



DCI/WAN Connectivity for LAN Fabric Deployments, Release 11.3(1)

First Published: 2020-05-29

Americas Headquarters

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2019 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1

Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - MPLS SR and LDP Handoff 1

Overview of VXLAN EVPN to SR-MPLS and MPLS LDP Interconnection 1

VXLAN MPLS Topology 3

Configuration Tasks for VXLAN MPLS Handoff 5

Editing Fabric Settings for MPLS Handoff 5

 Editing Easy Fabric Settings 5

 Editing External Fabric Settings 7

Creating an Underlay Inter-Fabric Connection 8

Creating an Overlay Inter-Fabric Connection 12

Deploying VRFs 13

Changing the Routing Protocol and MPLS Settings 15

CHAPTER 2

Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - VRF Lite 17

Prerequisites and Guidelines 17

Sample Scenarios 20

VRF Lite Through the DCNM GUI – From a BGW Device to a Nexus 7000 Series Edge Router 21

VRF Lite Through the DCNM GUI – From a BGW Device To a Non-Nexus Device 33

Automatic VRF Lite (IFC) Configuration 40

Deleting VRF Lite IFCs 44

Additional References 46

Appendix 46

 N9K-3-BGW Configurations 46

CHAPTER 3

Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - Multi-Site 49

Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - Multi-Site 49

Prerequisites 50

Limitations	51
Save & Deploy Operation in the MSD Fabric	51
EVPN Multi-Site Configuration	53
Configuring Multi-Site Underlay IFCs - DCNM GUI	54
Configuring Multi-Site Underlay IFCs - Autoconfiguration	55
Configuring Multi-Site Underlay IFCs Towards a Non-Nexus Device - DCNM GUI	56
Configuring Multi-Site Overlay IFCs	57
Configuring Multi-Site Overlay IFCs - Autoconfiguration	59
Configuring Multi-Site Overlay IFCs Towards a Non-Nexus Device - DCNM GUI	60
Overlay and Underlay Peering Configurations on the Route Server N7k1-RS1	61
Viewing, Editing and Deleting Multi-Site Overlays	61
Deleting Multi-Site IFCs	62
Creating and Deploying Networks and VRFs in the MSD Fabric	63
Deploying a Legacy Site BGW (vPC-BGWs)	66
Additional References	70
Appendix	70
Multi-Site Fabric Base Configurations – Box Topology	70
IBGP Configuration for the Box Topology in the Easy7200 Fabric	71
Route Server Configuration	72
Multi-Site Overlay IFC Configuration	73
Multi-Site Underlay IFC Configuration – Out-of-Box Profiles	73



CHAPTER 1

Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - MPLS SR and LDP Handoff

This chapter describes how to configure the MPLS handoff features.

- [Overview of VXLAN EVPN to SR-MPLS and MPLS LDP Interconnection, on page 1](#)
- [VXLAN MPLS Topology, on page 3](#)
- [Configuration Tasks for VXLAN MPLS Handoff , on page 5](#)
- [Editing Fabric Settings for MPLS Handoff, on page 5](#)
- [Creating an Underlay Inter-Fabric Connection , on page 8](#)
- [Creating an Overlay Inter-Fabric Connection, on page 12](#)
- [Deploying VRFs, on page 13](#)
- [Changing the Routing Protocol and MPLS Settings, on page 15](#)

Overview of VXLAN EVPN to SR-MPLS and MPLS LDP Interconnection

From Cisco DCNM Release 11.3(1), the following handoff features are supported:

- VXLAN to SR-MPLS
- VXLAN to MPLS LDP

These features are provided on the border devices, that is, border leaf, border spine, and border super spine in the VXLAN fabric using the **Easy_Fabric_11_1** template. Note that the devices should be running Cisco NX-OS Release 9.3(1) or later. These DCI handoff approaches are the one box DCI solution where no extra Provider Edge (PE) device is needed in the external fabric.



Note If the switch is running a Cisco NX-OS Release 7.0(3)I7(X), enabling the MPLS handoff feature causes the switch to remove the NVE related config-profile CLIs when the switch is reloaded.

In the DCNM DCI MPLS handoff feature, the underlay routing protocol to connect a border device to an external fabric is ISIS or OSPF, and the overlay protocol is eBGP. The N-S traffic between the VXLAN fabric and external fabric running SR-MPLS or MPLS LDP is supported. Though, you can use DCNM for connecting two Data Center VXLAN fabrics via SR-MPLS or MPLS LDP.

Supported Platforms and Configurations

The following table provides information about the supported platforms:

Feature	Supported Platforms
VXLAN to SR-MPLS	Cisco Nexus 9300-FX2, Jericho+ based Nexus 9000 R-Series, and Nexus 3600 R-Series switches
VXLAN to MPLS LDP	Jericho+ based Cisco Nexus 9000 R-series and Cisco Nexus 3600 R-series switches

The following features aren't supported as they aren't supported on a switch:

- Coexisting of MPLS LDP and SR-MPLS interconnections
- vPC

The VXLAN to SR-MPLS handoff feature comprises the following configurations:

- Base SR-MPLS feature configuration.
- Underlay configuration between the DCI handoff device and the device in the external fabric for the underlay connectivity. DCNM supports ISIS or OSPF as the routing protocol for the underlay connectivity.
- Overlay configuration between a DCI handoff device and a core or edge router in the external fabric, or another border device in another fabric. The connectivity is established through eBGP.
- VRF profile

The VXLAN to MPLS LDP handoff feature comprises the following configurations:

- Base MPLS LDP feature configuration.
- Underlay configuration between the DCI handoff device and the device in the external fabric for the underlay connectivity. DCNM supports ISIS or OSPF as the routing protocol for the underlay connectivity.
- Overlay configuration between a DCI handoff device and a core or edge router in the external fabric, or another border device in another fabric. The connectivity is established through eBGP.
- VRF profile

Inter-Fabric Connections for MPLS Handoff

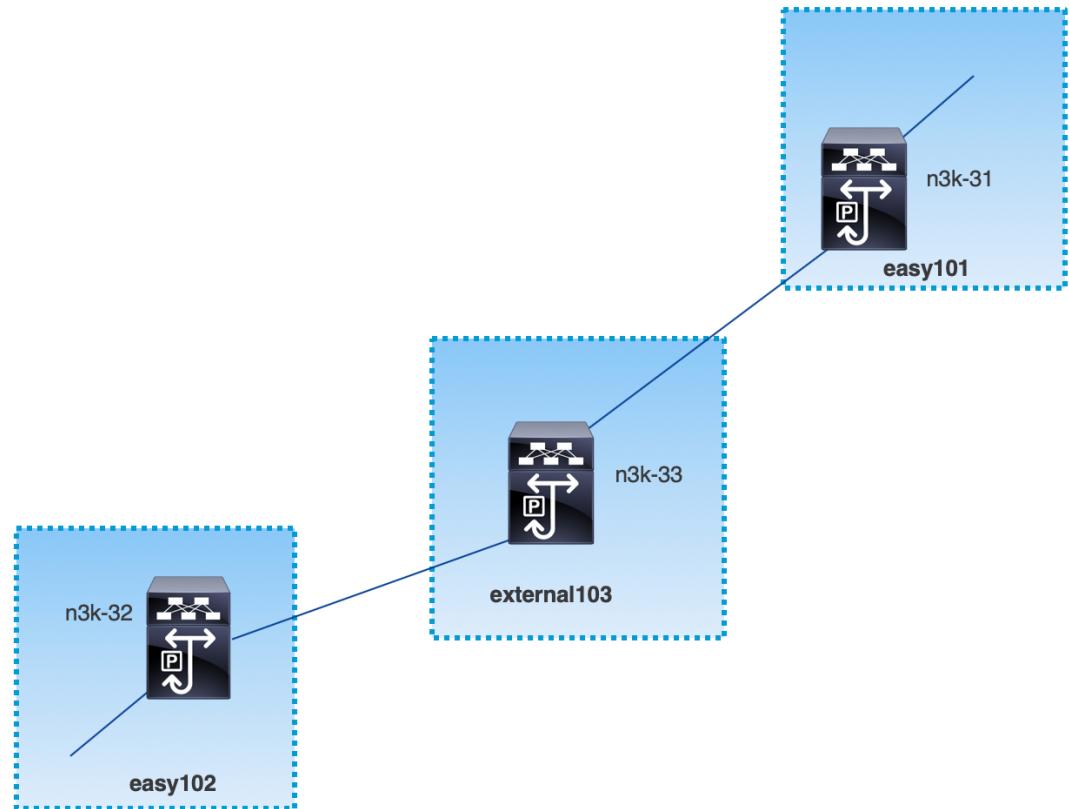
The following two inter-fabric connection links are introduced:

- **VXLAN_MPLS_UNDERLAY** for underlay configuration: This link corresponds to each physical link or Layer 3 port channel between the border and the external device (or a P router in MPLS or SR-MPLS). A border device can have multiple inter-fabric connection links as there could be multiple links connected to one or more external devices.
- **VXLAN_MPLS_OVERLAY** for eBGP overlay configuration: This link corresponds to the virtual link between a DCI handoff device and a core or edge router in the external fabric, or another border device in another fabric. This inter-fabric connection link can only be created on border devices which meet the image and platform requirement. A border device can have multiple of this type of IFC link as it could communicate to multiple core or edge routers.

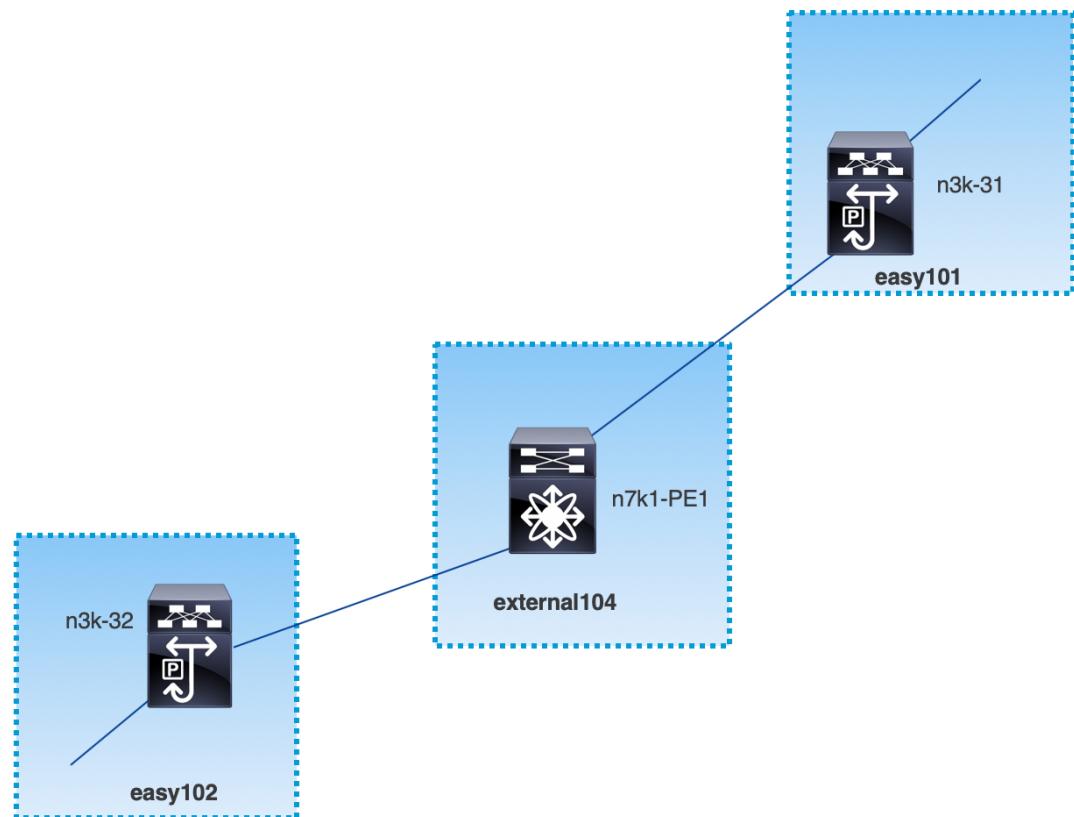
These inter-fabric connections can be manually created by using the DCNM Web UI or REST API. Note that the automatic creation of these inter-fabric connections isn't supported.

VXLAN MPLS Topology

MPLS-SR Topology



MPLS-LDP Topology



This topology shows only the border devices in the Easy Fabric and the core or edge router in the external fabric.

- The fabrics that are using the **Easy_Fabric_11_1** template are:
 - **easy101**
 - **easy102**
- The fabrics that are using the **External_Fabric_11_1** template are:
 - **external103**
 - **external104**
- The external fabric **external103** is running the MPLS SR protocol.
- The external fabric **external104** is running the MPLS LDP protocol.
- **n3k-31** and **n3k-32** are border devices performing VXLAN to MPLS handoff.
- **n7k-PE1** only supports MPLS LDP.
- **n3k-33** supports SR-MPLS.

Configuration Tasks for VXLAN MPLS Handoff

The following tasks are involved in configuring the MPLS handoff features:

1. Editing the fabric settings to enable MPLS handoff.
2. Creating an underlay inter-fabric connection link between the fabrics.
Specify whether you're using MPLS SR or LDP in the inter-fabric connection link settings.
3. Creating an overlay inter-fabric connection link between the fabrics.
4. Deploying a VRF for VXLAN to MPLS interconnection.

Editing Fabric Settings for MPLS Handoff

This section shows how to edit the fabric settings for the easy fabric and the external fabric to enable the MPLS handoff feature.

Editing Easy Fabric Settings

Procedure

- Step 1** Navigate to **Control > Fabric Builder**.
- Step 2** Click the **Edit Fabric** icon to edit the fabric settings.
- Step 3** Click the **Advanced** tab.

Editing Easy Fabric Settings

* Fabric Name : easy101

* Fabric Template : Easy_Fabric_11_1

General Replication vPC Protocols Advanced Resources Manageability Bootstrap Configuration Backup

Enable MPLS Handoff ?

* Underlay MPLS Loopback Id 101 Used for VXLAN to MPLS SR/LDP Handoff (Min:0, Max:1023)

Enable Default Queuing Policies ?

N9K Cloud Scale Platform Queuing Policy Queuing Policy for all 92xx, -EX, -FX, -FX2 series switches in the fabric

N9K R-Series Platform Queuing Policy Queuing Policy for all R-Series switches in the fabric

Other N9K Platform Queuing Policy Queuing Policy for all other switches in the fabric

Leaf Freeform Config

Note ! All configs should strictly match 'show run' output with respect to case and no spaces. Any mismatches will yield unexpected diffs during deployment.

Note ! All configs should strictly match 'show run' output with respect to case and no spaces. Any mismatches will yield unexpected diffs during deployment.

Save Cancel

Enable MPLS Handoff: Select the check box to enable the MPLS Handoff feature.

Note: For the brownfield import, select the **Enable MPLS Handoff** feature. Most of the IFC configuration will be captured in **switch_freeform**.

Underlay MPLS Loopback Id: Specifies the underlay MPLS loopback ID. The default value is 101.

Step 4 Click the **Resources** tab.

