



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 shows you how to extend a VRF to an external fabric. In DNCM, configuration parameters are enhanced as follows:

Configuration methods - You can configure VRF Lite through automatic configuration and through the DCNM GUI.

Supported destination devices - You can extend VRFs from a VXLAN fabric to Cisco Nexus and non-Nexus devices. A connected non-Cisco device can also be represented in the topology.

- [Prerequisites, on page 1](#)
- [Sample Scenarios, on page 4](#)
- [VRF Lite Through the DCNM GUI – From a BGW Device to a Nexus 7000 Series Edge Router , on page 5](#)
- [VRF Lite Through the DCNM GUI – From a BGW Device To a Non-Nexus Device , on page 17](#)
- [Automatic VRF Lite \(IFC\) Configuration, on page 24](#)
- [Deleting VRF Lite IFCs, on page 27](#)
- [Additional References, on page 29](#)
- [Appendix , on page 29](#)

Prerequisites

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 Configuration Guide, Release 11.2(1).
- Ensure that the role of the designated border device is Border, Border Spine, Border Gateway, or Border Gateway Spine (a switch on which Multi-Site and VRF Lite functions co-exist). 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 role is inappropriate for a border device, set the appropriate role.
- Create an external fabric. If you connect the VLXAN fabric border device to a Nexus 7000 Series switch (or other Nexus device) for external connectivity, add the Nexus 7000 series switch to the external fabric and set its role to **Edge Router**. In DCNM, you can import switches to an external fabric, and update selected configurations. For details, refer the Creating an External Fabric section in the Control chapter.
- To allow inter-subnet communication between end hosts in different VXLAN fabrics, where the subnets are present in both fabrics, you must disable the **Advertise Default Route** feature for the associated VRF. This will result in /32 routes for hosts being seen in both fabrics. For example, Host1 (VNI 30000, VRF 50001) in Fabric1 can send traffic to Host2 (VNI 30001, VRF 50001) in Fabric2 only if the host route is present in both fabrics. When a subnet is present in only one fabric, then default route is sufficient for inter-subnet communication. Steps:
 1. Go to the fabric's **VRFs** screen and select the VRF.
 2. Click the **Edit** option at the top left part of the screen.
 3. In the **Edit VRF** screen, click **Advanced** in the VRF Profile section.
 4. Clear the **Advertise Default Route** checkbox and click **Save**.

The screenshot displays the configuration interface for a VRF. It is divided into two main sections: 'VRF Information' and 'VRF Profile'.

VRF Information:

- VRF ID:** 50000
- VRF Name:** CNC-Prod-Pod2
- VRF Template:** Default_VRF_Universal
- VRF Extension Template:** Default_VRF_Extension_Universal
- VLAN ID:** (empty field) with a 'Propose VLAN' button and a help icon.

VRF Profile:

The 'Advanced' tab is selected. The fields include:

- RP Loopback ID:** (empty field) with a help icon and value '0-1023'.
- Underlay Mcast Add...:** (empty field) with a help icon and value 'IPv4 Multicast Address'.
- Overlay Mcast Groups:** (empty field) with a help icon and value '224.0.0.0/4 to 239.255.255.255/4'.
- Enable IPv6 link-loc...:** with a help icon and text 'Enables IPv6 link-local Option under VRF SVI'.
- Enable TRM BGW MSite:** with a help icon and text 'Enable TRM on Border Gateway Multisite'.
- Advertise Host Routes:** with a help icon and text 'Flag to Control Advertisement of /32 and /128 Routes to Edge Routers'.
- Advertise Default Route:** with a help icon and text 'Flag to Control Advertisement of Default Route Internally'.
- Config Static 0/0 Route:** with a help icon and text 'Flag to Control Static Default Route Configuration'.

At the bottom right, there are 'Save' and 'Cancel' buttons.

The following options apply only when VRF Lite connectivity is enabled on the border devices. By default, following Cisco best practices, DCNM uses eBGP over sub-interfaces for VRF Lite, Option-A peering. In other words, for each VRF Lite Inter-fabric connection (IFC), there is a per VRF per peer eBGP peering session established over IPv4/IPv6 respectively from the border device to the edge/WAN router. As applicable to this VRF Lite peering, there are 3 fields:

- **Advertise Host Routes** – By default, over the VRF Lite peering session, only non-host (/32 or /128) prefixes are advertised. But if host routes (/32 or /128) need to be enabled and advertised from the border device to the edge/WAN router, then the “**Advertise Host Routes**” check box can be enabled. Route-map does outbound filtering. By default, this check box is disabled.
- **Advertise Default Route** – This field controls whether a network statement 0/0 will be enabled under the vrf. This in turn will advertise a 0/0 route in BGP. By default, this field is enabled. When the check box is enabled, this will ensure that a 0/0 route is advertised inside the fabric over EVPN Route-type 5 to the leafs thereby providing a default route out of the leafs toward the border devices.
- **Config Static 0/0 Route** –The field controls whether a static 0/0 route to the edge/WAN router, should be configured under the VRF, on the border device. By default, this field is enabled. If WAN/edge routers are advertising a default route over the VRF Lite peering, to the border device in the fabric, then this field should be disabled. In addition, the “Advertise Default Route” field should also be disabled. This is because the 0/0 route advertised over eBGP will be sent over EVPN to the leafs without the need for any additional configuration. The clean iBGP EVPN separation inside the fabric with eBGP for external out-of-fabric peering, provides for this desired behavior.

Note that all of the options listed are per fabric fields. Hence, in Multi-Site deployments with MSD, these fields can be controlled at a per member fabric level.

5. Follow this procedure for all VRFs deployed on the VXLAN fabrics' border devices connected through VRF Lite.



Note If you create a new VRF, ensure that you clear the **Advertise Default Route** checkbox.



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

Sample Scenarios

Scenarios explained in this document:

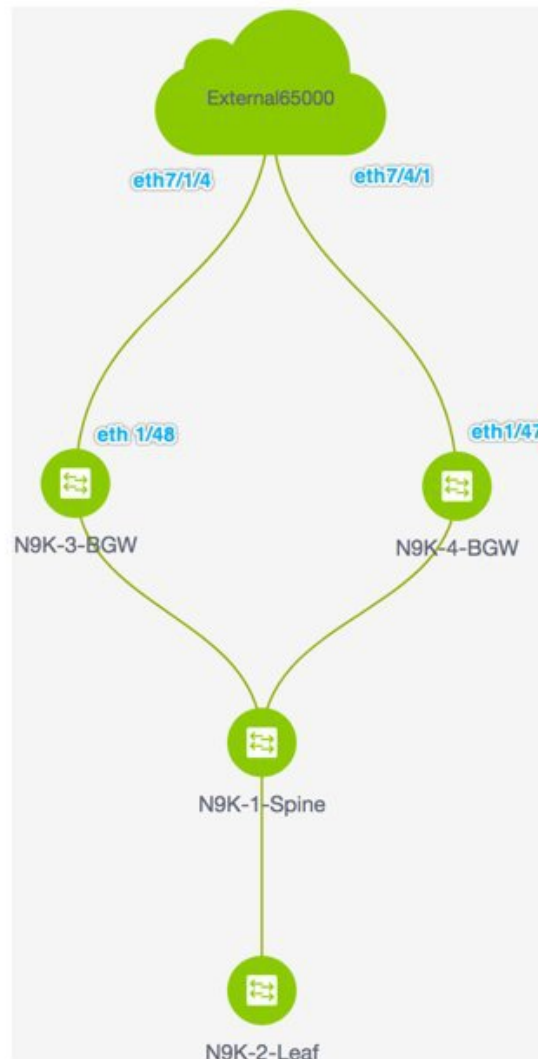
- VRF Lite through the DCNM GUI – From a BGW device to a Nexus 7000 Series edge router.
- VRF Lite through the DCNM GUI – From a BGW device to a non-Nexus device.
- Automatic VRF Lite (IFC) Configuration.



Note

- The sample scenarios are shown using a Border Gateway role but are equally applicable to the Border nodes as well.
- Anything that applies to Border or Border Gateway roles also applies to Border Spine and Border Gateway Spine roles.

VRF Lite Through the DCNM GUI – From a BGW Device to a Nexus 7000 Series Edge Router



- The topology displays the VXLAN BGP EVPN fabric **Easy7200** connected to the external fabric **External65000** (the cloud icon). The BGWs of the VXLAN fabric are connected to the edge router **n7k1-Edge1** (not visible in the image) in the external fabric.
- The BGWs 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. Network configurations for the VXLAN fabric are 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. This document shows you how to enable Layer 3 north-south traffic between the VXLAN fabric and the external fabric.

