part of the screen.

Undeploying VRF Instances on the Border Leafs

Switches Deploy

Fabric Name: 9K-FABRIC

MyVRF_50018 MyVRF_50019

Deploy Options:

ⓘ Select the row and click on the cell to edit and save changes

<input type="checkbox"/>	Switch	VLAN	Extend	Status
<input type="checkbox"/>	BL-1	2002	VRF_LITE	DEPLOYED
<input type="checkbox"/>	BL-2	2002	VRF_LITE	DEPLOYED

Save

Now, select *MyVRF_50019* and update similarly.

Switches Deploy

Fabric Name: 9K-FABRIC

MyVRF_50018 MyVRF_50019

Deploy Options:

ⓘ Select the row and click on the cell to edit and save changes

<input type="checkbox"/>	Switch	VLAN	Extend	Status
<input type="checkbox"/>	BL-1	2003	VRF_LITE	DEPLOYED
<input type="checkbox"/>	BL-2	2003	VRF_LITE	DEPLOYED

Save

Click on the **Save** button at the bottom right part of the Switches Deploy screen to undeploy all VRF configurations on the selected switches. The VRF Deployment screen (Topology view) appears.

Similar to the deployment process, the *BL-1* and *BL-2* switch icons will be displayed in blue color, indicating pending undeployment. You can preview the information by clicking the Preview (eye) icon.

The configurations for *MyVRF_50018* on *BL-1* switch will be removed as displayed in the following screen. You can select a switch and VRF to view corresponding configurations.

Preview Configuration

Select a Switch: ▼

Select a VRF: ▼

Generated Configuration:

```
configure terminal
no apply profile 9K-FABRIC-Default_VRF_Extension-50018
no configure profile 9K-FABRIC-Default_VRF_Extension-50018
```

After you verify that the configuration profiles that will be removed are correct for the selected switches, click the **Deploy** button (on the top right part of the screen) to undeploy the *MyVRF_50018* and *MyVRF_50019* configurations on *BL-1* and *BL-2*.

Resources Update

To access the Resource Allocation page, click **Control > Management > Resources**. After undeploying the VRFs *MyVRF_50018* and *MyVRF_50019* on the border leaves, the Resource Allocation page has the associated VLAN-VRF mapping removed.

SCOPE: Easy60000 admin

Control / Management / Resources

Resource Allocation Selected 0 / Total 47

Show All

<input type="checkbox"/>	Scope Type ▲	Scope	Allocated Resource	Allocated To	Resource Type	Is Allocated?	Allocated On
<input type="checkbox"/>	Device	FDO20401LB4	2004	MyVRF_50000	TOP_DOWN_VRF_VLAN	Yes	7/27/2018, 2:07...
<input type="checkbox"/>	Device	FDO20401LB4	2005	MyVRF_500017	TOP_DOWN_VRF_VLAN	Yes	7/27/2018, 2:37...
<input type="checkbox"/>	Device	FDO20401LB4	2000	MyVRF_50003	TOP_DOWN_VRF_VLAN	Yes	7/27/2018, 7:07...
<input type="checkbox"/>	Device	FDO20401LB4	2001	MyVRF_50016	TOP_DOWN_VRF_VLAN	Yes	7/30/2018, 3:49...
<input type="checkbox"/>	Device	SAL18432P6G	2	Loopback2	LOOPBACK_ID	Yes	7/23/2018, 11:2...
<input type="checkbox"/>	DeviceInterface	FDO20401LB4	2	MyVRF_50003	TOP_DOWN_L3_DOT1Q	Yes	7/30/2018, 3:22...
<input type="checkbox"/>	DeviceInterface	FDO20401LB4	3	MyVRF_50016	TOP_DOWN_L3_DOT1Q	Yes	7/30/2018, 3:54...

In the screenshot, it shows that *MyVRF_50018* and *MyVRF_50019* that was deployed on *BL-1* with VLAN *2002* and *2003* are now removed/unallocated.