* Fabric Name : easy101

* Fabric Template : Easy_Fabric_11_1

General Replication vPC Protocols Advanced Resources Manageability Bootstrap Configuration Backup

Manual Underlay IP Address Allocation Checking this will disable Dynamic Underlay IP Address Allocations

* Underlay Routing Loopback IP Range 10.2.0.0/22 Typically Loopback0 IP Address Range

* Underlay VTEP Loopback IP Range 10.3.0.0/22 Typically Loopback1 IP Address Range

* Underlay RP Loopback IP Range 10.254.254.0/24 Anycast or Phantom RP IP Address Range

* Underlay Subnet IP Range 10.4.0.0/16 Address range to assign Numbered and Peer Link SVI IPs

* Underlay MPLS Loopback IP Range 10.101.0.0/25 Used for VXLAN to MPLS SR/LDP Handoff

Underlay Routing Loopback IPv6 Range Typically Loopback0 IPv6 Address Range

Underlay VTEP Loopback IPv6 Range Typically Loopback1 and Anycast Loopback IPv6 Address Range

Underlay Subnet IPv6 Range IPv6 Address range to assign Numbered and Peer Link SVI IPs

BGP Router ID Range for IPv6 Underlay

* Layer 2 VXLAN VNI Range 30000-49000 Overlay Network Identifier Range (Min:1, Max:16777214)

* Layer 3 VXLAN VNI Range 50000-59000 Overlay VRF Identifier Range (Min:1, Max:16777214)

* Network VLAN Range 2300-2999 Per Switch Overlay Network VLAN Range (Min:2, Max:3967)

* VRF VLAN Range 2000-2299 Per Switch Overlay VRF VLAN Range (Min:2, Max:3967)

* Subinterface Dot1q Range 2-511 Per Border Dot1q Range For VRF Lite Connectivity (Min:2, Max:4093)

Save Cancel

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.

Step 5 Click **Save & Deploy** to configure the MPLS feature on each border device in the fabric.

For more information about remaining fields, see [Creating a New VXLAN BGP EVPN Fabric](#).

Editing External Fabric Settings

Procedure

Step 1 Navigate to **Control > Fabric Builder**.

Step 2 Click the **Edit Fabric** icon to edit the fabric settings.

Step 3 (Optional) Under the **General** tab, uncheck the **Fabric Monitor Mode** check box.

Step 4 Click the **Advanced** tab.

Creating an Underlay Inter-Fabric Connection

The screenshot shows the 'General' tab selected in the top navigation bar. Key configuration fields include:

- Fabric Name:** external103
- Fabric Template:** External_Fabric_11_1
- vPC Peer Link VLAN:** 3600
- Power Supply Mode:** ps-redundant
- Enable MPLS Handoff:**
- Underlay MPLS Loopback Id:** 101
- Enable AAA IP Authorization:**
- Enable DCNM as Trap Host:**

Enable MPLS Handoff: Select the check box to enable the MPLS Handoff feature.

Underlay MPLS Loopback Id: Specifies the underlay MPLS loopback ID. The default value is 101.

Step 5

Click the **Resources** tab.

The screenshot shows the 'Resources' tab selected in the top navigation bar. Key configuration fields include:

- Subinterface Dot1q Range:** 2-511
- Underlay Routing Loopback IP Range:** 10.1.0.0/22
- Underlay MPLS Loopback IP Range:** 10.102.0.0/25

Underlay MPLS Loopback IP Range: Specifies the underlay MPLS SR or LDP loopback IP address range.

Note that the IP range should be unique, that is, it should not overlap with IP ranges of the other fabrics.

Step 6

Click **Save & Deploy** to configure the MPLS feature on each edge or core router in the fabric.

For more information about remaining fields, see [Creating an External Fabric](#).

Creating an Underlay Inter-Fabric Connection

This procedure shows how to create an underlay inter-fabric connection link.

Procedure

Step 1

Choose **Control > Fabric Builder**.

Choose a VXLAN fabric from which you want to create an underlay inter-fabric connection to MPLS.

Click **Tabular view** in the **Actions** panel that is displayed at the upper left of the window.

Click the **Links** tab.

Step 5 Check the existing links that are already discovered for the fabric.

In this example, the link from **easy101** to **external103** is already discovered.

Step 6 Select the existing discovered link and click the **Update Link** icon.

	Name	Policy	Info	Admin State	Oper State
1	easy101	n3k-31-Ethernet1/3---n9k-1-spine-Ethernet2/1	Neighbor Present	Up:-	Up:-
2	easy101	n3k-31-Ethernet1/2---n9k-17-Ethernet2/5	Neighbor Present	Up:-	Up:-
3	easy101<->external...	n3k-31-Ethernet1/5---n3k-33-Ethernet1/5	ext_vxlan_mpls_underlay_setup	Link Present	Up:Up
4	easy101<->external...	n3k-31-Ethernet1/1---n7k1-PE1-Ethernet10/1	Link Present	Up:Up	Up:Up

If a link isn't discovered, click the **Add Link** icon and provide all the details for adding an inter-fabric link.

Step 7 In the **Link Management - Edit Link** window, **Link Type** should be **Inter-Fabric**. Choose **VXLAN_MPLS_UNDERLAY** from the **Link Sub-Type** drop-down list and choose **ext_vxlan_mpls_underlay_setup** from the **Link Template** drop-down list.

Step 8 Under **Link Profile**, provide all the required information for the **General** tab.

MPLS-SR Configuration Example for Inter Fabric Link

Creating an Underlay Inter-Fabric Connection

Link Management - Edit Link

Link Management - Edit Link

* Link Type	Inter-Fabric
* Link Sub-Type	VXLAN_MPLS_UNDERLAY
* Link Template	ext_vxlan_mpls_underlay_setu
* Source Fabric	easy101
* Destination Fabric	external103
* Source Device	n3k-31
* Source Interface	Ethernet1/5
* Destination Device	n3k-33
* Destination Interface	Ethernet1/5

Link Profile

General

* IP Address/Mask	10.33.101.1/30	?	IP address with mask (e.g. 192.168.10.1/24)
* Neighbor IP	10.33.101.2	?	Neighbor IP address
* MPLS Fabric	SR	?	MPLS LDP or Segment-Routing
* Source SR Index	2001	?	Unique SR SID index for the source border
* Destination SR Index	2003	?	Unique SR SID index for the destination border
* SR Global Block Range	16000-23999	?	For Segment Routing binding (Min:16000, Max:471804)
* DCI Routing Protocol	ospf	?	Routing protocol used on the DCI MPLS link
* OSPF Area ID	1.1.1.1	?	OSPF Area ID in decimal format or IP address format
* DCI Routing Tag	MPLS_UNDERLAY	?	Routing Process Tag (Max Size 20)

Advanced

Save

MPLS-LDP Configuration Example for Inter Fabric Link

Link Management - Edit Link

Link Type: Inter-Fabric

Link Sub-Type: VXLAN_MPLS_UNDERLAY

Link Template: ext_vxlan_mpls_underlay_setu

Source Fabric: easy102

Destination Fabric: external104

Source Device: n3k-32

Source Interface: Ethernet1/1

Destination Device: n7k1-PE1

Destination Interface: Ethernet10/3

Link Profile

General (Selected)

Advanced

IP Address/Mask: 10.33.102.1/30

Neighbor IP: 10.33.102.2

MPLS Fabric: LDP

Source SR Index: (disabled)

Destination SR Index: (disabled)

SR Global Block Range: (disabled)

DCI Routing Protocol: is-is

OSPF Area ID: (disabled)

DCI Routing Tag: MPLS_UNDERLAY

Save

IP address/Mask: Specifies the IP address with mask for the source interface.

Neighbor IP: Specifies the IP address of the destination interface.

MPLS Fabric: Specifies whether the external fabric is running SR or LDP.

Note MPLS SR and LDP can't coexist on a single device.

Source SR Index: Specifies a unique SID index for the source border. This field is disabled if you choose LDP in the **MPLS Fabric** field.

Destination SR Index: Specifies a unique SID index for the destination border. This field is disabled if you choose LDP for the **MPLS Fabric** field.

SR Global Block Range: Specifies the SR global block range. You need to have the same global block range across the fabrics. The default range is from 16000 to 23999. This field is disabled if you choose LDP for the **MPLS Fabric** field.

DCI Routing Protocol: Specifies the routing protocol used on the DCI MPLS underlay link. You can choose either **is-is** or **ospf**.

OSPF Area ID: Specifies the OSPF area ID if you choose OSPF as the routing protocol.

DCI Routing Tag: Specifies the DCI routing tag used for the DCI routing protocol.

Step 9

Click Save.

Creating an Overlay Inter-Fabric Connection

- Step 10** Click **Save & Deploy** to deploy the updated configurations.
- Step 11** In the **Config Deployment** window, click **Deploy Config**.
- Step 12** Navigate to the destination fabric from the **Fabric Builder** window and perform a **Save & Deploy**, that is, perform steps 10 and 11.
-

Creating an Overlay Inter-Fabric Connection

This procedure shows how to create an overlay inter-fabric connection after the underlay inter-fabric connection is created. The overlay inter-fabric connection is the same for MPLS SR and LDP because the overlay connection uses eBGP.

Procedure

- Step 1** Click the **Add Link** icon.
- Step 2** In the **Link Management - Add Link** window, specify all the details.

Link Management - Add Link X

* Link Type	Inter-Fabric
* Link Sub-Type	VXLAN_MPLS_OVERLAY
* Link Template	ext_vxlan_mpls_overlay_setup
* Source Fabric	easy101
* Destination Fabric	easy102
* Source Device	n3k-31
* Source Interface	Loopback101
* Destination Device	n3k-32
* Destination Interface	Loopback101

▼ Link Profile

General	* BGP Local ASN <input type="text" value="101"/> <small>BGP Local Autonomous System Number</small> * BGP Neighbor IP <input type="text" value="10.102.0.1"/> <small>Neighbor IP address for eBGP peering</small> * BGP Neighbor ASN <input type="text" value="102"/> <small>BGP Neighbor Autonomous System Number</small>
----------------	---

Save

Link Type: Choose **Inter-Fabric**.

Link-Sub Type: Choose **VXLAN_MPLS_OVERLAY** from the drop-down list.

Link Template: Choose **ext_vxlan_mpls_overlay_setup** from the drop-down list.

Source Fabric - This field is prepopulated with the source fabric name.

Destination Fabric - Choose the destination fabric from this drop-down box.

Source Device and Source Interface - Choose the source device and the MPLS loopback interface. The IP address of the loopback interface will be used for overlay eBGP peering.

Destination Device and Destination Interface: Choose the destination device and a loopback interface that connects to the source device.

General tab in the **Link Profile** section.

BGP Local ASN: In this field, the AS number of the source device is autopopulated.

BGP Neighbor IP: Fill up this field with the IP address of the loopback interface at the destination device for eBGP peering.

BGP Neighbor ASN: In this field, the AS number of the destination device is autopopulated.

Step 3 Click **Save**.

Step 4 Click **Save & Deploy** to deploy the updated configurations.

Step 5 In the **Config Deployment** window, click **Deploy Config**.

Step 6 Navigate to the destination fabric from the **Fabric Builder** window and perform a **Save & Deploy**, that is, perform steps 4 and 5.

Note If there is only one MPLS overlay IFC link on the switch, you can remove it only when there's no VRF attached to either end of the MPLS overlay link.

Deploying VRFs

This procedure shows how to deploy VRFs for VXLAN to MPLS interconnection.



Note When you use the 4 byte ASN and auto route target is configured, the route target that is automatically generated is 23456:VNI. If two different VRFs in two different fabrics have the same VNI value, the route-target of the two VRFs would be the same due to auto route target and the value 23456 is always constant. For two fabrics connected via VXLAN MPLS handoff, this could result in unintended route exchange. Therefore, for security reasons, if you want to disable auto route target, you can disable it by customizing the network template and network extension template.

Procedure

Step 1 Navigate to **Control > Fabrics > VRFs**.

Step 2 In the **VRFs** window, click the **Add** icon to create a VRF. For more information, see [Creating VRFs for the Standalone Fabric](#).

Step 3 Select the newly added VRF and click **Continue**.

Deploying VRFs

VRF Name	VRF ID	Status
MyVRF_50000	50000	NA

Step 4 In the **VRF Deployment** window, you can see the topology of the fabric. Select a border device to attach a VRF to the border device where the MPLS LDP IFC link is created.

In this example, **n3k-31** is the border device in the **easy101** fabric.

Step 5 In the **VRF Extension Attachment** window, select the VRF and click the **Freeform config** button under the CLI Freeform column.

Switch	VLAN	Extend	CLI Freeform	Status	Loopback Id	Loopback IPv4 Address	Loopback IPv6 Address
n3k-31	2000	NONE	Freeform config	NA			

Step 6 Add the following freeform config manually to the VRF:

```
vrf context $$VRF_NAME$$
address-family ipv4 unicast
  route-target import $$REMOTE_PE_RT$$
address-family ipv6 unicast
  route-target import $$REMOTE_PE_RT$$
```

In the freeform config, **REMOTE_PE_RT** refers to the neighbor's BGP ASN and VNI number in the **ASN:VNI** format if the neighbor is a border device in Easy Fabric managed by DCNM.

 All configs should strictly match 'show run' output, with respect to case and newlines. Any mismatches will yield unexpected diffs during deploy.

```
vrf context MyVRF_50000
  address-family ipv4 unicast
    route-target import 103:50000
  address-family ipv6 unicast
    route-target import 103:50000|
```

Save Config

Step 7 Click **Save Config**.

Step 8 (Optional) Enter the Loopback Id and Loopback IPv4 Address and IPv6 address for the border device.

Step 9 Click **Save**.

Step 10 (Optional) Click the **Preview** icon in the **VRF Deployment** window to preview the configuration that will be deployed.

Step 11 Click **Deploy**.

Perform the same task from Step 3 to Step 11 in the destination fabric if the neighbor is a border device in Easy Fabric managed by DCNM.

Changing the Routing Protocol and MPLS Settings

This procedure shows how to change the routing protocol of a device from using IS-IS to OSPF, or from using MPLS SR to LDP for underlay IFC.



Note MPLS SR and LDP cannot co-exist on a device, and using both IS-IS and OSPF for MPLS handoff on the same device is not supported.

Procedure

Step 1 Remove all the MPLS underlay and overlay IFCs from the device that needs the change of DCI routing protocol or MPLS fabric.

Step 2 Click **Save & Deploy** for each fabric that is involved in the removal of the IFCs.

This step deletes all global MPLS SR/LDP configurations and the MPLS loopback interface that was previously created.

Step 3 Create a new IFC using the preferred DCI routing protocol and MPLS settings. For more information, see [Creating an Underlay Inter-Fabric Connection , on page 8](#).

■ **Changing the Routing Protocol and MPLS Settings**



CHAPTER 2

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 DCNM, 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 and Guidelines, on page 17](#)
- [Sample Scenarios, on page 20](#)
- [VRF Lite Through the DCNM GUI – From a BGW Device to a Nexus 7000 Series Edge Router , on page 21](#)
- [VRF Lite Through the DCNM GUI – From a BGW Device To a Non-Nexus Device , on page 33](#)
- [Automatic VRF Lite \(IFC\) Configuration, on page 40](#)
- [Deleting VRF Lite IFCs, on page 44](#)
- [Additional References, on page 46](#)
- [Appendix , on page 46](#)

Prerequisites and Guidelines

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.

Prerequisites and Guidelines

- 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 shows the 'VRF Information' and 'VRF Profile' sections of the Cisco DCNM interface. In the 'VRF Information' section, the 'VRF ID' is set to 50000, 'VRF Name' is CNC-Prod-Pod2, and 'VRF Template' is Default_VRF_Universal. In the 'VRF Profile' section, the 'Advanced' tab is selected, showing options for RP Loopback ID (0-1023), Underlay Mcast Address (IPv4 Multicast Address), Overlay Mcast Groups (224.0.0.0/4 to 239.255.255.255/4), and several checkboxes for IPv6 link-local, TRM BGW MSite, Advertise Host Routes, Advertise Default Route, and Config Static 0/0 Route.