- Before VRF Lite configuration, end hosts associated with a specific VRF can send traffic to each other, but only within the fabric. After VRF Lite configuration, end hosts can send traffic outside the VXLAN fabric, towards other (VXLAN or classic LAN) fabrics

Enabling the VRF Lite feature

For this example, we will enable connectivity between Easy7200 and External65000. The steps:

Step 1 - Deploy IFC prototypes on physical interfaces, on N9K-3-BGW and N9K-4-BGW.

Step 2 - Deploy the individual VRF extensions on the BGWs N9K-3-BGW and N9K-4-BGW.

Step 3 - Deploy VRF extensions on the edge router n7k1-Edge1.

The third step completes the configuration between **Easy7200** and **External65000**.

Step 1 – Deploying IFC prototypes on physical interfaces on N9K-3-BGW and N9K-4-BGW

For VRF Lite configuration, you should enable eBGP peering between the fabric’s BGW interfaces and the edge router’s interfaces, through point-to-point connections. The BGW physical interfaces are:

- **eth 1/48** on **N9K-3-BGW**, towards **eth 7/1/4** on **n7k1-Edge1**.
- **eth 1/47** on **N9K-4-BGW**, towards **eth 7/4/1** on **n7k1-Edge1**.



Note You can also enable VRF Lite in a back-to-back topology wherein Border/Border Gateways are directly connected to each other.

1. Click **Control > Fabric Builder**. The Fabric Builder screen comes up.
2. Click the **Easy7200** box. The fabric topology comes up.
3. Click **Tabular view**. The **Switches | Links** screen comes up.

The **Links** tab lists fabric links. Each row either represents a link between two devices within **Easy7200** or a link from a device in **Easy7200** to an external fabric.



Note An inter-fabric link is a physical connection between two Ethernet interfaces or a virtual connection (such as a fabric overlay between two loopback interfaces). When you add a physical connection between devices, the new link appears in the **Links** tab by default.

4. Select the link checkbox (that represents the connection between **eth 1/48** on **N9K-3-BGW**, towards **eth 7/1/4** on **n7k1-Edge1**) and click the Edit icon at the top left part of the screen.

	Scope	Name	Policy	Info	Admin State	Oper State
	Easy7200	N9K-2-Leaf-Ethernet1/47---N9K-1-Spine-Ethernet1/47	int_intra_fabric_num_link_11_1	Link Present	Up:Up	Up:Up
1	Easy7200<->External65000	N9K-3-BGW-Ethernet1/48---n7k1-Edge1-Ethernet7/1/4	int_intra_fabric_num_link_11_1	Link Present	Up:Up	Up:Up
	Easy7200	N9K-3-BGW-Ethernet1/47---N9K-1-Spine-Ethernet1/43	int_intra_fabric_num_link_11_1	Link Present	Up:Up	Up:Up
4	Easy7200<->External65000	N9K-4-BGW-Ethernet1/47---n7k1-Edge1-Ethernet7/4/1		Link Present	Up:Up	Up:Up
5	Easy7200<->Easy60000	N9K-4-BGW-Ethernet1/2---N9K-15-BGW-Ethernet1/8		Link Present	Up:Up	Up:Up
6	Easy7200	N9K-4-BGW-Ethernet1/48---N9K-1-Spine-Ethernet1/42	int_intra_fabric_num_link_11_1	Link Present	Up:Up	Up:Up

The fields are:

Scope – The source and destination fabrics are displayed. For an intra-fabric link, only one fabric name (**Easy7200**) is displayed since the source and destination interfaces are part of the same fabric. An inter-fabric link is displayed as **Easy7200 <->External65000**.

Name – The name is formed with the following syntax:

source device ~ source interface --- destination device ~ destination interface.

So, the entry is **N9K-4-BGW ~ Ethernet1/47 --- n7k1-Edge1 ~ Ethernet7/4/1**.

Policy – The policy used for creating VRF Lite, ext_fabric_setup_11_1 is displayed.

Info – This displays the status of the link (Link Present, Neighbor Present, Neighbor Missing, etc).

Admin State – This displays the administrative state of the link (Up, Down, etc).

Oper State – This displays the operational state of the link (Up, Down, etc).

The **Link Management – Edit Link** comes up.

Link Management - Edit Link

- * Link Type: Inter-Fabric
- * Link Sub-Type: VRF_LITE
- * Link Template: ext_fabric_setup_11_1
- * Source Fabric: Easy7200
- * Destination Fabric: External65000
- * Source Device: n9k-2-leaf
- * Source Interface: Ethernet2/1
- * Destination Device: n9k-18-RS
- * Destination Interface: Ethernet1/2

Link Profile

- General
- Advanced

- * BGP Local ASN: 7200 ? Local BGP Autonomous System Number
- * IP Address/Mask: 2.2.2.2/24 ? IP address for sub-interface in each VRF
- * BGP Neighbor IP: 2.2.2.1 ? Neighbor IP address in each VRF
- * BGP Neighbor ASN: 65000 ? Neighbor BGP Autonomous System Number

Some fields are explained:

Link Sub-Type - By default, the **VRF_LITE** option is displayed.

Link Template – The default template for a VRF Lite IFC, **ext_fabric_setup_11_1**, is displayed. The template enables the source and destination interfaces as Layer 3 interfaces, configures the **no shutdown** command, and sets their MTU to 9216.

You can edit the **ext_fabric_setup_11_1** template or create a new one with custom configurations.

In the **General** tab, the BGP AS numbers of **Easy7200** and **External65000** are displayed. Fill in the other fields as explained.

▼ Link Profile

General

Advanced

* BGP Local ASN 7200 ? Local BGP Autonomous System Number

* IP Address/Mask 2.2.2.2/24 ? IP address for sub-interface in each VRF

* BGP Neighbor IP 2.2.2.1 ? Neighbor IP address in each VRF

* BGP Neighbor ASN 65000 ? Neighbor BGP Autonomous System Number

IP Address/Mask – Enter the IP address prefix to assign an IP address for the **Ethernet 1/48** sub interfaces, the source interface of the IFC. A sub-interface is created for each VRF extended over this IFC, and a unique 802.1Q ID is assigned to it. The IP address/Mask entered here, along with the BGP Neighbor IP field (explained below) will be used as the default values for the sub-interface created at VRF extension and can be overwritten.

For example, an 802.1Q ID of 2 is associated with subinterface Eth 1/48.2 for VRF 50000 traffic, and 802.1Q ID of 3 is associated with Eth 1/48.3 and VRF 50001, and so on.

(The VRF extension deployment is explained in a subsequent section).

The IP prefix is reserved with the DCNM resource manager. Ensure that you use a unique IP address prefix for each IFC you create in the topology.

BGP Neighbor IP – Enter the IP address of the eBGP neighbor for each VRF extension deployed on this IFC, on the **N9K-3_BGW** end.

Inter-fabric traffic from VRFs for an IFC will have the same source IP address (**2.2.2.2/24**) and destination IP address (**2.2.2.1**).


The **Advanced** tab has been added in the **Link Profile** section.

This tab contains the following fields:

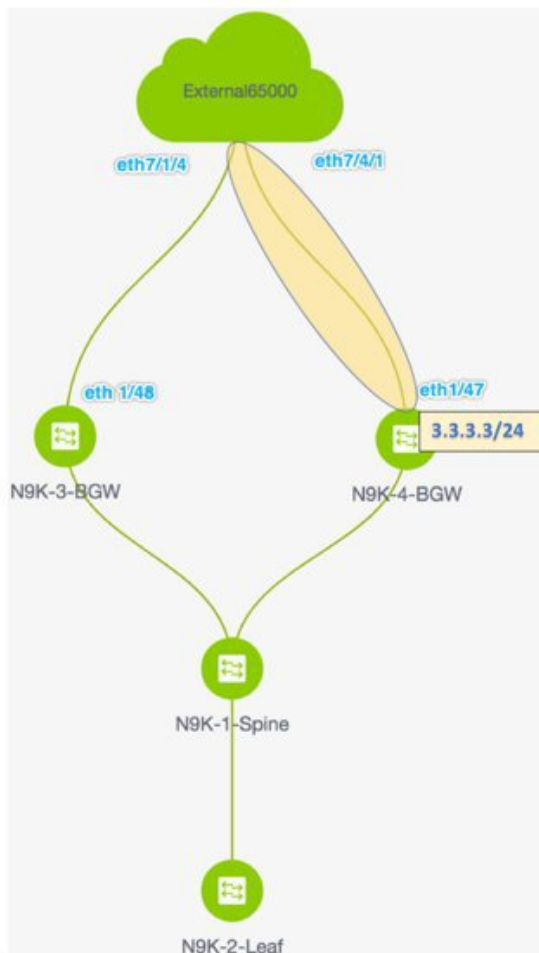
- **Source Interface Description**
- **Destination Interface Description**
- **Source Interface Freeform Config**
- **Destination Interface Freeform Config**

5. Click **Save** at the bottom right part of the screen.

The **Switches|Links** screen comes up again. You can see that the IFC entry is updated with the VRF Lite policy template used for creating the IFC, **ext_fabric_setup_11_1**. A representation of the topology is shown below.