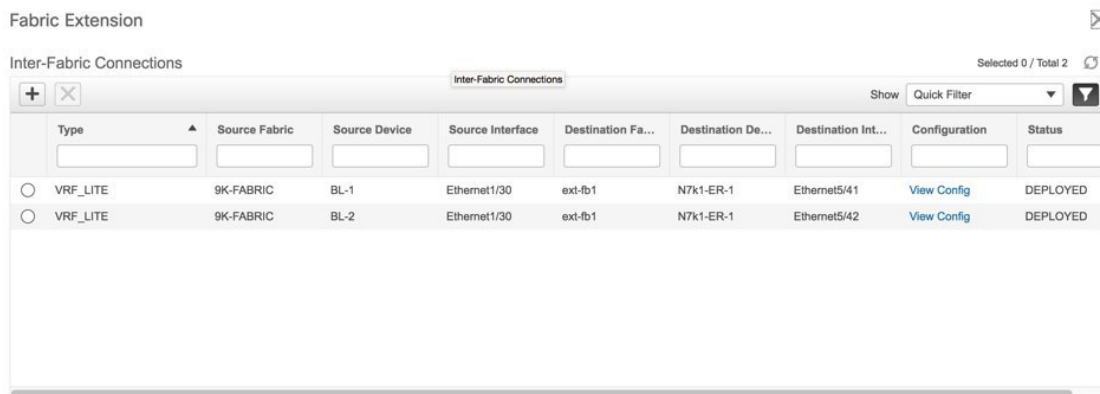
Remove VRF Lite Inter-fabric configuration on vPC border leaves

VRF Lite configuration can also be removed in a similar manner as long as there are no VRF extensions enabled over that connection. The following steps will demonstrate removal of *BL-1* and *BL-2* VRF Lite connections.

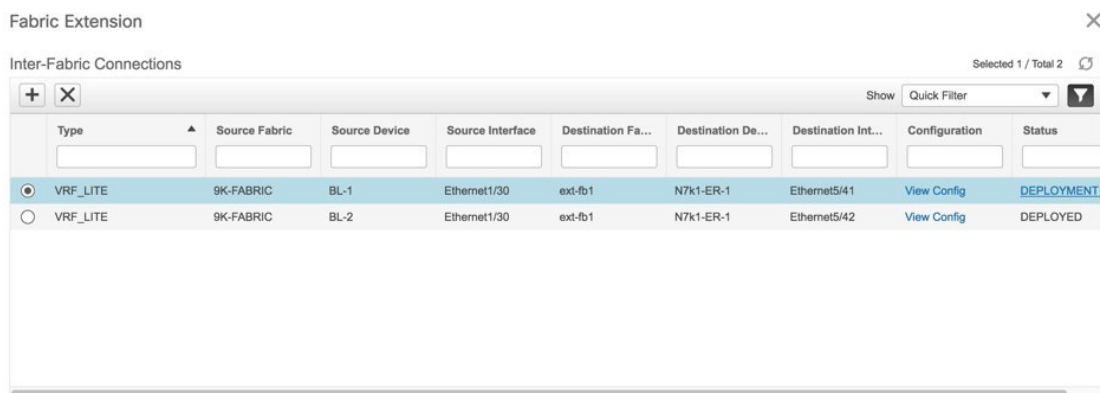
Follow similar steps as described in the VRF Lite configuration for *BL-1* in *9K-FABRIC*.

1. Click **Control > Networks & VRFs**.
2. Select *9K-FABRIC* from the drop-down box and click **Fabric Extension Setup**. The **Fabric Extension** screen comes up