Save **Cancel**

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.

Guideline

In a DCNM Release 10.4(2) setup where VRF-Lite IFCs are created, the required default prefix-lists or route-maps configs are added on the switch. When this DCNM Release 10.4(2) setup is upgraded to any of the DCNM 11.x releases, VRF-Lite related RPM configs might be saved as part of the `switch_freeform` policy.

The following route-map config is part of this `switch_freeform`:

```
route-map EXTCON-RMAP-FILTER-V6 deny 20
match ip address prefix-list host-route-v6
```

When this setup is upgraded from DCNM Release 11.x to 11.3(1), the route-map config is corrected with the following config:

```
route-map EXTCON-RMAP-FILTER-V6 deny 20
match ipv6 address prefix-list host-route-v6
```

Since RPM configs are saved in DCNM 11.x as `switch_freeform`, you need to manually delete the ip prefix-list match config in the `switch_freeform` policy so that `ipv6` match config is successful on the switch.

Sample Scenarios

Scenarios explained in this document:

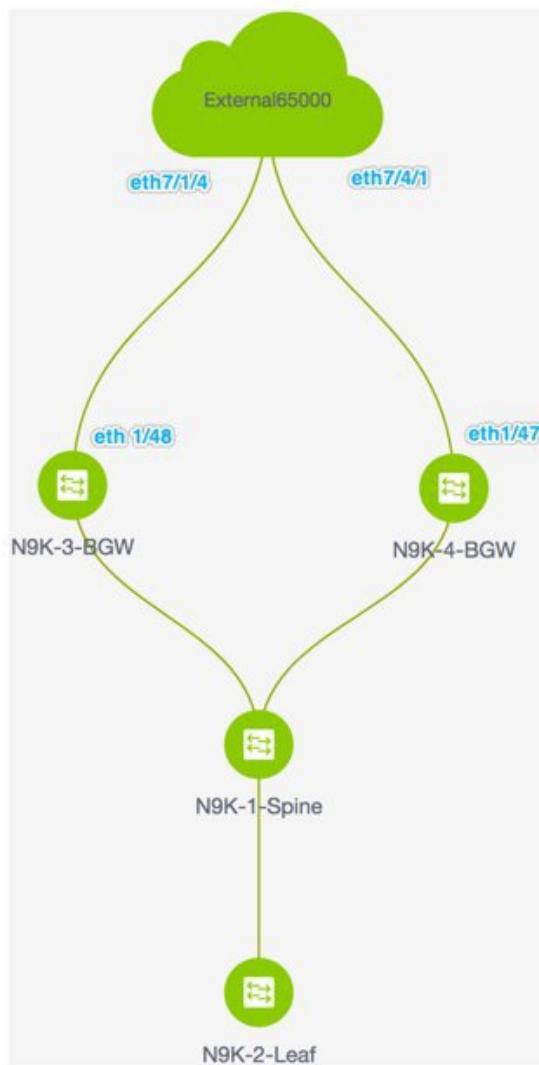
- VRF Lite through the DCNM GUI – From a BGW device to a Nexus 7000 Series edge router.
- VRF Lite through the DCNM GUI – From a BGW device to a non-Nexus device.
- Automatic VRF Lite (IFC) Configuration.



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. VRF Lite can be enabled on physical Ethernet interface or layer 3 port-channel. Sub-interface over physical interface or layer 3 port-channel interface is created by DCNM at the VRF extension moment for each VRF lite link the VRF is extended over.

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
Easy7200->External65000	N9K-3-BGW-Ethernet1/48--n7k1-Edge1-Ethernet7/14	ext_fabric_setup_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
Easy7200->External65000	N9K-4-BGW-Ethernet1/47--n7k1-Edge1-Ethernet7/4/1	ext_fabric_setup_11_1	Link Present	Up:Up	Up:Up
Easy7200->Easy60000	N9K-4-BGW-Ethernet1/2--N9K-15-BGW-Ethernet1/8	ext_fabric_setup_11_1	Link Present	Up:Up	Up:Up
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

* 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**).

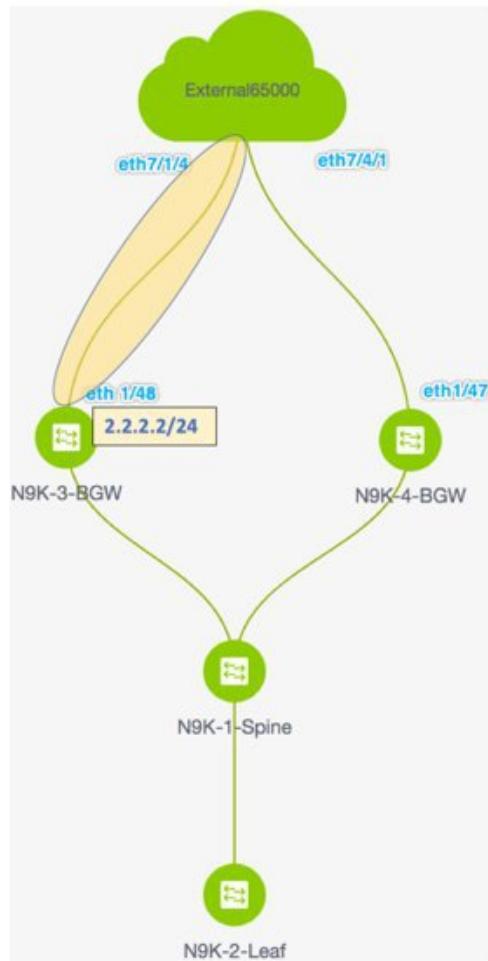
The **Advanced** tab has been added in the **Link Profile** section.

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.

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



- 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.

The screenshot shows the DCNM Links tab with the following table:

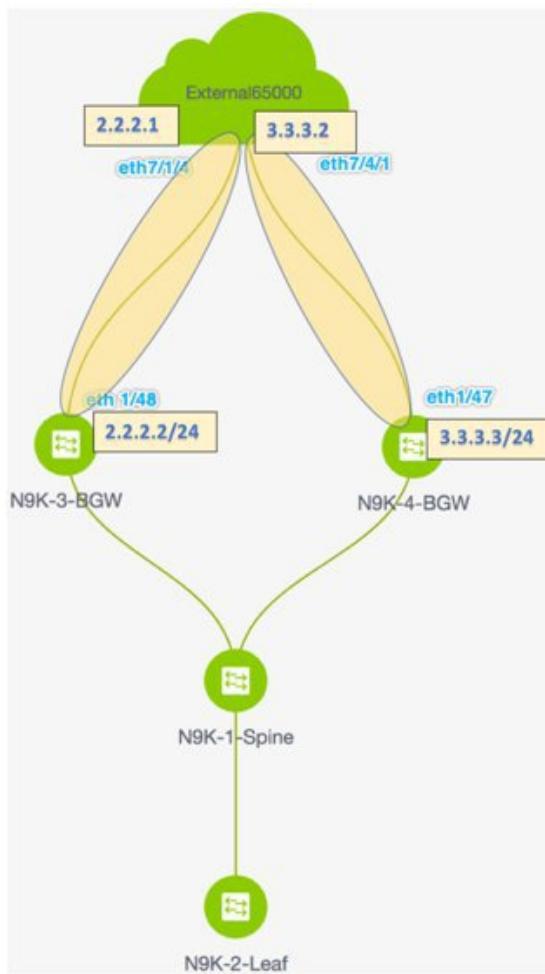
Scope	Name	Policy	Info	Admin State	Oper St
1	Easy7200->External65000	N9K-3-BGW-Ethernet1/48-n7k1-Edge1-Ethernet7/14	ext_fabric_setup_11_1	Link Present	Up:Up
2	Easy7200->External65000	N9K-4-BGW-Ethernet1/47-n7k1-Edge1-Ethernet7/4/1	ext_fabric_setup_11_1	Link Present	Up:Up
3	Easy7200	N9K-3-BGW-Ethernet1/47-N9K-1-Spine-Ethernet1/43	int_intra_fabric_num_link_11_1	Link Present	Up:Up
4	Easy7200	N9K-4-BGW-Ethernet1/48-N9K-1-Spine-Ethernet1/42	int_intra_fabric_num_link_11_1	Link Present	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

- 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_vrf lite_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_vrfite_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**.

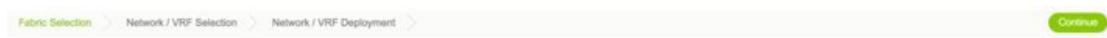
Step 2 - Deploy the individual VRF extensions on the BGWs N9K-3-BGW and N9K-4-BGW

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.

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

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
<input type="checkbox"/>	N9K-4-BGW	2000	NONE	Freeform config	NA

Save

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.

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 checked="" type="checkbox"/>	Switch	VLAN	Extend	CLI Freeform	Status			
<input checked="" type="checkbox"/>	N9K-3-BGW	2000	VRF_LITE	Freeform config	NA			
<input checked="" type="checkbox"/>	N9K-4-BGW	2000	VRF_LITE	Freeform config	NA			

Extension Details										
<input checked="" type="checkbox"/>	Source Switch	Type	IF_NAME	Dest. Switch	Dest. Interface	DOT1Q_ID	IP_MASK	NEIGHBOR_IP	NEIGHBOR ASN	IPV6_MASK
<input checked="" type="checkbox"/>	N9K-3-BGW	VRF_LITE	Ethernet1/48	Edge1	Ethernet7/1/4	2	2.2.2.2/24	2.2.2.1	65000	
<input checked="" type="checkbox"/>	N9K-4-BGW	VRF_LITE	Ethernet1/47	Edge1	Ethernet7/4/1	2	3.3.3.2/24	3.3.3.1	65000	

DCNM will create one sub-interface for each VRF-LITE link above using the values in DOT1Q_IP, IP_MASK and NEIGHBOR_IP fields. The IP_MASK and NEIGHBOR_IP fields for each VRF LITE extension are filled with the **IP Address/Mask** and **BGP Neighbor IP** values entered in VRF LITE link creation. The IP_MASK and NEIGHBOR_IP fields, along with the DOT1Q_ID field, can be overwritten.

IPV6_MASK and NEIGHBOR_IPV6 fields can be optionally entered if IPv6 eBGP session over the sub-interface is needed.

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.

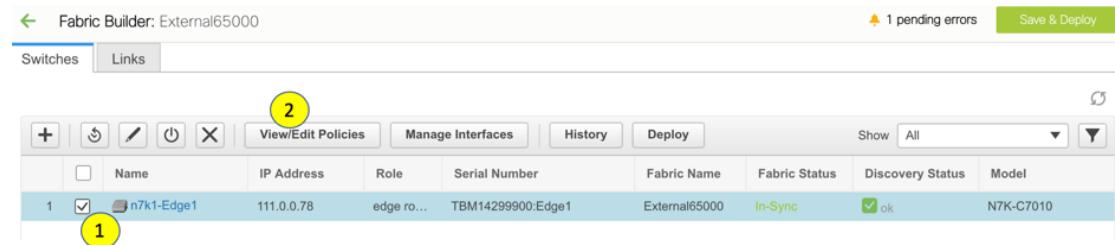
- **IP_MASK** - This will become the neighbor address at the edge router end and mask will be the local mask on the edge router. This is derived from the IFC prototype created in the earlier step.
- **Easy Fabric ASN** - This will become neighbor ASN from the edge router end. This is derived from the 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.

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



The View/Edit Policies screen comes up.

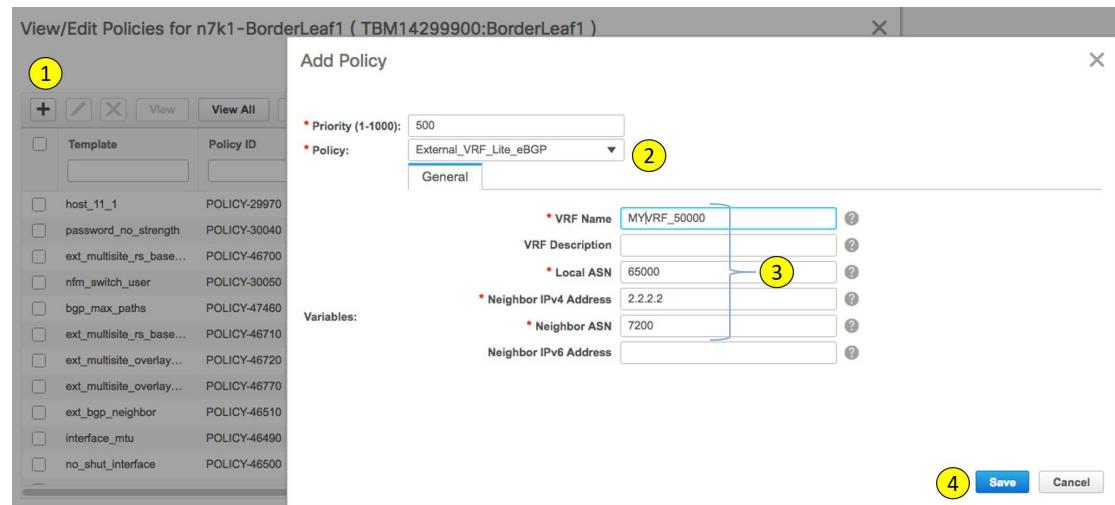
- Click + at the top left part of the screen to add a policy, and fill in the Add Policy screen as shown in the image.

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**.



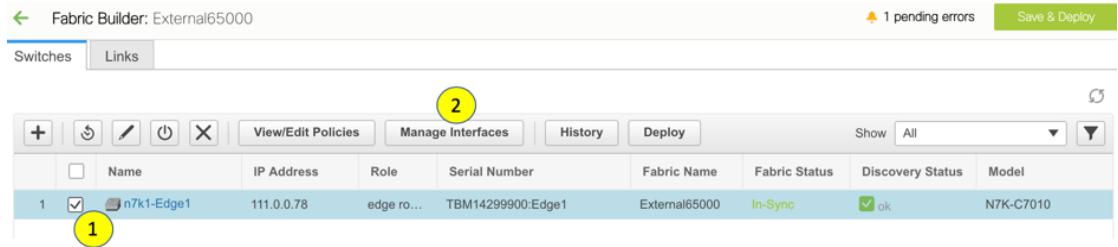
- As per the earlier steps, create a policy for the VRF extension towards **N9K-4-BGW**. The **Neighbor IPv4 Address** field for the second extension is updated with 3.3.3.3.

Sub interface policy on Edge Router

- In the **External65000** fabric topology screen, click **Tabular view**.