6. Similarly, create an IFC from **eth 1/47** on **N9K-4-BGW** towards **eth 7/4/1** on **n7k1-Edge1**. An entry is seen in the **Links** screen. A representation of the topology is shown below.



7. Click **Save and Deploy** at the top right part of the screen.

The **Links** tab after executing **Save and Deploy** looks like this. The links on which IFC has deployed have the relevant policy configured in the **Policy** column.

SCOPE: Easy7200 admin

Fabric Builder: Easy7200 Save & Deploy

Switches Links

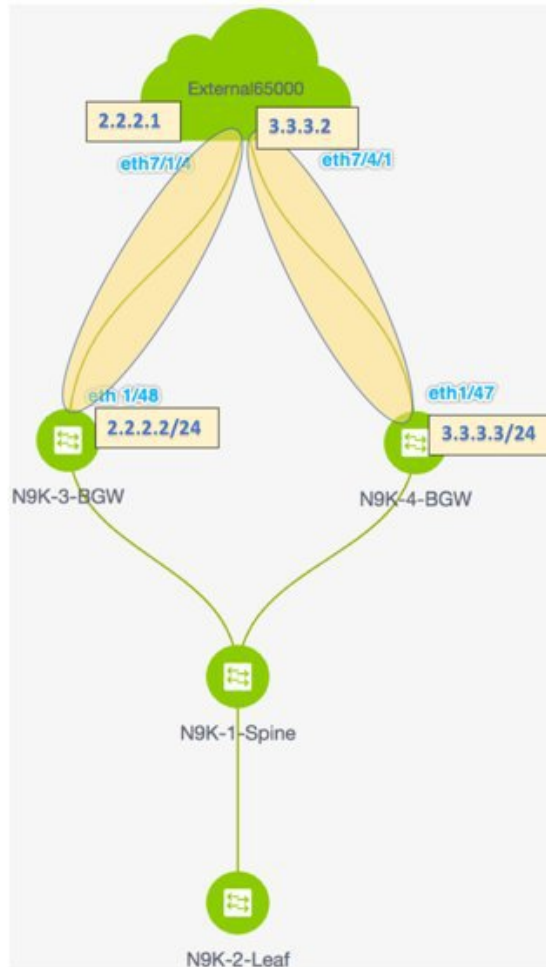
Scope	Name	Policy	Info	Admin State	Oper St
1	Easy7200->External65000 N9K-3-BGW-Ethernet1/48--n7k1-Edge1-Ethernet7/1/4	ext_fabric_setup_11_1	Link Present	Up:Up	Up:Up
2	Easy7200->External65000 N9K-4-BGW-Ethernet1/47--n7k1-Edge1-Ethernet7/4/1	ext_fabric_setup_11_1	Link Present	Up:Up	Up:Up
3	Easy7200 N9K-3-BGW-Ethernet1/47--N9K-1-Spine-Ethernet1/43	int_intra_fabric_num_link_11_1	Link Present	Up:Up	Up:Up
4	Easy7200 N9K-4-BGW-Ethernet1/48--N9K-1-Spine-Ethernet1/42	int_intra_fabric_num_link_11_1	Link Present	Up:Up	Up:Up
5	Easy7200 N9K-2-Leaf-Ethernet1/47--N9K-1-Spine-Ethernet1/47	int_intra_fabric_num_link_11_1	Link Present	Up:Up	Up:Up

8. Go to the **Scope** drop down box at the top right part of the screen and choose **External65000**. The external fabric **Links** screen is displayed. You can see that the two IFCs created from **Easy7200** to **External65000** is displayed here.



Note When you create an IFC or edit its setting in the VXLAN fabric, the corresponding entry is automatically created in the connected external fabric.

9. Click **Save and Deploy** to save the IFCs creation on **External65000**.



Base configurations – For VRF Lite to function, appropriate route maps and policies that apply to VRFs have to be deployed on the border devices **N9K-3-BGW** and **N9K-4-BGW**. You do not need to manually enable the base configurations. They are automatically deployed via a default template **ext_base_border_vrflite_11_1**.

For a device with a Border Leaf or Border Spine role, the base configurations are deployed when you execute the **Save and Deploy** operation (available in the fabric topology screen [via the **Fabric Builder** screen > Fabric Box]) for the first time in a fabric.

For a Border Gateway or Border Gateway Spine role, the base configurations are deployed when you deploy the first VRF Lite IFC on the device.

You need to modify the **ext_base_border_vrflite_11_1** template for specific needs before deployment or its policy should be deleted, template modified, and then deploy the template again. The configurations are noted in the **Appendix** section.

The first step in the VRF Lite configuration scenario, creating IFCs on the border devices and edge router, is complete. Next, the VRF extensions are deployed on the switches.

Step 1 - Deploy IFC prototypes on physical interfaces, on **N9K-3-BGW** and **N9K-4-BGW**.

Step 2 - Deploy the individual VRF extensions on the BGWs **N9K-3-BGW** and **N9K-4-BGW**.

Step 3 - Deploy VRF extensions on the edge router **n7k1-Edge1**.

The third step completes the configuration between **Easy7200** and **External65000**.


Step 2 - Deploy the individual VRF extensions on the BGWs **N9K-3-BGW** and **N9K-4-BGW**

During the IFC creation process, base configurations are created, and IP addresses are reserved for the interfaces that transport the inter-fabric traffic on **N9K-3-BGW** and **N9K-4-BGW**. In this step, the VRF and VRF extension configuration is deployed on the interfaces.

To extend VRFs beyond the fabric, the VRFs should have been created and deployed on relevant fabric devices, except the border devices.

The steps are:

1. Click **Control > Networks and VRFs**. The **Networks & VRFs** screen comes up.
2. Click **Continue**. The **Select a Fabric** screen comes up.
3. Select **Easy7200** and click **Continue** at the top right part of the screen.



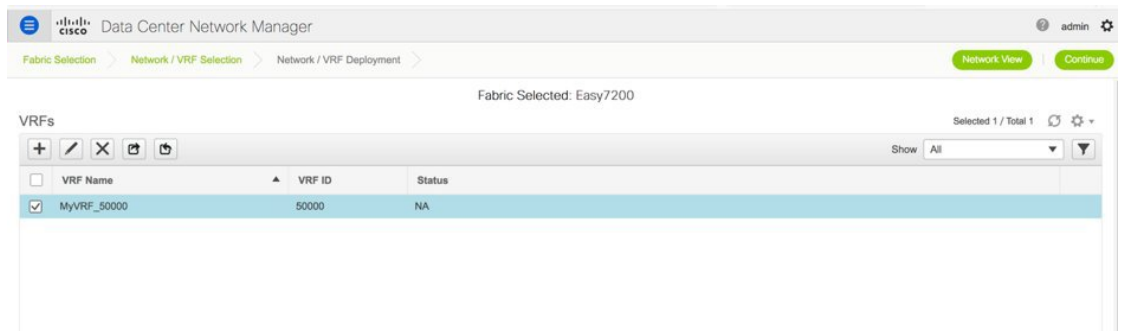
Select a Fabric

Choose a fabric with appropriate switches where you want the Top Down functionality to be enabled

Easy7200

The **Networks** screen comes up.

4. Click **VRFs** at the top right part of the screen. The **VRFs** screen comes up.
5. Select the VRF that you want to deploy (**MyVRF_5000** in this case) and click **Continue** at the top right part of the screen.