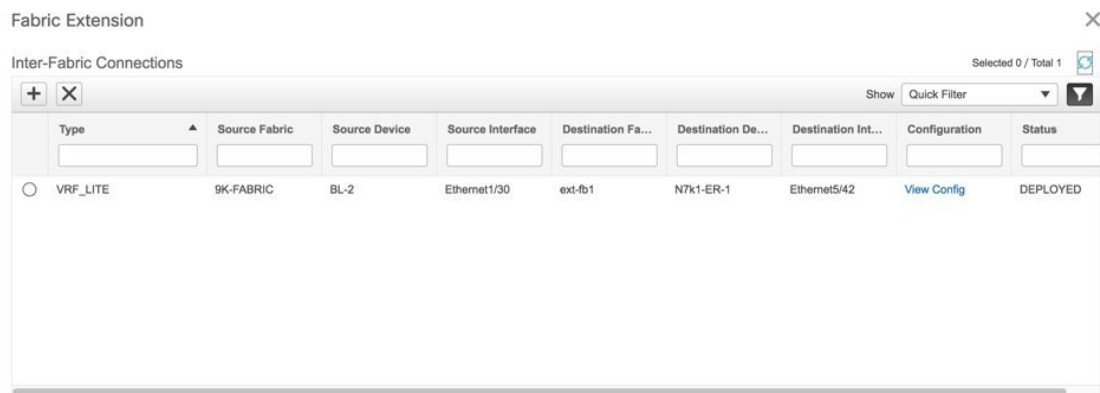
Remove VRF Lite Inter-fabric configuration on vPC border leafs



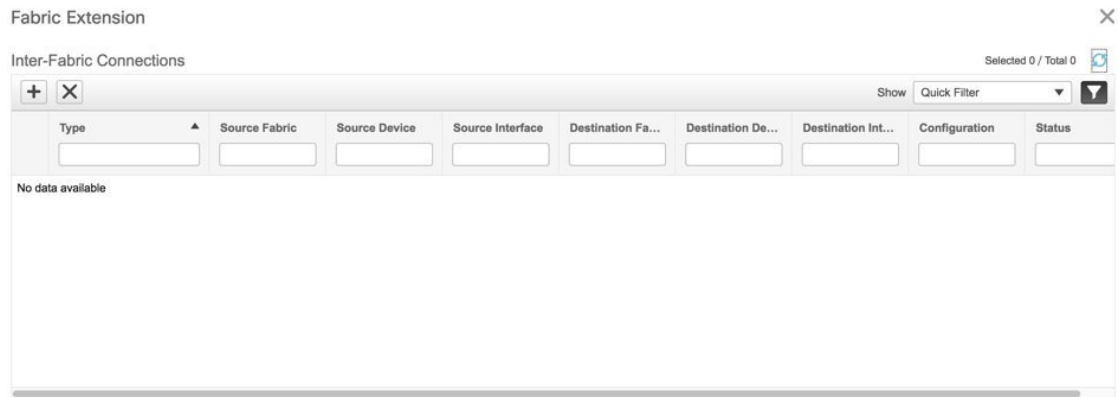
3. Click on the radio button next to **VRF_LITE** in the first row with Source Device *BL-1*.
4. Click the **X** button to delete this entry.



The next screen shows that the *BL-1* connection to *ER-1* is removed from the fabric extension list.



5. Similarly, select *BL-2* and click **X** to remove the *BL-2* connection to *ER-1*. After both *BL-1* and *BL-2* VRF Lite connections are removed, the Fabric Extension screen will have no entries.



Additional References

Document Title and Link	Document Description
Cisco Programmable Fabric with VXLAN BGP EVPN Configuration Guide	This document explains external connectivity using VRF Lite.

Appendix

Edge Router Configurations

ER-1 Configuration Example—The following configurations are enabled on *ER-1* to connect to *BL-1* and *BL-2* (border leaves), and reproduced here for reference.



Note *switch(config)#* refers to the global configuration mode. To access this mode, type the following on your switch: *switch# configure terminal*.

```
switch(config)#
interface Ethernet5/41  ## ER-1 interface to BL-1
 ip address 35.1.1.11/24
 no shutdown

interface Ethernet5/42  ## ER-1 interface to BL-2
 ip address 36.1.1.11/24
 no shutdown

router bgp 3000          ## eBGP sessions
 neighbor 35.1.1.10 remote-as 2000  ###Peering to BL-1 (eBGP)
  update-source Ethernet5/41
  address-family ipv4 unicast
  next-hop-self
 neighbor 36.1.1.10 remote-as 2000  ###Peering to BL-2 (eBGP)
  update-source Ethernet5/42
```

```
address-family ipv4 unicast
  next-hop-self
```