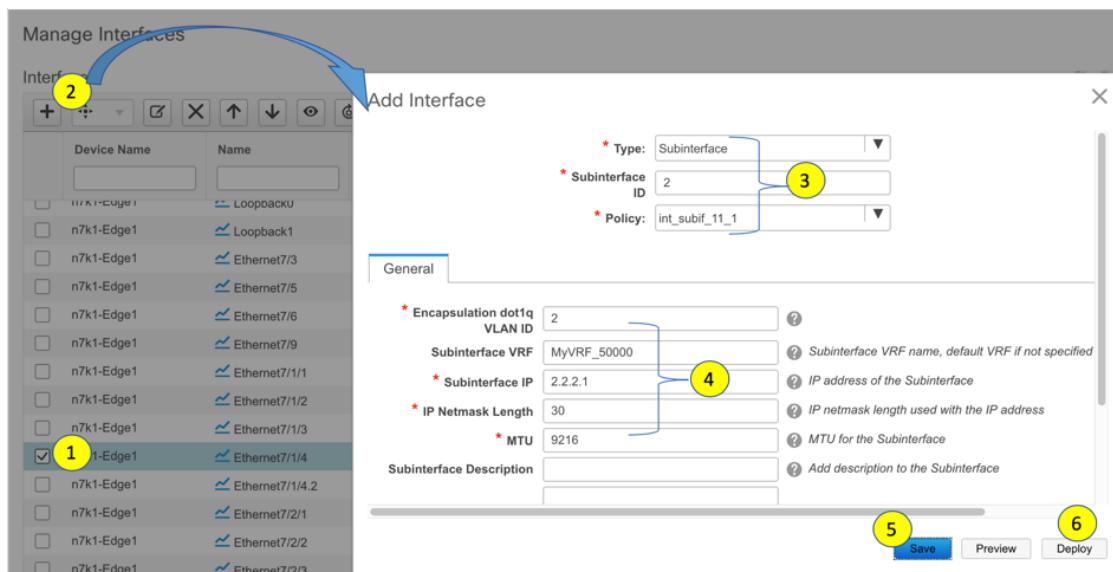
The **Switches | Links** screen comes up.

- 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 break out 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

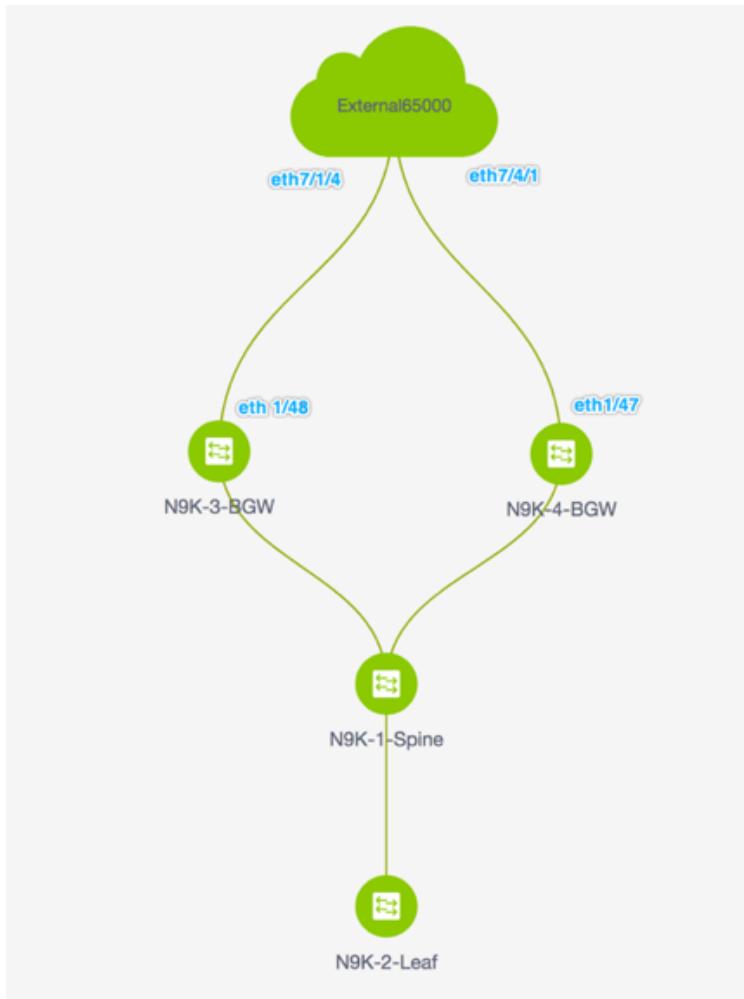
CDP or LLDP. To represent the non-Nexus device, you must create an external fabric. Refer the [Creating an External Fabric](#) topic to know how to create an external fabric. For this example, the external fabric **External65000** is created.

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.

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 shows the 'Link Management - Add Link' configuration screen. The 'Link Type' is set to 'Inter-Fabric'. Under 'Link Profile', the 'General' tab is selected, showing the following configuration:

- BGP Local ASN: 7200
- IP Address/Mask: 5.5.5.2/24
- BGP Neighbor IP: 5.5.5.1
- BGP Neighbor ASN: 65000

- 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

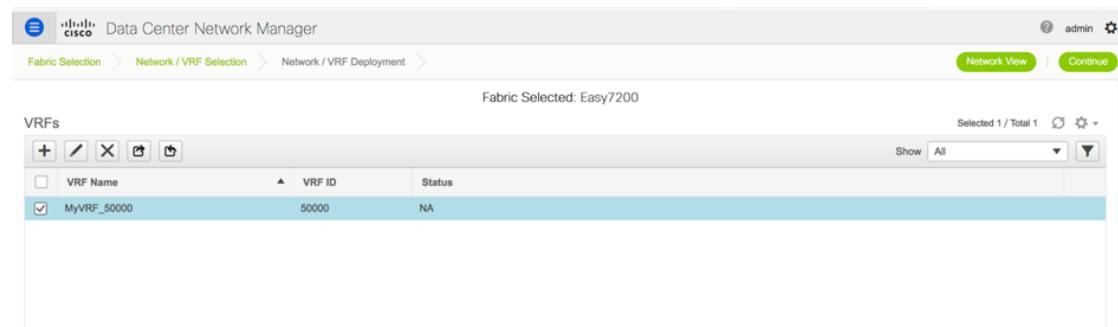
To extend VRFs beyond the fabric, the VRFs should have been created and deployed on relevant fabric devices, excepting the border devices.

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.



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

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

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
<input type="checkbox"/>	N9K-3-BGW		2000	NONE	Freeform config

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) X

Fabric Name: Easy7200

Deployment Options

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

MyVRF_50000		Switch	VLAN	Extend	CLI Freeform	▲	Loopback Id	Loopback IPv4 Address	Lo
<input type="checkbox"/>	N9K-3...	2000	VRF_LITE	<small>Freeform config</small>					

Extension Details

Sourc...	Type	IF_NAME	Dest. Switch	Dest. Interface	DOT1Q_ID
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.

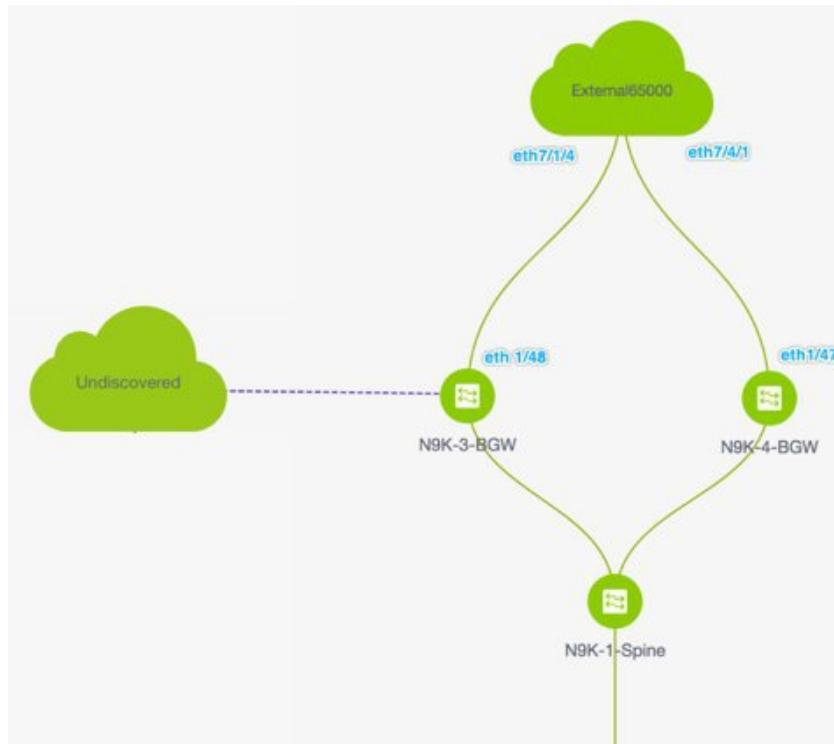
At the bottom right part of the screen, the color codes that represent different stages of deployment are displayed. The color of the switch icons changes accordingly (Blue for Pending state, yellow for In Progress state when the provisioning is in progress, red for failure state, green when successfully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

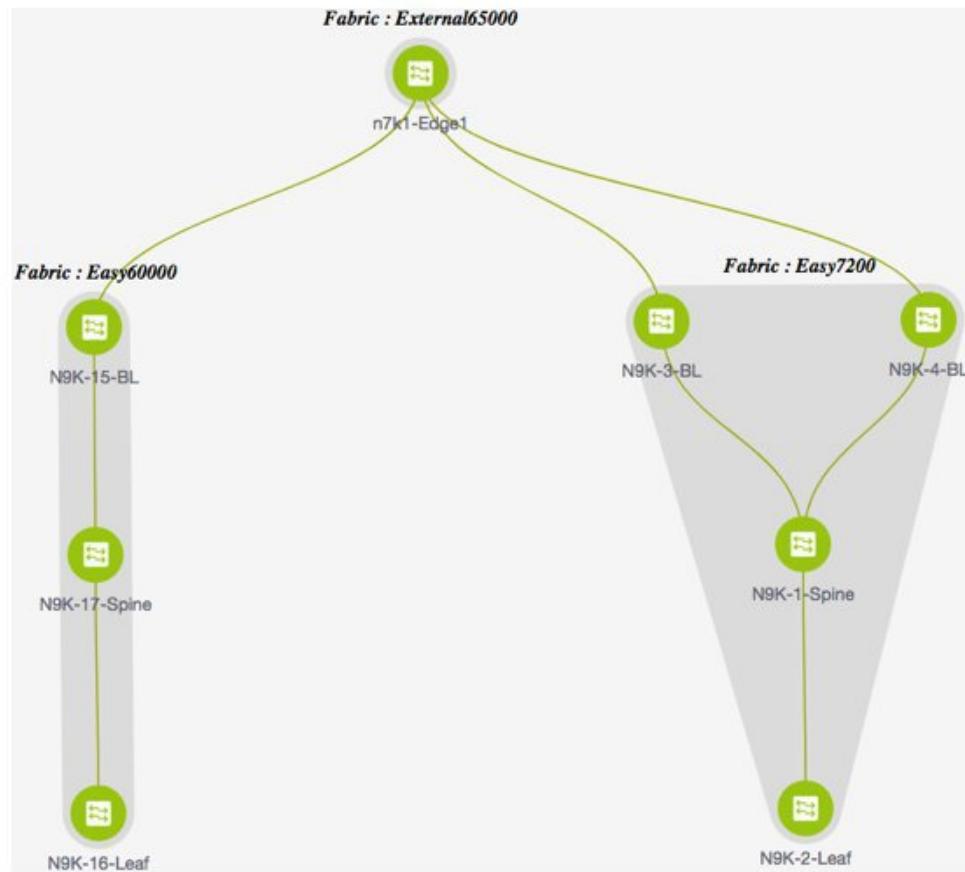


Automatic VRF Lite (IFC) Configuration

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	
<input checked="" type="radio"/> General <input type="radio"/> Replication <input type="radio"/> vPC <input type="radio"/> Advanced <input checked="" type="radio"/> Resources <input type="radio"/> Manageability <input type="radio"/> Bootstrap <input type="radio"/> Configuration Backup		
* Layer 2 VXLAN VNI Range	30000-49000	? Overlay Network Identifier Range (Min:1, Max:1677)
* Layer 3 VXLAN VNI Range	50000-59000	? Overlay VRF Identifier Range (Min:1, Max:1677)
* Network VLAN Range	2300-2999	? Per Switch Overlay Network VLAN Range (Min:1, Max:1677)
* VRF VLAN Range	2000-2299	? Per Switch Overlay VRF VLAN Range (Min:2, Max:1677)
* Subinterface Dot1q Range	2-511	? Per Border Dot1q Range For VRF Lite Connectivity
* VRF Lite Deployment	ToExternalOnly	? VRF Lite Inter-Fabric Connection Deploy Option: ToExternalOnly
* VRF Lite Subnet IP Range	10.33.0.0/16	? Address range to assign P2P DCI Links
* VRF Lite Subnet Mask	30	? Mask for Subnet Range

Similarly, update the settings for the Easy60000 fabric too.

- Check the **Auto Deploy Flag** check box in the **Link Management** dialog box. Checking this check box enables VRF lite deployment, including VRF Lite sub-interface and BGP peering configuration, on both ends of the link for managed devices.

Link Management - Edit Link

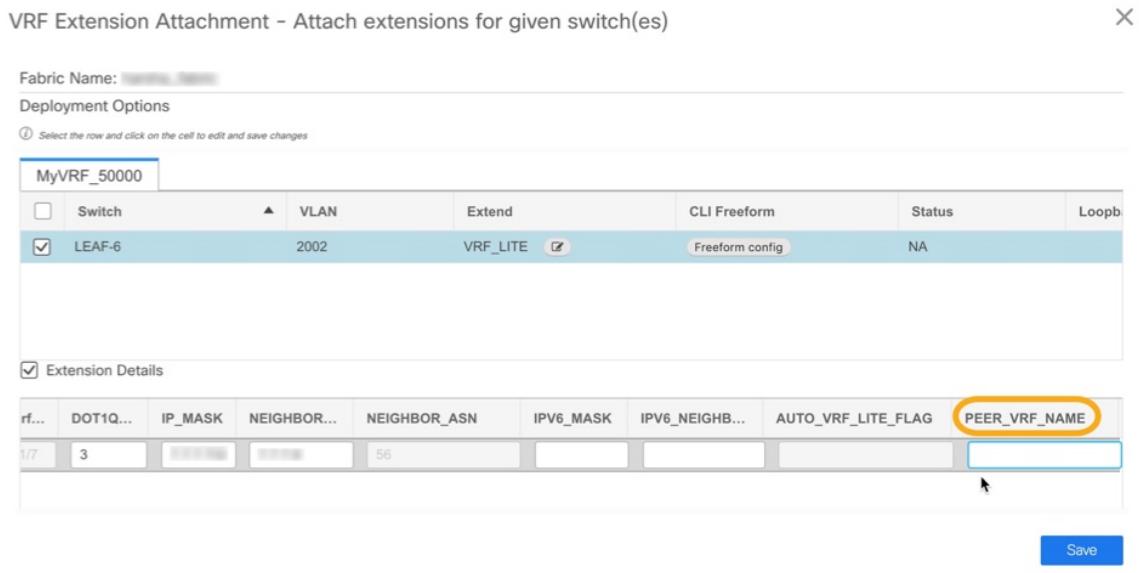
* Link Type	Inter-Fabric
* Link Sub-Type	VRF_LITE
* Link Template	ext_fabric_setup_11_1
* Source Fabric	Top_Down_ABC
* Destination Fabric	External
* Source Device	BL-2
* Source Interface	Port-channel901
* Destination Device	CORE-2
* Destination Interface	Port-channel901

Link Profile

<input checked="" type="radio"/> General <input type="radio"/> Advanced		
* BGP Local ASN	3000.3000	? Local BGP Autonomous System Number
* IP Address/Mask	10.33.0.1/30	? IP address for sub-interface in each VRF
* BGP Neighbor IP	10.33.0.2	? Neighbor IP address in each VRF
* BGP Neighbor ASN	5000.5000	? Neighbor BGP Autonomous System Number
Link MTU	9216	? Interface MTU on both ends of VRF Lite IFC
Auto Deploy Flag	<input checked="" type="checkbox"/>	? Flag that controls Auto VRF Lite Deployment on both ends for Managed devices

Deleting VRF Lite IFCs

- When you extend the VRF lite in a consecutive scenario, the VRF must be present in the peer fabric and the VRF name must be same. An error message appears, if the VRF is not present in the peer fabric and if you try to extend the VRF lite.
- When you extend the VRF lite between an easy fabric and an external fabric, the VRF name can be same as the source fabric, default, or another VRF name. Enter the VRF name used in the external fabric in the **PEER_VRF_NAME** field. The child PTIs for the subinterface, the VRF creation and BGP peering on the external fabric has the non-empty source. Hence, the policies cannot be edited or deleted from the **View/Edit policies** window.
- Enter the IP mask under the **IP_Mask** column as well to deploy VRF from topdown using symmetric VRF lite.
- Deploy configurations in both the fabrics. Perform **Save & Deploy** on the external fabrics to deploy the configurations. The easy fabric configuration can be deployed either from the topdown VRFs page or from the **Fabric Builder** window.