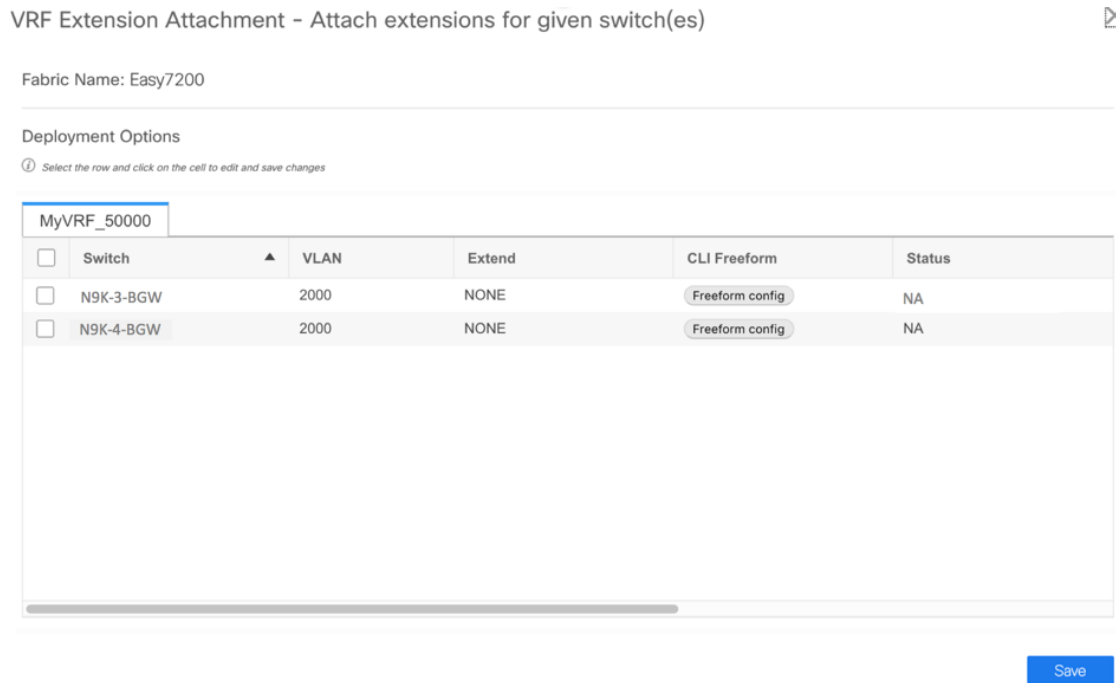


The **Easy7200** fabric topology comes up.

6. Select the **Multi-Select** checkbox at the top right part of the screen and drag the cursor across the BGWs on which you want to deploy the VRF and VRF extension configuration.



The **VRF Extension Attachment** screen comes up. Each row represents a switch and each tab a VRF. Update settings for each tab as explained.



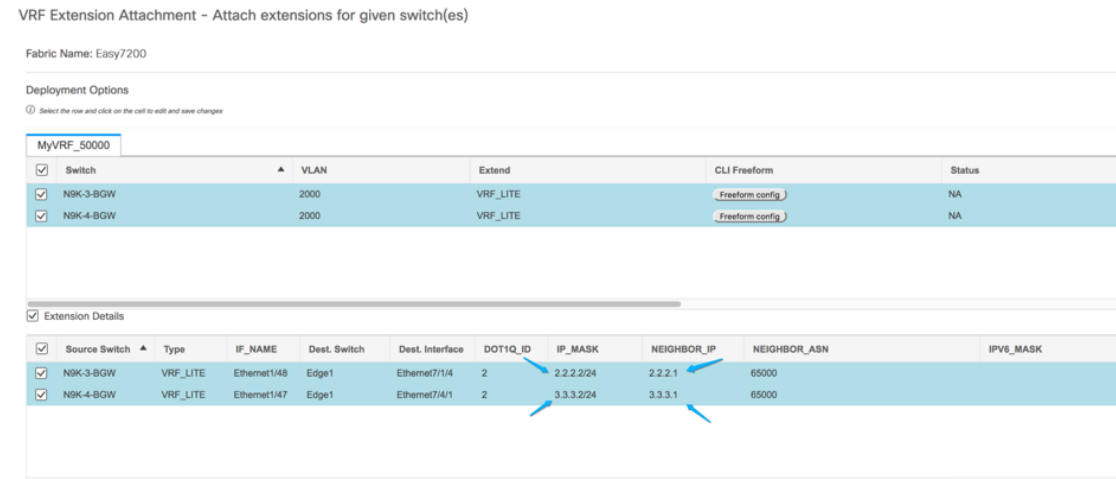
In the **Extend** column, click on **NONE** and choose the **VRF_LITE** option from the drop down box. Do this for the second row too.

Select the checkboxes in both rows.

The **Extension Details** section comes up at the bottom of the screen. It displays the IFCs created on the selected switches, wherein each row represents an IFC.

Select the IFC check boxes in both rows.

After selecting the IFCs, the screen looks like this.



Click **Save** at the bottom right part of the screen.

The fabric topology screen comes up.

7. Click the **Preview** option at the top right part of the screen to preview VRF and VRF extension configuration.
8. Click **Deploy** at the top right part of the screen.

At the bottom right part of the screen, the color codes that represent different stages of deployment are displayed. The color of the switch icons changes accordingly (Blue for Pending state, yellow for In Progress state when the provisioning is in progress, red for failure state, green when successfully deployed).

When the switch icons turn green, it means that the VRFs are successfully deployed.

The second step in the VRF Lite configuration scenario, deploying VRF extensions on the border devices is complete. Next, the VRF extensions are deployed on the edge router **n7k1-Edge1**.

Step 1 - Deploy IFC prototypes on physical interfaces, on **N9K-3-BGW** and **N9K-4-BGW**.

Step 2 - Deploy the individual VRF extensions on the BGWs **N9K-3-BGW** and **N9K-4-BGW**.

Step 3 - Deploy VRF extensions on the edge router n7k1-Edge1.

The third step completes the configuration between **Easy7200** and **External65000**.

Step 3 - Deploy VRF extensions on the edge router n7k1-Edge1


In order to extend VRFs on the edge router, keep a note of the following fields. VRF extension on the border device is on a per interface basis.

- **IP_MASK** - This will become the neighbor address at the edge router end and mask will be the local mask on the edge router. This is derived from the IFC prototype created in the earlier step.
- **Easy Fabric ASN** - This will become neighbor ASN from the edge router end. This is derived from the IFC prototype created in the earlier step.
- **Dot1Q tag** - This will be same on the edge router. This is derived from the VRF extension table.
- **Neighbor ASN** - This will become LOCAL ASN on the edge router. IFC prototype.
- **Neighbor IP** - This will become Local IP for sub-interface on the edge router. IFC prototype.
- **Destination port** - Will be local port on edge router upon which extension will be deployed.

You have deployed VRF extensions for **MyVRF_50000** from the BGWs **N9K-3-BGW** and **N9K-4-BGW**. Now, you should deploy the VRF extensions on the other end of the links, on **n7k1-Edge1**. In DCNM, the CLI template used for this is **External_VRF_Lite_eBGP**.

eBGP configuration on the edge router

1. In the **External65000** fabric topology screen, click **Tabular view**.
The **Switches | Links** screen comes up.
2. Select the switch checkbox and click the **View/Edit Policies** button.



The **View/Edit Policies** screen comes up.

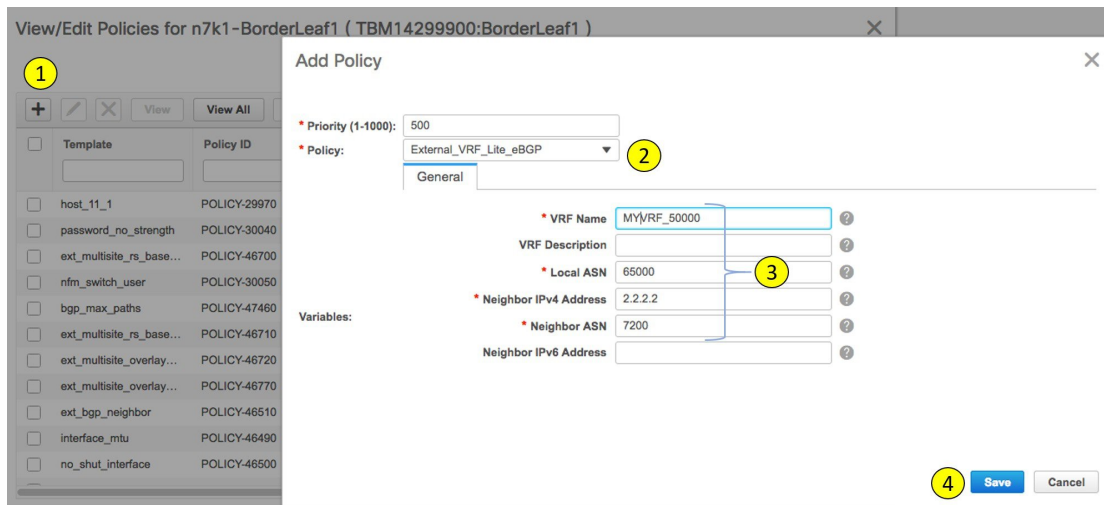
3. Click + at the top left part of the screen to add a policy, and fill in the **Add Policy** screen as shown in the image.