The following configurations are manually enabled on *ER-1* for VRF extension to the border leafs:

```
configure profile 9K-FABRIC-Default_VRF_Extension-50016
vrf context MyVRF_50016
  address-family ipv4 unicast
    route-target import 3000:3
    route-target export 3000:3
  rd 3000:3
interface Ethernet5/41.3
  encapsulation dot1Q 3
  vrf member MyVRF_50016
  ip address 35.1.1.11/24
  ipv6 address 35:1:1:1::2/64
  no shutdown
interface Ethernet5/42.3
  encapsulation dot1Q 3
  vrf member MyVRF_50016
  ip address 36.1.1.11/24
  ipv6 address 36:1:1:1::2/64
  no shutdown
router bgp 3000
  vrf MyVRF_50016
    address-family ipv4 unicast
      maximum-paths ibgp 2
      neighbor 35.1.1.10 remote-as 2000
      address-family ipv4 unicast
        send-community both
      neighbor 36.1.1.10 remote-as 2000
      address-family ipv4 unicast
        send-community both
```

Configurations Pushed to *BL-1* Through DCNM:

VRF extension pushed to *BL-1* through DCNM

```
### Route map
ip prefix-list default-route seq 5 permit 0.0.0.0/0 le 1
ip prefix-list host-route seq 5 permit 0.0.0.0/0 eq 32
route-map EXTCON-RMAP-FILTER deny 10
  match ip address prefix-list default-route
route-map EXTCON-RMAP-FILTER deny 20
  match ip address prefix-list host-route
route-map EXTCON-RMAP-FILTER permit 1000

ipv6 prefix-list default-route-v6 seq 5 permit 0::/0
ipv6 prefix-list host-route-v6 seq 5 permit 0::/0 eq 128
route-map EXTCON-RMAP-FILTER-V6 deny 10
  match ipv6 address prefix-list default-route-v6
route-map EXTCON-RMAP-FILTER-V6 deny 20
  match ip address prefix-list host-route-v6
route-map EXTCON-RMAP-FILTER-V6 permit 1000

### VRF-Lite interface of BL-1
interface Ethernet1/30
  no switchport
  ip address 35.1.1.10/24
  no shutdown
```

```

### External BGP (eBGP) session of BL-1
router bgp 2000
  address-family ipv4 unicast
    redistribute direct route-map RMAP-REDIST-DIRECT
  neighbor 35.1.1.11 remote-as 3000
  update-source Ethernet1/30
  address-family ipv4 unicast
    next-hop-self

```

The following configuration profile is pushed through DCNM when *MyVRF_50016* is deployed on *BL-1*:

```

configure profile 9K-FABRIC-Default_VRF_Extension-50016
vlan 2001
  vn-segment 50016
  interface vlan 2001
    vrf member MyVRF_50016
    ip forward
    ipv6 forward
    no ip redirects
    no ipv6 redirects
    mtu 9216
    no shutdown

interface nve 1
  member vni 50016 associate-vrf

vrf context MyVRF_50016
  vni 50016
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
    ip route 0/0 35.1.1.11
  address-family ipv6 unicast
    route-target both auto
    route-target both auto evpn
    ipv6 route 0::/0 35.1.1.1.2

router bgp 2000
  vrf MyVRF_50016 ## bgp VRF configured
  address-family ipv4 unicast
    advertise l2vpn evpn
    redistribute direct route-map FABRIC-RMAP-REDIST-SUBNET
    maximum-paths ibgp 2
    network 0/0
  address-family ipv6 unicast
    advertise l2vpn evpn
    redistribute direct route-map FABRIC-RMAP-REDIST-SUBNET
    maximum-paths ibgp 2
    network 0::/0
  neighbor 35.1.1.11 remote-as 3000
  address-family ipv4 unicast
    send-community both
    route-map EXTCON-RMAP-FILTER out
  neighbor 35.1.1.1.2 remote-as 3000
  address-family ipv6 unicast
    send-community both
    route-map EXTCON-RMAP-FILTER-V6 out

interface Ethernet1/30.3 #sub interface member of VRF deployed
  encapsulation dot1q 3
  vrf member MyVRF_50016
  ip address 35.1.1.10/24

```

```
    ipv6 address 35:1:1:1::1/64
    no shutdown

configure terminal
    apply profile 9K-FABRIC-Default_VRF_Extension-50016
```