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.

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

Scope	Name	Policy	Info	Admin State	Oper State
1	Easy7200<->External...	ext_fabric_setup_11_1	Link Present	Up/Up	Up/Up
2	Easy7200<->External...	ext_fabric_setup_11_1	Link Present	Up/Up	Up/Up
3	Easy7200	int_intra_fabric_num_link_11_1	Link Present	Up/Up	Up/Up
4	Easy7200	int_intra_fabric_num_link_11_1	Link Present	Up/Up	Up/Up
5	Easy7200	int_intra_fabric_num_link_11_1	Link Present	Up/Up	Up/Up

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 the click the delete button.

Device Name	Name	Admin	Oper	Reason	Policy	Overlay N
n7k1-Edge1	mgmt0	↑	↑	ok	NA	NA
n7k1-Edge1	Vlan1	↓	↓	Administratively down	NA	NA
n7k1-Edge1	Loopback0	↑	↑	ok	NA	NA
n7k1-Edge1	Loopback1	↑	↓		NA	NA
n7k1-Edge1	Ethernet7/3	↓	↓	Administratively down	NA	NA
n7k1-Edge1	Ethernet7/5	↑	↓	Link not connected	int_trunk_host_11_1	NA
n7k1-Edge1	Ethernet7/6	↑	↑	ok	routed_host	NA
n7k1-Edge1	Ethernet7/9	↓	↓	Administratively down	NA	NA
n7k1-Edge1	Ethernet7/1/1	↓	↓	Administratively down	NA	NA
n7k1-Edge1	Ethernet7/1/2	↑	↓	Link not connected	NA	NA
n7k1-Edge1	Ethernet7/1/3	↓	↓	Administratively down	NA	NA
n7k1-Edge1	Ethernet7/1/4	↑	↑	ok	ext_int_routed_host_11_1	NA
<input checked="" type="checkbox"/> n7k1-Edge1	<input checked="" type="checkbox"/> Ethernet7/1/4.2	↑	↑	ok	int_subif_11_1	NA

Deleting the eBGP policy

Additional References

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)									
	+ / Edit	X	View	View All	Deploy	Show	Quick Filter	X	
	Template	Priority	Fabric Name	Serial Number	Editable	Entity Type	Entity Name	Source	Policy ID
<input checked="" type="checkbox"/>	External_VRF_Lite_eBGP	500	External65000	TBM14299900:Edge1	true	SWITCH	SWITCH		POLICY-33350
<input type="checkbox"/>	base_external_router	500	External65000	TBM14299900:Edge1	true	SWITCH	SWITCH		POLICY-33360
<input type="checkbox"/>	breakout_interface	500	External65000	TBM14299900:Edge1	true	SWITCH	SWITCH		POLICY-33960
<input type="checkbox"/>	routed_interface	350	External65000	TBM14299900:Edge1	false	INTERFACE	Ethernet7/1/4	LINK-UUID-4770	POLICY-32770
<input type="checkbox"/>	routed_interface	350	External65000	TBM14299900:Edge1	false	INTERFACE	Ethernet7/4/1	LINK-UUID-4810	POLICY-32870
<input type="checkbox"/>	interface_vrf	350	External65000	TBM14299900:Edge1	false	INTERFACE	Ethernet7/4/1.2	Ethernet7/4/1.2	POLICY-33370
<input type="checkbox"/>	interface_vrf	350	External65000	TBM14299900:Edge1	false	INTERFACE	Ethernet7/1/4.2	Ethernet7/1/4.2	POLICY-33410
<input type="checkbox"/>	routed_host	350	External65000	TBM14299900:Edge1	false	INTERFACE	Ethernet7/1/6		POLICY-33900
<input type="checkbox"/>	trunk_interface	350	External65000	TBM14299900:Edge1	false	INTERFACE	Ethernet7/5	Ethernet7/5	POLICY-34170
<input type="checkbox"/>	interface_mtu	352	External65000	TBM14299900:Edge1	false	INTERFACE	Ethernet7/1/4	LINK-UUID-4770	POLICY-32780
<input type="checkbox"/>	no_shut_interface	360	External65000	TBM14299900:Front1	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 12vpn evpn
```

N9K-3-BGW Configurations