You can use a user defined template too in the **Policy** field.



Note Note the policy ID for this VRF extension. It is useful when deleting the policy to remove the extension, when applicable.


This defines a policy from the edge router towards **N9K-3-BGW**.



4. As per the earlier steps, create a policy for the VRF extension towards **N9K-4-BGW**. The **Neighbor IPv4 Address** field for the second extension is updated with 3.3.3.3.

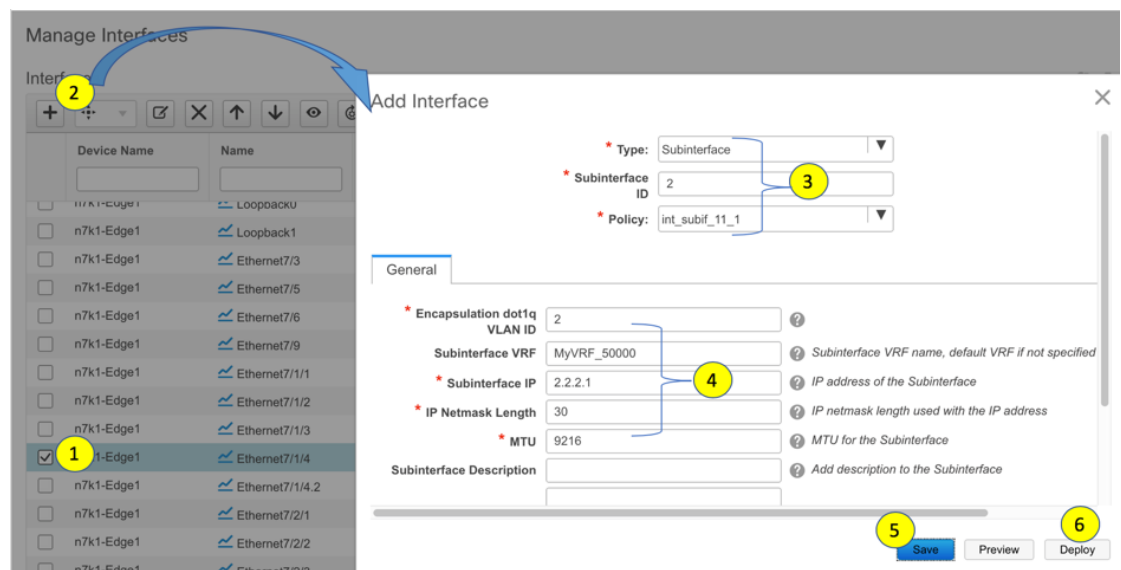
Sub interface policy on Edge Router

1. In the **External65000** fabric topology screen, click **Tabular view**.
The **Switches | Links** screen comes up.
2. Select the switch checkbox and click the **Manage Interfaces** button.



The **Manage Interfaces** screen comes up.

- As shown in the image, select the interface connected to the border device (in this case **Eth7/1/4**), and click + at the top left part of the screen. Then, fill the **Add Interface** screen from corresponding IFC and VRF extensions on the border device.



The example shows a breakout port on the Cisco Nexus 7000 Series switch. This breakout must be performed using the DCNM breakout policy (the template name is **breakout_interface**). If this is not done, the subinterface deletion is blocked by DCNM.

- Click **Save** to save the settings, and **Deploy** to deploy the settings onto the switch.
- As explained in the earlier steps, create another subinterface policy for the VRF extension towards **N9K-4-BGW**. The **Subinterface IP** field for the second extension is updated with 3.3.3.1.

The third step in the VRF Lite configuration scenario, deploying VRF extensions on the edge router **N7k1-Edge1** is complete. This step completes the configuration between **Easy7200** and **External65000**.

VRF Lite Through the DCNM GUI – From a BGW Device To a Non-Nexus Device

In this case, the non-Nexus device is an ASR 9000 Series router, **ASR9K-1-Edge** which is connected to the BGW **N9K-3-BGW** in the **Easy7200** fabric. The router is not imported through DCNM nor discovered via

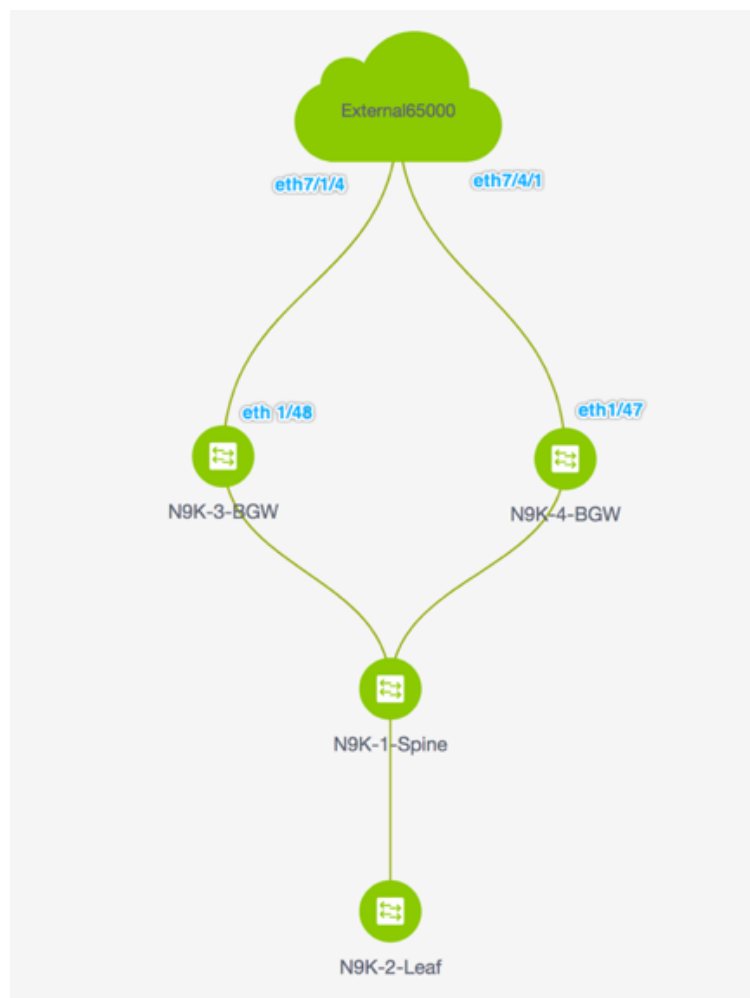
CDP or LLDP. To represent the non-Nexus device, you must create an external fabric. Refer the **Creating an External Fabric** topic to know how to create an external fabric. For this example, the external fabric **External65000** is created.

The device and connection are displayed in the DCNM topology after the IFC creation between **ASR9K-1-Edge** and **N9K-3-BGW**.



Note A connected non-Cisco device can also be represented in the topology.

The topology:



The steps are:

Step 1 - Deploy an IFC prototype on the N9K-3-BGW physical interface that connects to ASR9K-1-Edge.

Step 2 - Deploy the individual VRF extensions on N9K-3-BGW.

This step completes the configuration between **Easy7200** and the non-Nexus device.

Step 1 - Deploy an IFC prototype on the N9K-3-BGW physical interface that connects to ASR9K-1-Edge

For VRF Lite configuration, you should enable eBGP peering between the fabric's BGW interface and the **ASR9K-1-Edge** interface, through a point-to-point link.

1. Click **Control > Fabric Builder**. The **Fabric Builder** screen comes up.
2. Click the rectangular box that represents the **Easy7200** fabric. The fabric topology screen comes up.
3. Click **Tabular view**. The **Switches | Links** screen comes up.

The **Links** tab lists fabric links. Each row either represents a link between two devices within **Easy7200** or a link from a device in **Easy7200** to an external fabric.