```
 redistribute direct route-map fabric-rmap-redist-subnet
 maximum-paths ibgp 2
 network 0.0.0.0/0
 address-family ipv6 unicast
   advertise l2vpn evpn
   redistribute direct route-map fabric-rmap-redist-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
```



CHAPTER 3

Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - Multi-Site

This chapter explains LAN Fabric border provisioning using EVPN Multi-Site feature.

- [Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - Multi-Site , on page 49](#)
- [Prerequisites , on page 50](#)
- [Limitations, on page 51](#)
- [Save & Deploy Operation in the MSD Fabric , on page 51](#)
- [EVPN Multi-Site Configuration , on page 53](#)
- [Viewing, Editing and Deleting Multi-Site Overlays , on page 61](#)
- [Deleting Multi-Site IFCs, on page 62](#)
- [Creating and Deploying Networks and VRFs in the MSD Fabric , on page 63](#)
- [Deploying a Legacy Site BGW \(vPC-BGWs\), on page 66](#)
- [Additional References, on page 70](#)
- [Appendix , on page 70](#)

Border Provisioning Use Case in VXLAN BGP EVPN Fabrics - Multi-Site

This section explains how to connect two Virtual eXtensible Local Area Network (VXLAN) Border Gateway Protocol (BGP) Ethernet VPN (EVPN) fabrics through DCNM using the EVPN Multi-Site feature. The EVPN Multi-Site configurations are applied on the Border Gateways (BGWs) of the two fabrics. Also, you can connect two member fabrics of a Multi-Site Domain (MSD).

Introduced in DCNM 11.0(1) release, MSD is a multifabric container that is created to manage multiple member fabrics. It is a single point of control for definition of overlay networks and VRFs that are shared across member fabrics. See Multi-Site Domain for VXLAN BGP EVPN Fabrics section in the Control chapter for more information on MSD.

For a detailed explanation on the EVPN Multi-Site feature, see the [VXLAN BGP EVPN Multi-Site Design and Deployment](#) document.

Configuration methods - You can create underlay and overlay Inter-Fabric Connections (IFCs) between member fabrics through auto-configuration and through the DCNM GUI.

vPC configuration is supported for BGWs with the role **Border Gateway** from Cisco DCNM Release 11.1(1).

Prerequisites

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

Prerequisites

- The EVPN Multi-Site feature requires Cisco Nexus 9000 Series NX-OS Release 7.0(3)I7(1) or later.
- Familiarity with VXLAN BGP EVPN data center fabric architecture and configuration through DCNM.
- Familiarity with MSD fabrics, if you are connecting member fabrics of an MSD fabric.
- Fully configured VXLAN BGP EVPN fabrics that are ready to be connected using the EVPN Multi-Site feature, external fabric(s) configuration through DCNM, and relevant external fabric devices' configuration (for example, route servers).
 - VXLAN BGP EVPN fabrics (and their interconnection) can be configured manually or using DCNM. This document explains the process to connect the fabrics through DCNM. So, you should know how to configure and deploy a VXLAN BGP EVPN fabric, and how to create an external fabric through DCNM. For more details, see the VXLAN BGP EVPN Fabrics Provisioning section in the **Control** chapter.
- When you enable the EVPN Multi-Site feature on a BGW, ensure that there are no prior overlay deployments on it. Remove existing overlay profiles and then start provisioning Multi-Site extensions through DCNM.
- Execute the **Save & Deploy** operation in the member fabrics and external fabrics, and then in the MSD fabric.



Note The **Save & Deploy** button appears at the top right part of the fabric topology screen (accessible through the **Fabric Builder** window and clicking the fabric).

- Ensure that the role of the designated BGW is **Border Gateway** (or **Border Gateway Spine** for spine switches). To verify, right-click the BGW and click **Set role**. You can see that **(current)** is added to the current role of the switch.
- To ensure consistency across fabrics, ensure the following:



Note These checks are done for member fabrics of an MSD when the fabrics are moved under the MSD fabric.

- The underlay IP addresses across the fabrics, the loopback 0 address and the loopback 1 address subnets should be unique. Ensure that each fabric has a unique IP address pool to avoid duplicates.
- Each fabric should have a unique site ID and BGP AS number associated and configured.
- All fabrics should have the same Anycast Gateway MAC address.
- While the MSD provisions a global range of network and VRF values, some parameters are fabric-specific and some are switch-specific. You should specify fabric instance values for each

fabric (for example, multicast group subnet address) and switch instance values for each switch (for example, VLAN ID).


Note

Case 1 - During network creation, if a VLAN is specified, then for every switch, when you attach the network to the switch, automatically the VLAN will be autopopulated with the same VLAN that was specified during network creation. The network listing screen shows the VLAN a network level which applies for all the switch (has to be the same). The other thing to keep in mind is that even if one specified a VLAN during network creation, this can still be overwritten on a per switch basis.

Case 2 - During network creation, if a VLAN is not specified, then for every switch, when you attach the network to the switch, the next free VLAN from the per-switch VLAN pool is autopopulated. This means that on a per-switch basis, the VLAN may be different. The user can always overwrite the autopopulated VLAN and DCNM will honor it. For this case, it is possible that VNI 10000 may use VLAN 10 on leaf1 and VLAN 11 on leaf2. Hence, in the network listing, in this case, the VLAN will not be showcased.

DCNM always keeps track of VLANs on a per switch basis in its resource manager. This is true for either of the 2 cases mentioned above.

Limitations

- vPC configuration is not supported for the **Border Gateway Spine** role.
- The VXLAN OAM feature in Cisco DCNM is only supported on a single fabric or site.
- FEX is not supported on a Border Gateway or a Border Leaf with vPC or anycast.

Save & Deploy Operation in the MSD Fabric

These are some operations performed when you execute **Save & Deploy**:

- **Duplicate IP address check:** The MSD fabric checks if any BGW has a duplicate IP address. If so, an error message is displayed.



Change the BGP peering loopback IP address of the BGW(s).

After duplicate IP address issues are resolved, execute the **Save & Deploy** operation again in the MSD fabric.

Save & Deploy Operation in the MSD Fabric

- BGW base configuration:** When you execute Save and Deploy for the first time in the MSD fabric (assuming there are currently no IFCs or overlays to deploy), appropriate base configurations are deployed on the BGWs. They are given below:

Configuration	Description
<pre>evpn multisite border-gateway 7200 delay-restore time 300</pre>	<p>7200 is the site ID of the member fabric Easy7200. BGP ASN value is used to auto populate the site ID field. You can override this value. Even if you change the BGP ASN value, the site ID is still set to the first BGP ASN value.</p>
<pre>interface nve1 multisite border-gateway interface loopback100</pre>	<p>The loopback interface 100 is the configuration set in the MSD fabric settings. Once a loopback ID is chosen and Save and Deploy is executed, the loopback ID cannot be changed.</p> <p>To modify the role of the BGW in the MSD fabric, perform the following steps:</p> <ol style="list-style-type: none"> 1. In the easy fabric, modify the role of the BGW to leaf or border. 2. Save and deploy the changes. <p>This will remove the loopback 100 from the switch</p> <ol style="list-style-type: none"> 3. Change role back to BGW, and do a save and deploy. 4. In the MSD fabric, change the loopback ID setting to a desired value, and do a save and deploy.
<pre>interface ethernet1/47 evpn multisite fabric-tracking</pre>	<p>The evpn multisite fabric-tracking command is configured on all ports on a Border Gateway that have a connection to a switch with a Spine role.</p> <p>In case of a Border Gateway Spine role, all ports facing the leaf switch have this command configured</p>
<pre>interface loopback100 ip address 10.10.0.1/32 tag 54321 ip router ospf UNDERLAY area 0.0.0.0 ip pim sparse-mode no shutdown</pre>	<p>The Multi-Site loopback interface. This is configured on all Border Gateway (Spines).</p> <p>All BGWs in the same fabric get the same IP address. Each fabric gets its own unique IP address.</p> <p>It is not possible to change this address or ID without first changing role of the BGW.</p>

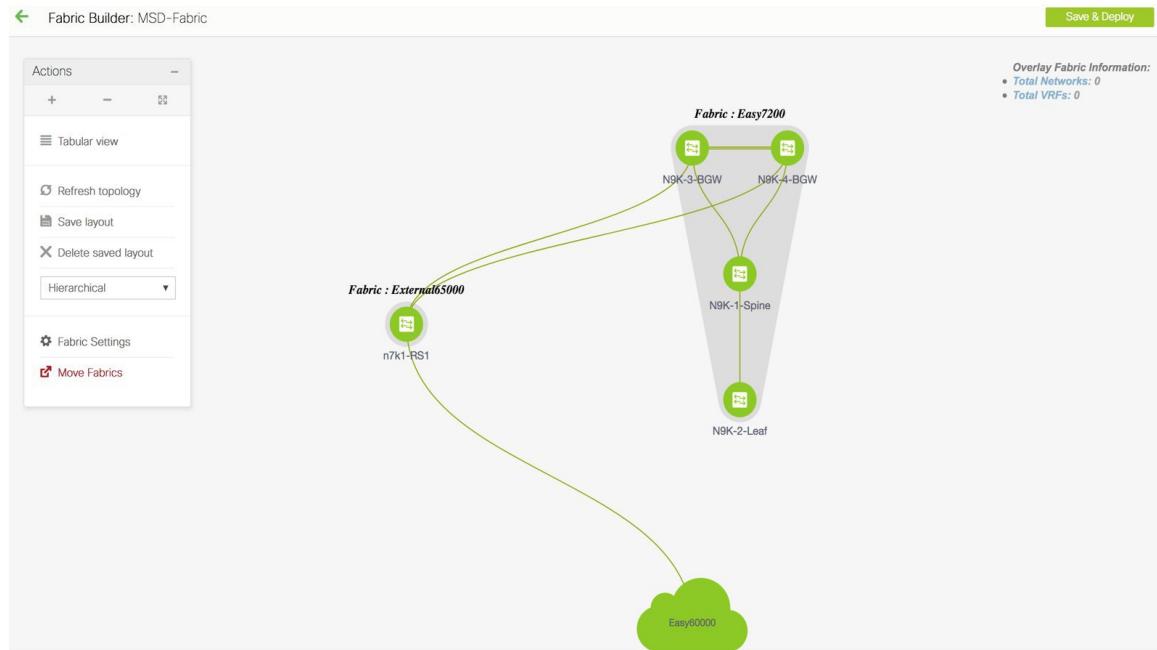
Configuration	Description
route-map rmap-redist-direct permit 10 match tag 54321	This is the configuration to redistribute the BGP peering loopback IP address (commonly loopback0), the VTEP primary (in case of vPC, the loopback secondary IP address), commonly loopback1, and the Multi-Site loopback IP address into the Multi-Site eBGP underlay sessions.

- When you execute the **Save & Deploy** operation in the MSD fabric, it works on all the BGW (or BGW Spine) devices in the member fabrics of an MSD.

After completing the EVPN Multi-Site specific prerequisites, start EVPN Multi-Site configuration. A sample scenario is explained.

EVPN Multi-Site Configuration

The EVPN Multi-Site feature is explained through an example scenario. Consider two VXLAN BGP EVPN fabrics, **Easy60000** and **Easy7200**, and an external fabric, **External65000**. The three fabrics are member fabrics of the MSD fabric **MSD-Fabric** and identified by a unique AS number. Easy60000 and Easy7200 are connected to a route server in External65000 (each fabric is). This document shows you how to enable end-to-end Layer 3 and Layer 2 traffic between hosts in Easy60000 and Easy7200, through the route server.



VXLAN BGP EVPN intra-fabric configurations, including network and VRF configurations are provisioned on the switches through DCNM software, 11.1(1) release. However, server traffic between the fabrics is only possible through the following configurations:

- A Data Center Interconnect (DCI) function like the Multi-Site feature is configured on the BGWs of both the fabrics (N9K-3-BGW and N9K-4-BGW in Easy7200, and the BGW in Easy60000). As part of

the configuration, since the BGWs of the fabrics are connected to the route server N7k1-RS1 in the external fabric External65000, appropriate eBGP peering configurations are enabled on the BGWs.

- As of now, overlay networks and VRFs are enabled on the non-BGW leaf and spine switches. For a fabric's traffic to go beyond the BGW, networks and VRFs should be deployed on all the BGWs too.

In a nutshell, the EVPN Multi-Site feature configuration comprises of setting up the BGW base configuration (enabled during the Save & Deploy operation), the eBGP underlay and overlay peering from the three BGWs to the route server N7k1-RS1. Both the underlay and overlay peering are established over eBGP through DCNM release 11.1(1).

You can create Multi-Site Inter-Fabric Connections (IFCs) between the fabrics through the DCNM GUI or through automatic configuration. First, underlay IFC creation is explained, followed by the overlay IFC creation.

Configuring Multi-Site Underlay IFCs - DCNM GUI

The end-to-end configurations can be split into these 2 high-level steps.

Step 1 - EVPN Multi-Site configurations on the BGWs in Easy7200

Step 2 - EVPN Multi-Site configurations on the BGW in Easy60000



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

Step 1 - EVPN Multi-Site configurations on the BGWs in Easy7200

For Multi-Site connectivity from Easy7200 to the external fabric, N9K-3-BGW and N9K-4-BGW are connected to the route server N7k1-RS1 in the external fabric. Follow these steps:

Deploying underlay IFCs between Easy7200 and External65000

- Deploying Underlay IFC from N9K-3-BGW to N7k1-RS1.
- Deploying Underlay IFC from N9K-4-BGW to N7k1-RS1.

Deploying Underlay IFC from N9K-3-BGW to N7k1-RS1

For the Multi-Site DCNM GUI configuration option, the **Deploy Border Gateway Method** field in the MSD fabric's settings (**DCI** tab) is set to **Manual**.

- Navigate to the **Links** tab and select the physical link connecting N9K-3-BGW to N7k1-RS1.
- Click the link edit icon as shown in the figure below to bring up the pop up.
- Select the MS underlay IFC sub type and fill in the required fields.



Note Enter the value as 1 in the **BGP Maximum Paths** field to allow DCNM to pick maximum path value. Enter a value between 2 and 64 to decide the maximum path value.

4. Save and deploy in the MSD will deploy the configuration on the N9K-3-BGW and N7k1-RS1.

Similar steps can be used to edit already created IFCs via the Links tab.

5. Similarly, create the underlay IFC from N9K-4-BGW to N7k1-RS1.

This completes Step 1 of the following.

Step 1 - EVPN Multi-Site configurations on the BGWs in Easy7200.

Step 2 - EVPN Multi-Site configurations on the BGW in Easy60000.

Next, configurations are enabled on the BGW in Easy60000.

Step 2 - EVPN Multi-Site configurations on the BGW in Easy60000

For Multi-Site connectivity between the Easy6000 fabric and the external fabric, EVPN Multi-Site configurations are enabled on the BGW interfaces in Easy60000 that are connected to the route server (N7k1-RS1) in the external fabric. Follow the steps as per the explanation for the connections between Easy7200 and External65000.

Configuring Multi-Site Underlay IFCs - Autoconfiguration

An underlay IFC is a physical link between the devices' interfaces.

- For underlay connectivity from Easy7200 to the external fabric, N9K-3-BGW and N9K-4-BGW are connected to the route server N7k1-RS1 in the external fabric.
- For underlay connectivity from Easy60000 to the external fabric, its BGW is connected to the route server N7k1-RS1.

Deploying Multi-Site Underlay IFCs Through Autoconfiguration

The underlay generated by DCNM is an eBGP session in the default IPv4 unicast routing table, in order to distribute the three loopback addresses needed for the Multi-Site control plane and data plane to function correctly.

For the Multi-Site autoconfiguration option, the underlay IFCs are automatically deployed by the MSD fabric.

The following rules apply to Multi-Site underlay IFC creation:

1. Check the **Multi-Site Underlay IFC Auto Deployment Flag** check box to enable the multi-site underlay autoconfiguration. Uncheck the check box to disable autoconfiguration. The check box is unchecked by default.
2. An IFC is deployed on every physical connection between the BGWs of different member fabrics that are physically connected.
3. An IFC is deployed on every physical connection between a BGW and a router with the role Core Router imported into an external fabric which is a member of the MSD fabric.
- If you do not want an IFC to be auto generated on a connection, then shut the link, execute the Save & Deploy operation, and delete the undesired IFCs. Also, ensure that there is no existing policy or pre-configured IP address on the interface. Else, use the Manual mode.
4. The IP addresses used to deploy the underlay are derived from the IP address range in the DCI Subnet IP Range field (DCI tab) of the MSD fabric.

Configuring Multi-Site Underlay IFCs Towards a Non-Nexus Device - DCNM GUI

Just like overlay IFCs, Multi-Site underlay IFCs can be viewed via the MSD, external and member fabrics. Also, the underlay IFCs can be edited and deleted via the VXLAN or MSD fabrics.

Configuring Multi-Site Underlay IFCs Towards a Non-Nexus Device - DCNM GUI

In this case, the non-Nexus device is not imported into DCNM, or discovered through Cisco Discovery Protocol or Link Layer Discovery Protocol (LLDP). For example, a Cisco ASR 9000 Series router or even a non-Cisco device.

The steps are similar to the **Configuring Multi-Site Underlay IFCs - DCNM GUI** task.

1. In the **Fabric Builder** window, choose the **Easy7200** fabric.

The **Easy7200** topology window appears.

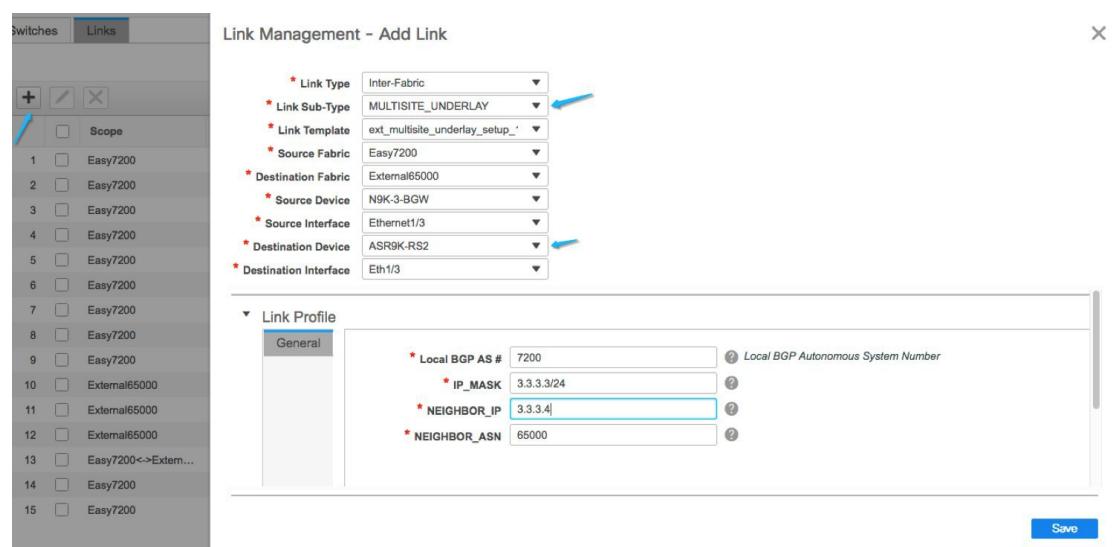
2. From the **Actions** panel at the left, click **Tabular view**.

The **Switches | Links** window appears.

3. Click the **Links** tab and click **+**.

The **Add Link** window appears.

4. Fill in the fields.



Link Type – Choose **Inter-Fabric**.

Link Sub-Type – Choose **MULTISITE_UNDERLAY**.

Link Template - By default, the **ext_multisite_underlay_setup_11_1** template is populated.

Source Fabric - **Easy7200** is selected by default since the IFC is created from **Easy7200** to the ASR device.

Destination Fabric – Select the external fabric. In this case, **External65000** is selected.

Source Device and **Source Interface** - Choose the border device and the interface that is connected to the ASR device.

Destination Device - Type any string that identifies the device. The destination device **ASR9K-RS2** does not appear in the drop-down list when you create an IFC for the first time. Once you create an IFC towards **ASR9K-RS2** and associate it with the external fabric **External65000**, **ASR9K-RS2** appears in the list of devices displayed in the **Destination Device** field.

Also, after the first IFC creation, **ASR9K-RS2** is displayed in the **External65000** external fabric topology, within Fabric Builder.

Destination Interface - Type any string that identifies the interface.

You have to manually enter the destination interface name each time.

General tab in the **Link Profile** section.

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

Source IP Address/Mask - Enter the IP address and mask that is used as the local interface for the Multi-Site underlay IFC.

Destination IP - Enter the IP address of the **ASR9K-RS2** interface used as the eBGP neighbor.

Destination BGP ASN - In this field, the AS number of the external fabric **External65000** is autopopulated since it is chosen as the external fabric.

5. Click **Save** at the bottom right of the window.

The **Switches|Links** window appears again. You can see that the IFC entry is updated.

6. Click **Save and Deploy** at the top right of the window.

The link on which the IFC is deployed has the relevant policy configured in the **Policy** column.

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

Configuring Multi-Site Overlay IFCs

An overlay IFC is a link between the devices' loopback0 interfaces.

Deploying Overlay IFCs in Easy7200 and Easy60000 comprises of these steps:

- Deploying Overlay IFC from N9K-3-BGW to N7k1-RS1.
- Deploying Overlay IFC from N9K-4-BGW to N7k1-RS1.
- Deploying the Overlay IFC from the BGW in Easy60000 to N7k1-RS1.

Deploying Overlay IFCs between Easy7200 and External65000

- Deploying Overlay IFC from N9K-3-BGW to N7k1-RS1.
- Deploying Overlay IFC from N9K-4-BGW to N7k1-RS1.

Deploying Overlay IFCs - from N9K-3-BGW to N7k1-RS1

1. Click Control > Fabric Builder. The Fabric Builder window appears.
2. Choose the MSD fabric, **MSD-Fabric**. The fabric topology comes up.

Configuring Multi-Site Overlay IFCs

3. Click Tabular view. The Switches | Links screen comes up.
4. Click the Links tab. It lists links within the MSD fabric. Each row either represents an intra-fabric link within Easy7200 or Easy60000, or a link between border devices of member fabrics, including External65000.

5. Click the Add Link icon at the top left part of the screen.

The Link Management – Add Link screen comes up.

Some fields are explained:

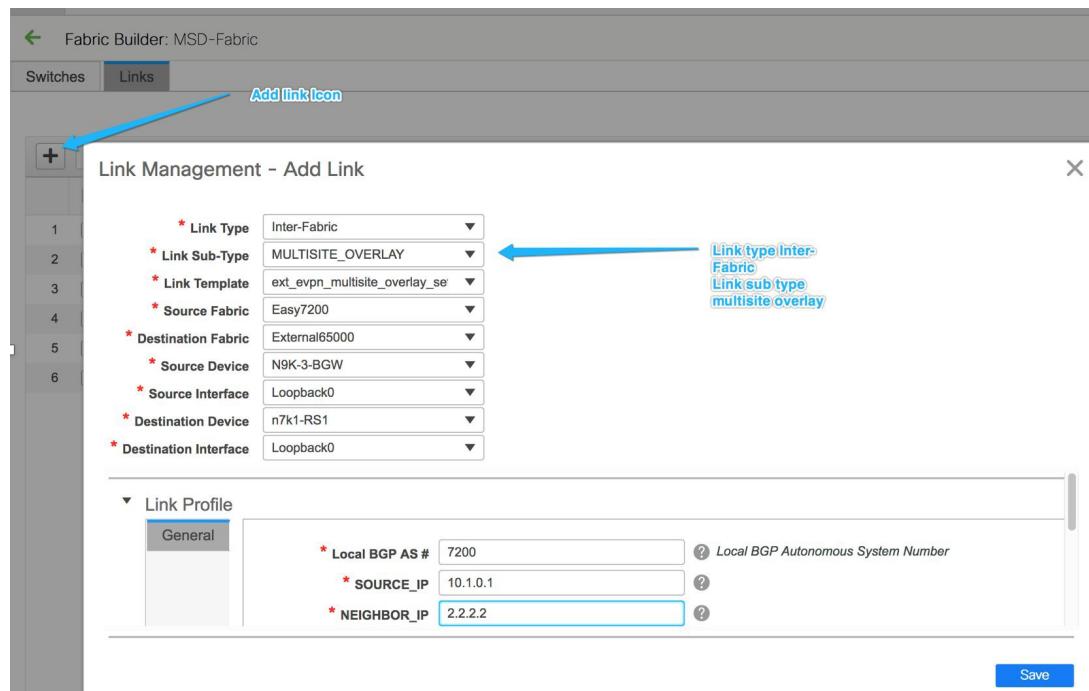
Link Type – Inter-Fabric is autopopulated.

Link Sub-Type – Choose MULTISITE_OVERLAY.

Link Template – The default template for creating an overlay is displayed.

You can edit the 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. The BGP AS numbers are derived based on fabric values.



6. 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.

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

8. Go to the Scope drop-down list at the top right of the window 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.

10. Similarly, create an overlay IFC from N9K-4-BGW to N7k1-RS1.

After the overlay IFCs from N9K-3-BGW and N9K-4-BGW to N7k1-RS1 are deployed, the fabric overlay traffic can flow between Easy7200 and External65000.

11. Similarly, deploy the overlay IFC from the BGW in the Easy60000 fabric to N7k1-RS1.

Configuring Multi-Site Overlay IFCs - Autoconfiguration

An overlay IFC is a link between the devices' loopback0 interfaces. For overlay connectivity from the Easy7200 and Easy60000 fabrics to the route server N7k1-RS1 in External65000, a link is enabled between the BGW devices and the N7k1-RS1's loopback0 interfaces.

Deploying Overlay IFCs in Easy7200 and Easy60000

- Deploying Overlay IFC from N9K-3-BGW to N7k1-RS1.
- Deploying Overlay IFC from N9K-4-BGW to N7k1-RS1.
- Deploying the Overlay IFC from the BGW in Easy60000 to N7k1-RS1.

Deploying Multi-Site Overlay IFCs Through Autoconfiguration

You can automatically configure the Multi-Site overlay through one of these options:

1. Route Server - The BGW forms an overlay to the route server. This option is explained in the example.
2. Direct to BGW: A full mesh of Multi-Site Overlay IFC from every BGW in a fabric to every BGW in other member fabrics.

To choose one of the above options, go to the MSD fabric's settings, select the DCI tab, and set the Deploy Border Gateway Method field to Route_Server (such as for this example) or Direct to BGW. By default, the Manual option is selected.

The IFCs needed for deployment of Networks and VRFs at the BGW nodes can be auto configured via the MSD fabric template. The settings to enable that are in MSD fabric template.

The default mode for the **Deploy Border Gateway Method** field is **Manual**, which implies that the IFCs have to be created via the link tab in MSD fabric. It must be changed to the **Route_Server** or **Direct to BGW** mode for autoconfiguration.

The IFCs created via auto configuration can only be edited or deleted via the link tab in MSD or member fabrics (except external fabric). As long as an IFC exists, or there is any user defined policy on the physical or logical link, auto configuration will not touch the IFC configuration.

You can see that **Route_Server** is selected in the Deploy Border Gateway Method field in the above image.

Route Server

Configuring Multi-Site Overlay IFCs Towards a Non-Nexus Device - DCNM GUI

This implies that all BGW devices in all member fabrics will create a Multi-Site overlay BGP connection to one or more route servers in one or more external fabrics which are members of the MSD fabric.

In this topology, there is one route server n7k1-RS1, and its BGP peering address (1.1.1.1) is shown in the route server list. This peering address can be configured out of band or with create interface tab in DCNM. N7k1-RS1 must be imported into the DCNM (in the external fabric, in this example) and the peering address configured before executing the Save & Deploy option.

You can edit the route server peering IP address list at any time, but you can delete a configured Multi-Site overlay only through the Links tab.

The BGP AS number of each route server should be specified in the MSD fabric settings. Note that the route server AS number can be different than the fabric AS number of the external fabric.

Configuring Multi-Site Overlay IFCs Towards a Non-Nexus Device - DCNM GUI

In this case, the non-Nexus device is not imported into DCNM, or discovered through Cisco Discovery Protocol or Link Layer Discovery Protocol (LLDP). For example, a Cisco ASR 9000 Series router or even a non-Cisco device.

The steps are similar to the **Configuring Multi-Site Overlay IFCs - DCNM GUI** task.

1. In the **Fabric Builder** window, choose the **Easy7200** fabric.

The **Easy7200** topology window appears.

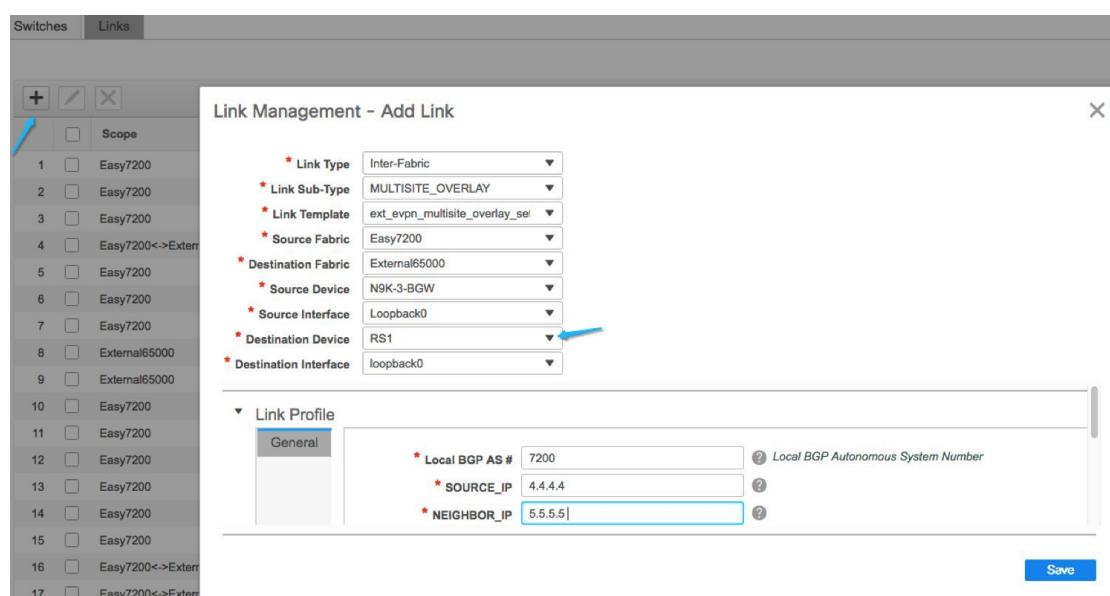
2. From the **Actions** panel, click **Tabular view**.

The **Switches | Links** window appears.

3. Click the **Links** tab and click **+**.

The **Add Link** screen comes up.

4. Fill in the fields.



Link Type – Choose **Inter-Fabric**.

Link Sub-Type – Choose **MULTISITE_OVERLAY**.

Link Template – By default, the **ext_evpn_multisite_overlay_setup** template is populated.

Source Fabric – **Easy7200** is selected by default since the IFC is created from **Easy7200** to the ASR device.

Destination Fabric – Select the external fabric. In this case, **External65000** is selected.

Source Device and Source Interface - Choose the border device and the loopback0 interface that is the source interface of the overlay.

Destination Device: Type any string that identifies the device. The destination device **ASR9K-RS1** does not appear in the drop-down list when you create an IFC for the first time. Once you create an IFC towards **ASR9K-RS1** and associate it with the external fabric **External65000**, **ASR9K-RS1** appears in the list of devices displayed in the **Destination Device** field.

Also, after the first IFC creation, **ASR9K-RS1** is displayed in the **External65000** topology screen, within Fabric Builder.

Destination Interface: Type any string that identifies the interface.

You have to manually enter the destination interface name each time.

General tab in the **Link Profile** section.

Source BGP ASN: In this field, the AS number of the source fabric **Easy7200** is autopopulated.

Source IP Address/Mask: Enter the IP address of the loopback0 interface for the Multi-Site overlay IFC.

Destination IP: Enter the IP address of the **ASR9K-RS1** loopback interface used for this Multi-Site overlay IFC.

Destination BGP ASN: In this field, the AS number of the external fabric **External65000** is autopopulated since it is chosen as the external fabric.

- 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 link on which the IFC is deployed has the relevant policy configured in the **Policy** column.

- Go to the **Scope** drop-down list at the top right of the window and choose **External65000**. The external fabric **Links** screen is displayed. You can see that the overlay IFC is displayed here.

Overlay and Underlay Peering Configurations on the Route Server N7k1-RS1

When you execute the Save and Deploy operation in the MSD fabric during the IFCs creation, peering configurations are enabled on the router server N7k1-RS1 towards the BGWs in the VXLAN fabrics.

Viewing, Editing and Deleting Multi-Site Overlays

The overlay IFCs can be viewed via the MSD and member fabrics Links tab as shown below.

Deleting Multi-Site IFCs

The IFCs can be edited and deleted in the member fabric or in the MSD fabric.

Multi-Site overlay IFCs can also be created by the links tab in MSD fabric.

Once the IFC is deleted, you should execute the Save & Deploy operation in the external and VXLAN fabric (or MSD fabric) to undeploy the IFC on the switches and remove the intent from DCNM.



Note Until a particular IFC is completely deleted from DCNM, auto configuration will not regenerate it on a Save & Deploy operation in the MSD fabric.