4. Click + to add a new link. The **Link Management – Add Link** screen comes up.

Link Management - Add Link

* Link Type: Intra-Fabric

* Link Sub-Type: Fabric

* Link Template: int_intra_fabric_num_link_11_1

* Source Fabric: Easy7200

* Destination Fabric:

* Source Device:

* Source Interface:

* Destination Device:

* Destination Interface:

Link Profile

General

* Source IP: ? IP address of the source interface

* Destination IP: ? IP address of the destination interface

Interface Admin State: ? Admin state of the interface

Save

Fill or choose the fields as noted:

Link Type – Choose **Inter-Fabric**.

Link Sub-Type – **VRF_Lite** is displayed by default.

Link Template - By default, the **ext_fabric_setup_11_1** template is populated.



Note You can add, edit, or delete user-defined templates. See **Template Library** section in the **Control** chapter for more details.

Source Fabric - **Easy7200** is selected by default.

Destination Fabric – Select **External65000**.

Source Device and **Source Interface** - Choose the BGW and the interface that connects to the ASR device.

Destination Device and **Destination Interface**— Destination device and interface do not appear in the drop down box. Type any string here that will help identify the device. This name appears in the external fabric topology screen in the **Fabric builder** screen.

General tab in the Link Profile section.

BGP Local ASN - In this field, the AS number of the source fabric Easy7200 is autopopulated.

IP Address/Mask - Enter the IP address and mask that is used in the VRF Extension Sub-interfaces.

BGP Neighbor IP - Enter the IP address that is used on the External box as local interface address for the VRF Extensions.

BGP Neighbor ASN - In this field, the AS number of the external fabric External65000 is autopopulated since we selected it as the external fabric.

After filling up the **Add Link** screen, it looks like this:

The screenshot displays the 'Link Management - Add Link' configuration window. The form includes the following fields:

- Link Type:** Inter-Fabric
- Link Sub-Type:** VRF_LITE
- Link Template:** ext_fabric_setup_11_1
- Source Fabric:** Easy7200
- Destination Fabric:** External65000
- Source Device:** N9K-3-BGW
- Source Interface:** Ethernet1/5
- Destination Device:** ASR9K-1-Edge
- Destination Interface:** Ethernet1/5

The 'Link Profile' section is expanded to the 'General' tab, showing the following BGP configuration:

- BGP Local ASN:** 7200 (Local BGP Autonomous System Number)
- IP Address/Mask:** 5.5.5.2/24 (IP address for sub-interface in each VRF)
- BGP Neighbor IP:** 5.5.5.1 (Neighbor IP address in each VRF)
- BGP Neighbor ASN:** 65000 (Neighbor BGP Autonomous System Number)

A 'Save' button is located at the bottom right of the window.

- Click **Save** at the bottom right part of the screen.

The **Switches|Links** screen comes up again. You can see that the IFC entry is updated.

- Click **Save and Deploy** at the top right part of the screen.

The links on which the IFC is deployed has the relevant policy (**ext_fabric_setup_11_1**) configured in the **Policy** column.

- Go to the **Scope** drop down box at the top right part of the screen and choose **External65000**. The external fabric **Links** screen is displayed. You can see that the IFC created from **Easy7200** to the ASR device is displayed here.


- Click **Save and Deploy**.

The first step in the VRF Lite configuration scenario from a BGW to a non-Nexus device is complete. Next, the VRF extensions are deployed on the BGW towards the ASR device.

Step 2 - Deploy the individual VRF extensions on N9K-3-BGW

To extend VRFs beyond the fabric, the VRFs should have been created and deployed on relevant fabric devices, excepting the border devices.

1. Click **Control** > **Networks and VRFs**. The **Networks & VRFs** screen comes up.
2. Click **Continue**. The **Select a Fabric** screen comes up.
3. Select **Easy7200** and click **Continue** at the top right part of the screen.




Select a Fabric

Choose a fabric with appropriate switches where you want the Top Down functionality to be enabled

Easy7200

The **Networks** screen comes up.

4. Click **VRFs** at the top right part of the screen. The **VRFs** screen comes up.
5. Select the VRF that you want to deploy (**MyVRF_5000** in this case) and click **Continue** at the top right part of the screen.



The Easy7200 fabric topology comes up.

6. Double-click the **N9K-3-BGW** icon on which you want to deploy the VRF and VRF extension configuration.

The **VRF Extension Attachment** screen comes up. Each row represents a switch and each tab a VRF. Only one VRF is extended in this example.

VRF Extension Attachment - Attach extensions for given switch(es)

Fabric Name: Easy7200

Deployment Options

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

MyVRF_50000

<input type="checkbox"/>	Switch ▲	VLAN	Extend	CLI Freeform	Status
<input type="checkbox"/>	N9K-3-BGW	2000	NONE	Freeform config	NA

[Save](#)

In the **Extend** column, click on **NONE**. A drop down box appears. Choose the **VRF_LITE** option, and click outside the row.

Select the checkbox next to the switch.

The **Extension Details** section comes up at the bottom of the screen. It displays the IFCs created on the selected switches, wherein each row represents an IFC.

Select the IFC check box. After selecting the IFCs, the screen looks like this.

VRF Extension Attachment - Attach extensions for given switch(es) ✕

Fabric Name: Easy7200

Deployment Options

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

<input type="checkbox"/>	Switch	VLAN	Extend	CLI Freeform	Loopback Id	Loopback IPv4 Address	Lo
<input checked="" type="checkbox"/>	N9K-3...	2000	VRF_LITE	Freeform config			

Extension Details

<input checked="" type="checkbox"/>	Sourc...	type	IF_NAME	Dest. Switch	Dest. Interface	DOT1Q_I
<input checked="" type="checkbox"/>	N9K-3...	VRF_LITE	Ethernet1/48	Edge1	Ethernet7/1/4	2

[Save](#)

Click **Save** at the bottom right part of the screen.

The fabric topology screen comes up.

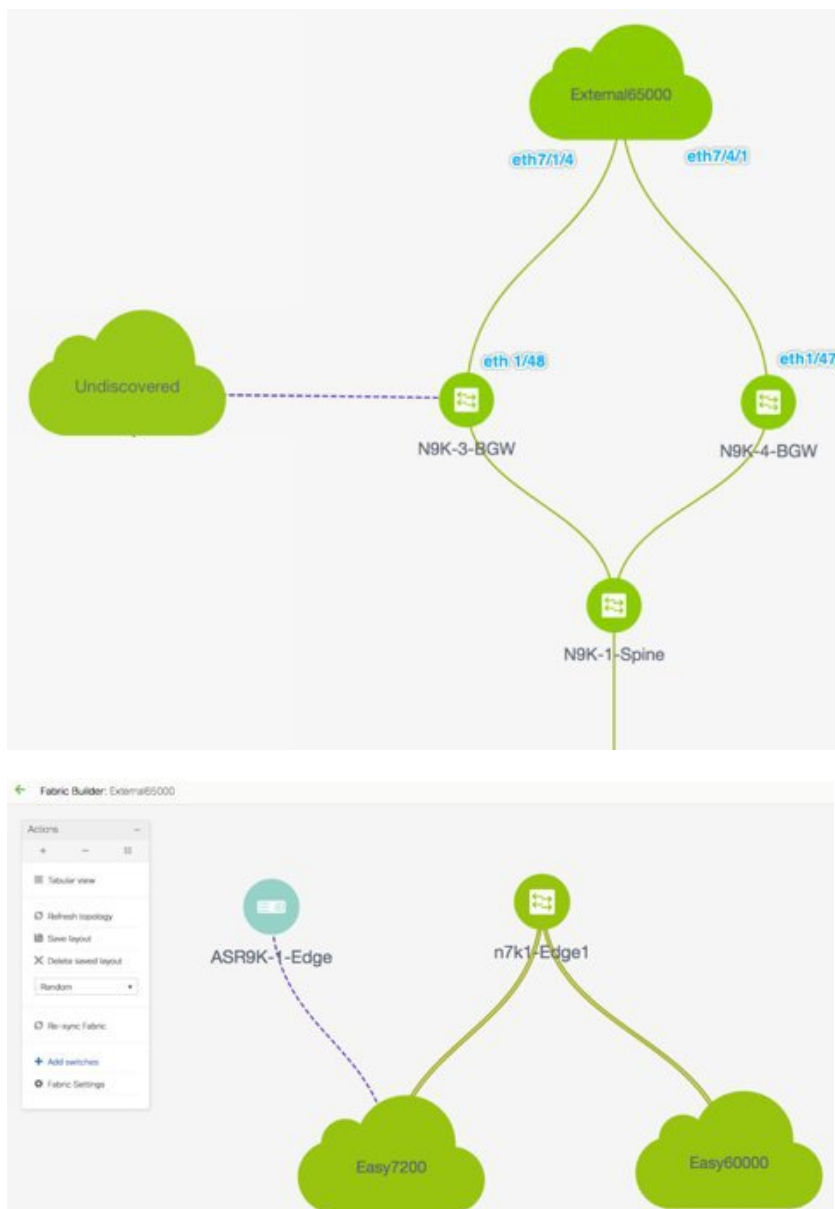
7. Click the **Preview** option at the top right part of the screen to preview VRF and VRF extension configuration.
8. Click **Deploy** at the top right part of the screen.

At the bottom right part of the screen, the color codes that represent different stages of deployment are displayed. The color of the switch icons changes accordingly (Blue for Pending state, yellow for In Progress state when the provisioning is in progress, red for failure state, green when successfully deployed, and so on).

When the switch icons turn green, it means that the VRF is successfully deployed.

The second step in the VRF Lite configuration scenario, deploying VRF extensions on the border device towards the non-Nexus ASR device is complete.

The device and connection will display in the **Easy7200** and **External65000** fabrics.

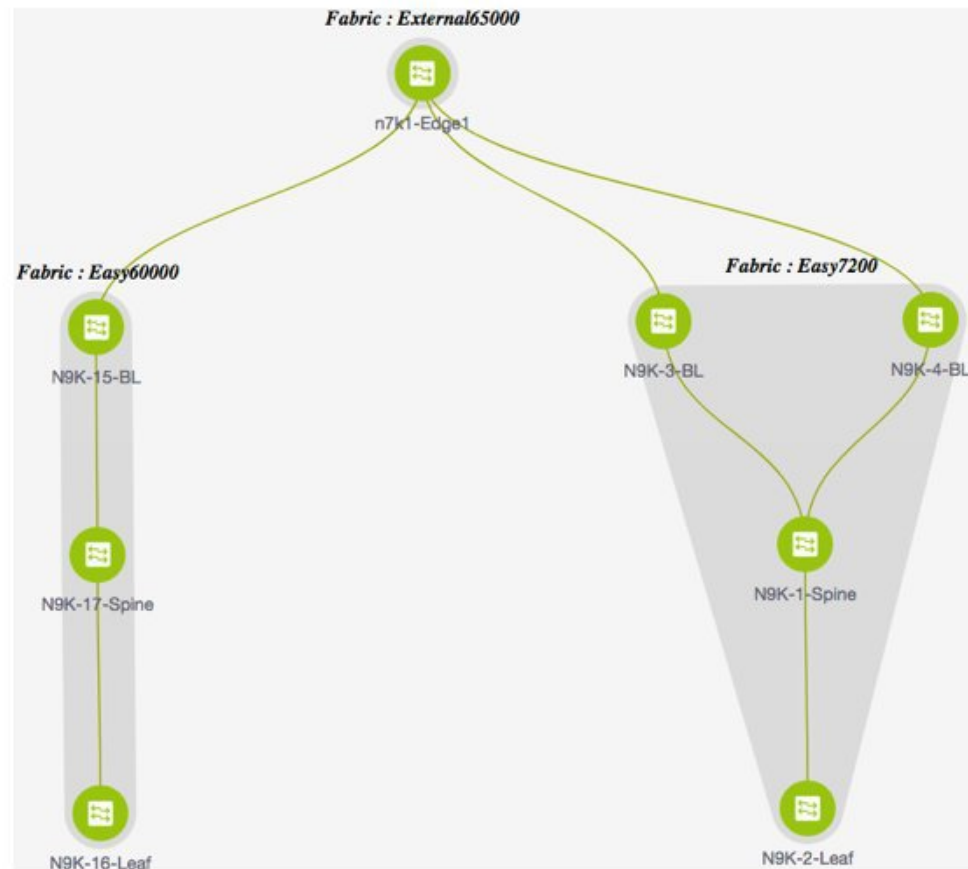


Automatic VRF Lite (IFC) Configuration

You can enable VRF Lite auto-configuration by changing the fabric settings of the **VRF Lite Deployment** field under the **Resources** tab from **Manual** to any of the auto-configuration settings.



Note In the fabric topology screen within **Fabric Builder**, you can view only the individual fabric and the external fabric connected.



- The topology displays VXLAN BGP EVPN fabrics **Easy60000** (at the left) and **Easy7200** (at the right) and external fabric **External65000** (at the top). The border leaf of one VXLAN fabric is connected to the border leaf of the other through the edge router **n7k1-Edge1** in the external fabric.
- The border leafs are special devices that allow clear control and data plane segregation from the fabric to the external Layer 3 domain while allowing for policy enforcement points for any inter-fabric traffic. Multiple border devices in the fabric ensure redundancy in the case of failures and effective load distribution. This document shows you how to enable Layer 3 north-south traffic between the VXLAN fabrics and the external fabric.
- Before VRF Lite configuration, end hosts associated with a specific VRF can send traffic to each other, but only within the fabric. After VRF Lite configuration, end hosts can send traffic across fabrics.
- Network configurations for the VXLAN fabric are provisioned through DCNM.

The template used for VRF Lite IFC auto configuration is **ext_fabric_setup_11_1**. You can edit the **ext_fabric_setup_11_1** template or create a new one with custom configurations.

Automatic VRF Lite Creation Rules

- Ensure that no user policy is enabled on the interface that connects to the edge router. If a policy exists, then the interface will not be configured.
- Auto configuration is provided for the following cases:

- **Border** role in the VXLAN fabric and **Edge Router** role in the connected external fabric device
- **Border Gateway** role in the VXLAN fabric and **Edge Router** role in the connected external fabric device
- **Border** role to another **Border** role directly

Note that auto configuration is not provided between two BGWs.

If you need a VRF Lite between any other roles, then you have to deploy it manually through the DCNM GUI.

- To deploy configurations in the external fabric, ensure that the **Fabric Monitor Mode** check box is cleared in the external fabric settings of the **External65000** fabric. When an external fabric is set to **Fabric Monitor Mode Only**, you cannot deploy configurations on its switches.

There are four modes available for VRF Lite IFC creation.

1. **Manual** - Use the GUI to deploy the VRF Lite IFCs as shown in the earlier section.
2. **To External Only** - Configure a VRF Lite IFC on each physical interface of a border leaf (Spine) device in the VXLAN fabric that is connected to a device with the **Edge Router** role in the external fabric .
3. **Back to Back Only** - Configure VRF Lite IFCs between directly connected border leaf (Spine) device interfaces of different VXLAN fabrics.
4. **Both** - Use this option to configure IFCs for the modes **To External Only** and **Back to Back Only**.



Note DCI subnet is required, even if the VRF Lite mode is **Manual**. This helps with the DCNM resource handling.

The default mode in fabric settings is Manual Mode. In order to change the mode to any of the others, edit the fabric settings. Under the Resources Tab, modify the VRF Lite Deployment field to one of the above mentioned auto config modes. In this example, ToExternalOnly option is chosen.

VRF Lite Subnet IP Range: The IP address for VRF Lite IFC deployment is chosen from this range. The default value is 10.33.0.0/16. Best practice is to ensure that each fabric has its own unique range and distinct from any underlay range in order to avoid possible duplication. These addresses are reserved with the Resource Manager.

VRF Lite Subnet Mask: By default its set to /30 which is best practice for P2P links.

Edit Fabric

* Fabric Name : Easy7200

* Fabric Template : Easy_Fabric_11_1

General | Replication | vPC | Advanced | **Resources** | Manageability | Bootstrap | Configuration Backup

* Layer 2 VXLAN VNI Range : 30000-49000

* Layer 3 VXLAN VNI Range : 50000-59000

* Network VLAN Range : 2300-2999

* VRF VLAN Range : 2000-2299

* Subinterface Dot1q Range : 2-511

* VRF Lite Deployment : ToExternalOnly

* VRF Lite Subnet IP Range : 10.33.0.0/16

* VRF Lite Subnet Mask : 30

Save Cancel

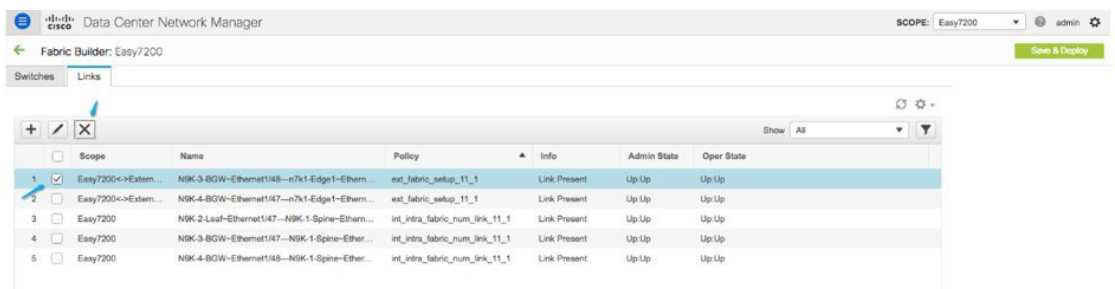
Similarly, update the settings for the Easy60000 fabric too.

Once the fields are set, execute the **Save and Deploy** option in the VXLAN and external fabrics.

Deleting VRF Lite IFCs

Before deleting the IFC, remove all VRF extensions enabled on the IFC. Else, an error message is reported.

1. Go to the Links tab of the fabric.
2. Select the links with VRF Lite policy configured and click the delete button.



3. Click OK to confirm deletion.
4. Execute the Save and Deploy option in the fabric to reset the VRF Lite policy.

Deleting VRF Extensions deployed in External Fabric

This is a two part process:

1. Delete the sub interface created using interface TAB.



Note Skip this step if the VRF extension is to a non-Nexus device.

2. Delete the policy created for eBGP external connection.

Deleting the sub-interface

Navigate to the Control->Interfaces page as shown below, select the sub-interface(s) to be deleted and click the delete button.

Control / Fabrics / Interfaces

Interfaces

	Device Name	Name	Admin	Oper	Reason	Policy	Overlay N
<input type="checkbox"/>	n7k1-Edge1	mgmt0	↑	↑	ok	NA	NA
<input type="checkbox"/>	n7k1-Edge1	Vlan1	↓	↓	Administratively down	NA	NA
<input type="checkbox"/>	n7k1-Edge1	Loopback0	↑	↑	ok	NA	NA
<input type="checkbox"/>	n7k1-Edge1	Loopback1	↑	↓		NA	NA
<input type="checkbox"/>	n7k1-Edge1	Ethernet7/3	↓	↓	Administratively down	NA	NA
<input type="checkbox"/>	n7k1-Edge1	Ethernet7/5	↑	↓	Link not connected	int_trunk_host_11_1	NA
<input type="checkbox"/>	n7k1-Edge1	Ethernet7/6	↑	↑	ok	routed_host	NA
<input type="checkbox"/>	n7k1-Edge1	Ethernet7/9	↓	↓	Administratively down	NA	NA
<input type="checkbox"/>	n7k1-Edge1	Ethernet7/1/1	↓	↓	Administratively down	NA	NA
<input type="checkbox"/>	n7k1-Edge1	Ethernet7/1/2	↑	↓	Link not connected	NA	NA
<input type="checkbox"/>	n7k1-Edge1	Ethernet7/1/3	↓	↓	Administratively down	NA	NA
<input type="checkbox"/>	n7k1-Edge1	Ethernet7/1/4	↑	↑	ok	ext_int_routed_host_11_	NA
<input checked="" type="checkbox"/>	n7k1-Edge1	Ethernet7/1/4.2	↑	↑	ok	int_subif_11_1	NA

Deleting the eBGP policy

Navigate to fabric builder page and select the relevant external fabric (External65000 in this example). Select the device and using the second mouse button select view edit policy.

Select the row for the policy ID used in eBGP policy create. Click the “X” as shown below to delete the policy.

Issue a save and deploy in external fabric to deploy the policy change.

View/Edit Policies for n7k1-Edge1 (TBM14299900:Edge1)

Selected 1 / Total 30

+ / ✖ / View / View All / Deploy

Show Quick Filter

Template	Priority	Fabric Name	Serial Number	Editable	Entity Type	Entity Name	Source	Policy ID
<input checked="" type="checkbox"/> External_VRF_Lite_eBGP	500	ExternalE5000	TBM14299900:Edge1	true	SWITCH	SWITCH		POLICY-33350
<input type="checkbox"/> base_external_router	500	ExternalE5000	TBM14299900:Edge1	true	SWITCH	SWITCH		POLICY-33360
<input type="checkbox"/> breakout_interface	500	ExternalE5000	TBM14299900:Edge1	true	SWITCH	SWITCH		POLICY-33960
<input type="checkbox"/> routed_interface	350	ExternalE5000	TBM14299900:Edge1	false	INTERFACE	Ethernet7/1/4	LINK-UUID-4770	POLICY-32770
<input type="checkbox"/> routed_interface	350	ExternalE5000	TBM14299900:Edge1	false	INTERFACE	Ethernet7/4/1	LINK-UUID-4810	POLICY-32870
<input type="checkbox"/> interface_vrf	350	ExternalE5000	TBM14299900:Edge1	false	INTERFACE	Ethernet7/4/1.2	Ethernet7/4/1.2	POLICY-33370
<input type="checkbox"/> interface_vrf	350	ExternalE5000	TBM14299900:Edge1	false	INTERFACE	Ethernet7/1/4.2	Ethernet7/1/4.2	POLICY-33410
<input type="checkbox"/> routed_host	350	ExternalE5000	TBM14299900:Edge1	false	INTERFACE	Ethernet7/6		POLICY-33900
<input type="checkbox"/> trunk_interface	350	ExternalE5000	TBM14299900:Edge1	false	INTERFACE	Ethernet7/5	Ethernet7/5	POLICY-34170
<input type="checkbox"/> interface_mtu	352	ExternalE5000	TBM14299900:Edge1	false	INTERFACE	Ethernet7/1/4	LINK-UUID-4770	POLICY-32780
<input type="checkbox"/> no shut interface	362	ExternalE5000	TBM14299900:Edge1	false	INTERFACE	Ethernet7/1/4	LINK-UUID-4770	POLICY-32800

Deleting IFCs Created By Automatic VRF Lite creation

Editing and deleting IFCs are done through the Link tab in the VXLAN fabric. The extra consideration for auto configured IFCs is that, in order to prevent the regeneration of IFC on next save and deploy, the mode should be changed back to manual mode, or Save config should be done only on the relevant devices.

- In a consecutive scenario, if you delete the VRF lite IFC on one of the fabrics, the VRF lite is deleted from the peer fabric as well.
- When you want to delete a VRF lite between an easy fabric and an external fabric, delete the extension in the easy fabric using the top-down approach. The extension will be automatically deleted from the external fabric.
- Deploy the configurations in both the fabrics.

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

N9K-3-BGW Configurations

N9K-3-BGW (base border configurations) generated by template ext_base_border_vrflite_11_1



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

```
(config) #
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
route-map extcon-rmap-filter-allow-host deny 10
    match ip address prefix-list default-route
route-map extcon-rmap-filter-allow-host 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
route-map extcon-rmap-filter-v6-allow-host deny 10
    match ipv6 address prefix-list default-route-v6
route-map extcon-rmap-filter-v6-allow-host permit 1000
```

N9K-3-BGW VRF extension configuration

```
(config) #
configure profile MyVRF_50000
    vlan 2000
        vn-segment 50000
    interface vlan2000
        vrf member myvrf_50000
            ip forward
            ipv6 forward
            no ip redirects
            no ipv6 redirects
            mtu 9216
            no shutdown

(config) #

vrf context myvrf_50000
    vni 50000
    rd auto
    address-family ipv4 unicast
        route-target both auto
        route-target both auto evpn

    ip route 0.0.0.0/0 2.2.2.1
    address-family ipv6 unicast
        route-target both auto
        route-target both auto evpn

router bgp 7200
    vrf myvrf_50000
        address-family ipv4 unicast
            advertise l2vpn evpn
            redistribute direct route-map fabric-rmap-redis-subnet
            maximum-paths ibgp 2
            network 0.0.0.0/0
        address-family ipv6 unicast
            advertise l2vpn evpn
            redistribute direct route-map fabric-rmap-redis-subnet
            maximum-paths ibgp 2
```

```
neighbor 2.2.2.1 remote-as 65000
  address-family ipv4 unicast
  send-community both
  route-map extcon-rmap-filter out

(config) #

interface ethernet1/48.2
  encapsulation dot1q 2
  vrf member myvrf_50000
  ip address 2.2.2.2/24
  no shutdown
interface nve1
  member vni 50000 associate-vrf
configure terminal
  apply profile MyVRF_50000
```