	Scope	Name	Policy	Info	Admin State	Oper State
1	Easy7200<->External65000	N9K-4-BGW-Loopback0--n7k1-RS1-Loopback0	ext_evpn_multisite_overlay_setup	Neighbor Missing	--	--
2	Easy7200<->External65000	N9K-3-BGW-Loopback0--n7k1-RS1-Loopback0	ext_evpn_multisite_overlay_setup	Neighbor Missing	--	--
3	Easy60000<->External65000	N9K-15-BGW-Ethernet1/3--n7k1-RS1-Ethernet7/10/1		Link Present	Up:Up	Up:Up
4	Easy7200	N9K-2-Leaf-Ethernet1/47--N9K-1-Spine-Ethernet1/47	int_intra_fabric_num_link_11_1	Link Present	Up:Up	Up:Up
5	Easy7200<->External65000	N9K-3-BGW-Ethernet1/48--n7k1-RS1-Ethernet7/1/4	ext_multisite_underlay_setup_11_1	Link Present	Up:Up	Up:Up
6	Easy7200	N9K-3-BGW-Ethernet1/47--N9K-1-Spine-Ethernet1/43	int_intra_fabric_num_link_11_1	Link Present	Up:Up	Up:Up
7	Easy7200<->External65000	N9K-4-BGW-Ethernet1/47--n7k1-RS1-Ethernet7/4/1	ext_multisite_underlay_setup_11_1	Link Present	Up:Up	Up:Up
8	Easy7200	N9K-4-BGW-Ethernet1/22--N9K-3-BGW-Ethernet1/22	int_intra_fabric_num_link_11_1	Link Present	Up:Up	Up:Up
9	Easy7200	N9K-4-BGW-Ethernet1/21--N9K-3-BGW-Ethernet1/21	int_intra_fabric_num_link_11_1	Link Present	Up:Up	Up:Up
10	Easy7200	N9K-4-BGW-Ethernet1/48--N9K-1-Spine-Ethernet1/42	int_intra_fabric_num_link_11_1	Link Present	Up:Up	Up:Up

Deleting Multi-Site IFCs

1. Navigate to the Links tab, select the IFCs to be deleted and click the Delete icon as shown below.
2. Perform a Save & Deploy in the MSD fabric to complete deletion.



Note If auto configuration of IFCs is enabled in the MSD fabric settings, then performing a Save & Deploy may regenerate the IFC intent.

If all or large number of IFCs are to be deleted, then temporarily change the BGW deploy mode to Manual setting before performing Save & Deploy.

- Deleting IFC in a non-Nexus Switch: If you delete the last IFC in a non-Nexus switch, the switch is removed from the topology. From Cisco DCNM Release 11.2(1), you can remove non-Nexus switches and neighbor switches like a physical switch from the **Tabular view** window or from the fabric topology window by right-clicking the switch and choosing **Discovery > Remove from fabric** in the drop-down menu.
- Removing a fabric from an MSD fabric: Before removing a fabric from an MSD fabric, remove all the multisite overlays from all BGWs in that fabric. Otherwise, you will not be able to remove the fabric. After the following save and deploy in the easy fabric, all the multisite configurations, such as IFC, multisite loopback address configured in MSD are removed from BGWs.

- Device role change: You can change the device role from Border to Border Gateway, but the role change from Border Gateway to Border is allowed only if there are no multisite IFCs or overlays deployed on the device.

Creating and Deploying Networks and VRFs in the MSD Fabric

Networks and VRFs can be created from the MSD context in the Networks and VRF page, these can be deployed on BGW nodes for all member fabrics of that MSD.

The following screenshots show how to select networks and deploy them. From the MSD fabric context, any device can be selected for network or VRF deployment. However, networks or VRFs can be deployed only on BGWs from the MSD context in the network deployment screen. The leaf deployment can be done from the fabric context or from the Fabric Builder context.

The image consists of two screenshots of the Cisco Data Center Network Manager interface, illustrating the steps for network deployment:

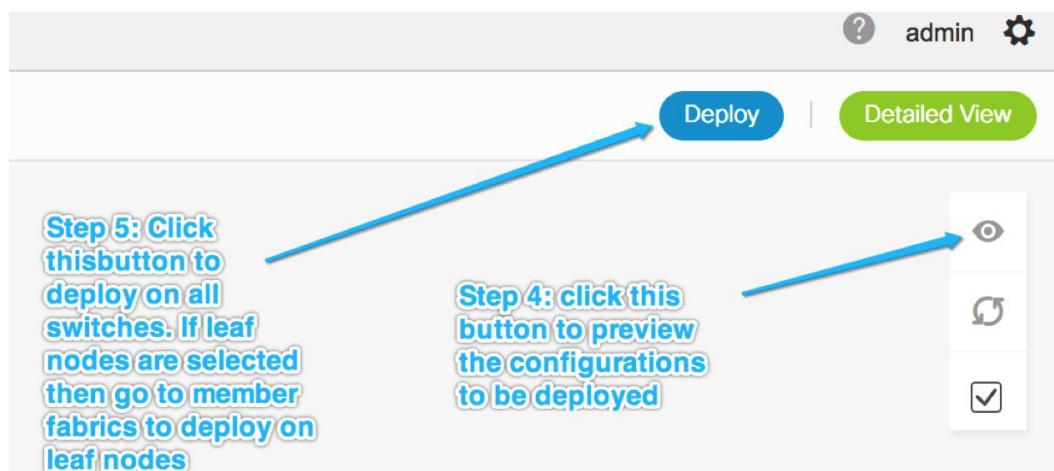
Screenshot 1: Network / VRF Selection

This screenshot shows the 'Network / VRF Selection' step. The 'Fabric Selected: MSD-Fabric' is chosen. A table lists networks, with one row selected: 'MyNetwork_30000'. A blue arrow points to this row with the text 'Step 2: select NW(s) to be deployed'. Another blue arrow points to the 'Continue' button at the top right with the text 'Step 3: press the continue button to go to the deployment page'.

Network Name	Network ID	VRF Name	IPv4 Gateway/Subnet	IPv6 Gateway/Prefix	Status	VLAN ID
MyNetwork_30000	30000	MyVRF_50000	11.0.0.1/24		NA	111

Screenshot 2: Network Extension Attachment - Attach extensions for given switch(es)

This screenshot shows the 'Network Extension Attachment' step. It lists switches under 'Deployment Options': 'N9K-3-BGW' and 'N9K-4-BGW', both checked. A blue arrow points to these checkboxes with the text 'Step 2: Check this boxes to select BGWs on which to deploy the NW(s)'. Another blue arrow points to the 'Save' button at the bottom right with the text 'Step 3: click this to move to deployment screen, repeat till all nodes on which NW(s) are to be deployed'.



Deploying Networks with a Layer 3 Gateway on a BGW

Perform the following steps:



Note Selecting an interface to deploy SVI is only available on vPC BGW setups. This is a device limitation not a DCNM limitation.

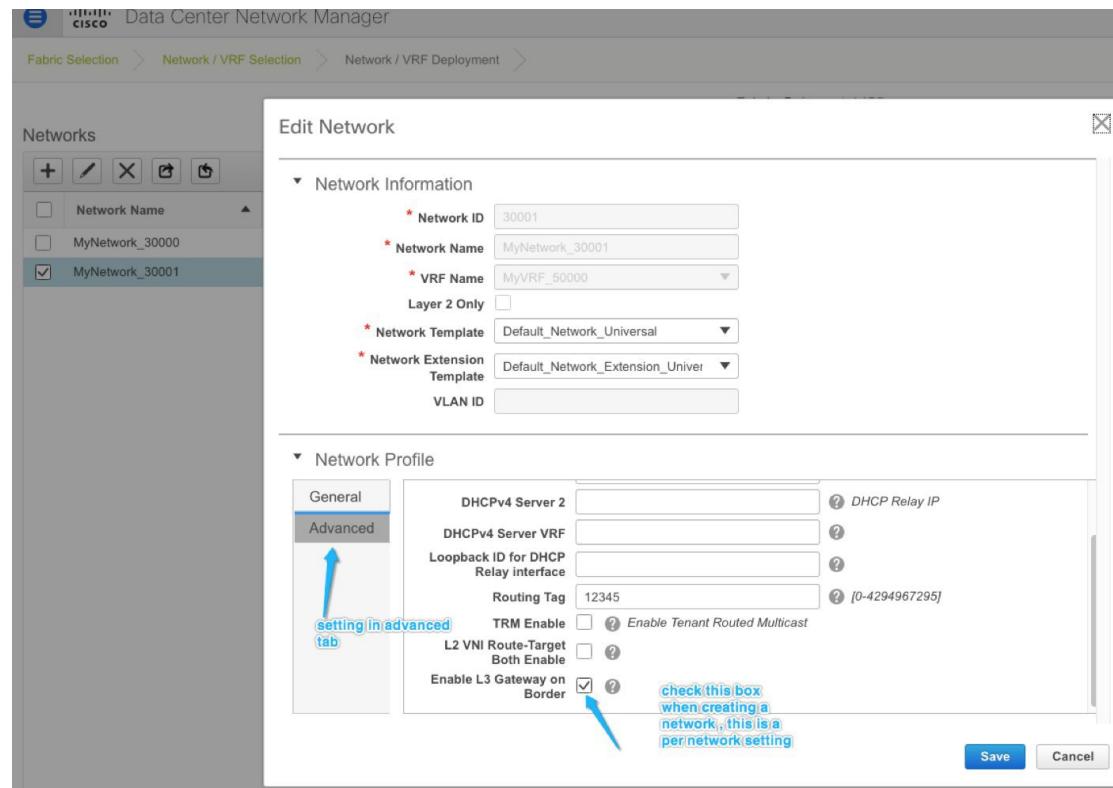
- In order to deploy a network with a Layer 3 gateway on a Border device (Border, Border spine, Border Gateway, Border Gateway spine), perform these steps.

When creating the network, check the **Enable L3 gateway on Border** check box, as shown in the figure below. Note that this is a network wide setting, so whenever this network is deployed on the Border device, the Layer 3 gateway will be deployed. If this is required on only a subset of the Borders, then a custom template is required.

When deploying the network on the Border device, select the interface(s) in the **Interface** column in case of vPC BGW.

Just like the leaf switch, the candidate ports should have **int_trunk_host_policy_11_1**, otherwise they will not be in the interface list.

The interface policy can be modified through the **Control > Interfaces** tab.



- When deploying the network on the vPC pair of BGWs, select the interface(s) in the Interfaces column. Only vPC port channel interfaces are the candidate interfaces.

Deploying a Legacy Site BGW (vPC-BGWs)

Network Extension Attachment - Attach extensions for given switch(es) X

Fabric Name: MSD

Deployment Options

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

MyNetwork_30001						
	Switch	VLAN	Extend	Interfaces	CLI Freeform	Status
<input checked="" type="checkbox"/>	BL-1	2300	MULTISITE	<input type="button"/> Port-channel500	Freeform config	DEPLOYED
<input checked="" type="checkbox"/>	BL-2	2300	MULTISITE	<input type="button"/> Port-channel500	Freeform config	DEPLOYED

Save

Interfaces X

	Interface/Ports	Port Type
<input checked="" type="checkbox"/>	Port-channel500	trunk

Save

Deploying a Legacy Site BGW (vPC-BGWs)

The recommended way of integrating non-VXLAN BGP EVPN (legacy) and VXLAN BGP EVPN fabrics is by using a pair of VPC BGWs. For more information about this method, see [NextGen DCI with VXLAN EVPN Multi-Site Using vPC Border Gateways White Paper](#).

The vPC BGW method replaces the Pseudo-Border Gateway method recommended in the DCNM release 11.1(1).

In this section, the tasks from the white paper that can be accomplished by DCNM are explained with an example topology.

Prerequisites

- Legacy network is already setup by a method. This is out of the scope for this document.
- Familiarity with fabric creation and Multi-Site use case.

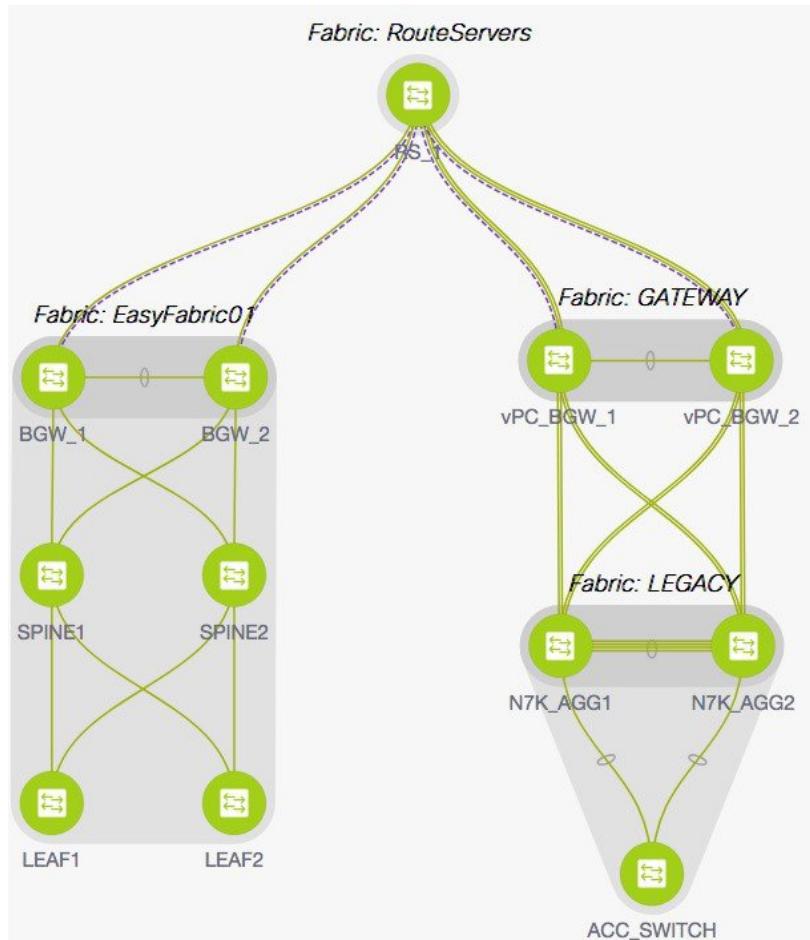
Tasks Overview

The following information is covered as part of this section:

1. Fabrics to be created using DCNM:
 - a. VXLAN fabric with vPC border gateways.
 - b. Easy Fabric for VXLAN.
 - c. External fabric for Route Server. Note that this fabric is optional if you are using Direct to BGWs topology.
 - d. External fabric to monitor the legacy devices.
 - e. MSD fabric as a container for all fabrics.
2. vPC connection from vPC BGWs towards the legacy site. It is expected that vPC from legacy towards BGWs is done out of band.
3. Multi-Site underlay eBGP inter-fabric connection (IFC) creation.
4. Multi-Site overlay eBGP IFC creation.

Topology Overview

Let us look at the example topology.



This topology contains the following five fabrics:

1. GATEWAY

This fabric is created for the vPC border gateways.

This fabric is an Easy fabric without any Spine nodes and it is set up as a regular Easy fabric with the following characteristics:

- Under the **Replication** tab, the **Replication Mode** is set to **Ingress**.
- The vPC Border gateways role are set as BGW.
- The IFC create method will be set to Manual or auto configuration as per user preference.
- Gateway fabric has a vPC interface configuration towards the Legacy Fabric.
- A member fabric of MSD.
- Save and deploy operation is performed in Easy fabric and MSD fabric.

2. LEGACY

This fabric is created for the Legacy network. The fabric type is External and could be kept in the monitor mode. Fully configured devices are imported into this fabric as shown in External Fabric procedure.

3. EasyFabric01

This represents a fully functional VXLAN fabric. The Border Gateway switches of this fabric are connected via IFC's to Route Servers or Direct to BGWs of Legacy fabric as per your topology. Both models are supported as shown in the Multi-Site use case.

4. RouteServers

In this topology, Centralized to Route Server topology is used. Typically, there would be more than one Route Server for redundancy reasons. This fabric is of type External as shown in the Multi-Site use case.

5. MSD

The MSD fabric is created to configure the base multi-site for the member fabrics. All the above four fabrics are imported into the MSD fabric for the BGW base. Optionally, you can enable auto-configuration of all underlay and overlay IFCs.

Configuring vPC from vPC Border Gateways to Legacy Network

In the **Manage Interfaces** window for the GATEWAY fabric, click the **Add (+)** icon and enter the information for the fields as shown in the following image. From the **Policy** drop-down list, select the vPC policy and fill in the fields for your topology.

Edit Configuration

Name: vPC_BGW_2~vPC_BGW_1:vPC1

Policy: int_vpc_trunk_host_11_1

Note : PeerOne = vPC_BGW_2 & PeerTwo = vPC_BGW_1

General

Peer-1 Port-Channel ID	1	Peer-1 VPC port-channel number (Min:1, Max:4096)
Peer-2 Port-Channel ID	1	Peer-2 VPC port-channel number (Min:1, Max:4096)
Peer-1 Member Interfaces	E1/21-24	A list of member interfaces for Peer-1 [e.g. e1/5,eth1/7-9]
Peer-2 Member Interfaces	E1/21-24	A list of member interfaces for Peer-2 [e.g. e1/5,eth1/7-9]
* Port Channel Mode	active	Channel mode options: on, active and passive
* Enable BPDU Guard	no	Enable spanning-tree bpduguard
Enable Port Type Fast	<input checked="" type="checkbox"/>	Enable spanning-tree edge port behavior
* MTU	jumbo	MTU for the Port Channel
* Peer-1 Trunk Allowed...	all	Allowed values: 'none', 'all', or vlan ranges (ex: 1-200,500-2000,3000)
* Peer-2 Trunk Allowed...	all	Allowed values: 'none', 'all', or vlan ranges (ex: 1-200,500-2000,3000)
Peer-1 PO Description	Add description to Peer-1 VPC port-channel (Max Size 254)	

Save Preview Deploy

Additional References

After entering all the information, click **Preview** to preview the configurations that are deployed, and then click **Deploy**.

Multi-Site Underlay eBGP IFC Creation

The Multi-Site underlay configuration is same as MSD shown in the Multi-Site use case. Choose GUI or autoconfiguration based method to create IFCs to the Core router or directly to BGW of other fabric, as per your topology.

In this topology, vPC Border Gateways are physically connected to Route Server (RS1), one MS underlay IFC is configured from each BGW (in GATEWAY and EasyFabric01) to RS1. Both methods are detailed in the Multi-Site use case.

Configuring Multi-Site Overlay IFCs

Multi-Site overlay IFCs need to be created between vPC BGWs to either a centralized route server or Direct to each BGW in **EasyFabric01**. In the example topology, there is one Overlay IFC from each BGW to RS1.

The summary of the IFCs for this topology are shown in the following image.

	Fabric Name	Name	Policy	Info	Admin State	Oper State
1	EasyFabric01<->RouteServers	BGW_1~loopback0—RS_1~Loopback0	ext_evpn_multisite_overlay_setup	Neighbor Missing	-:-	-:-
2	EasyFabric01<->RouteServers	BGW_2~loopback0—RS_1~Loopback0	ext_evpn_multisite_overlay_setup	Neighbor Missing	-:-	-:-
3	GATEWAY<->RouteServers	vPC_BGW_1~loopback0—RS_1~Loopback0	ext_evpn_multisite_overlay_setup	Neighbor Missing	-:-	-:-
4	GATEWAY<->RouteServers	vPC_BGW_2~loopback0—RS_1~Loopback0	ext_evpn_multisite_overlay_setup	Neighbor Missing	-:-	-:-
5	EasyFabric01<->RouteServers	BGW_1-Ethernet4/3—RS_1-Ethernet5/5	ext_multisite_underlay_setup_11_1	Link Present	Up:Up	Up:Up
6	EasyFabric01<->RouteServers	BGW_2-Ethernet1/51—RS_1-Ethernet5/6	ext_multisite_underlay_setup_11_1	Link Present	Up:Up	Up:Up
7	GATEWAY<->RouteServers	vPC_BGW_1-Ethernet1/14—RS_1-Ethernet5/7/2	ext_multisite_underlay_setup_11_1	Link Present	Up:Up	Up:Up
8	GATEWAY<->RouteServers	vPC_BGW_2-Ethernet1/13—RS_1-Ethernet5/7/3	ext_multisite_underlay_setup_11_1	Link Present	Up:Up	Up:Up

Additional References

Document Title and Link	Document Description
VXLAN EVPN Multi-Site Design and Deployment White Paper	This document explains Multi-Site design and deployment in detail.
Configuring VXLAN EVPN Multi-Site	This document explains manual configurations for the Multi-Site solution.

Appendix

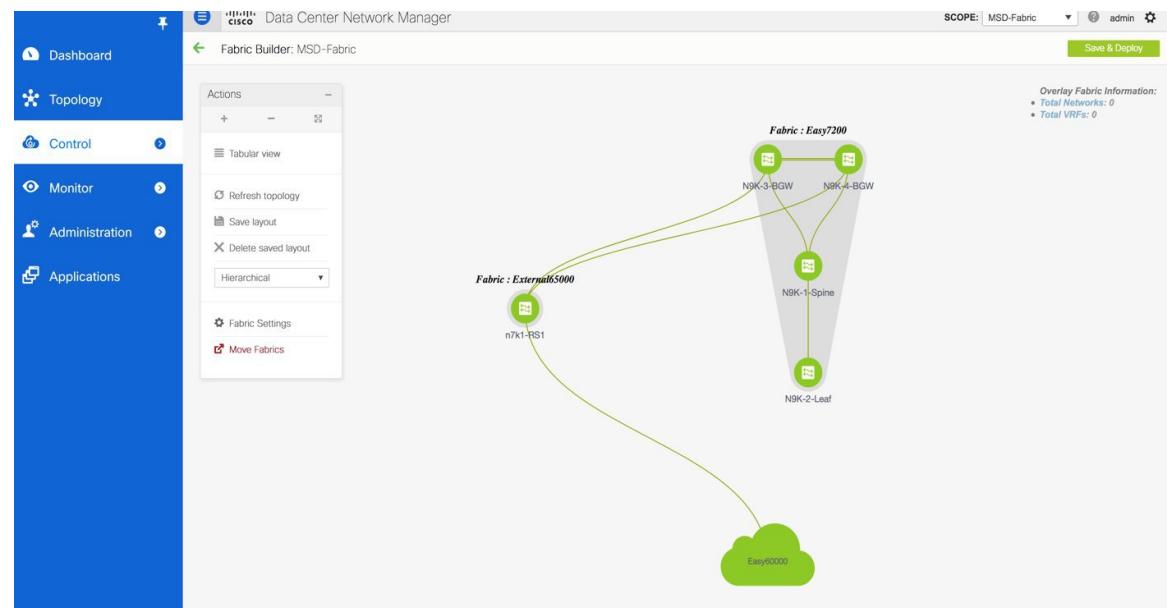
Multi-Site Fabric Base Configurations – Box Topology

In the Easy7200 fabric, N9K-3-BGW and N9K-4-BGW are connected to each other over two physical interfaces, and the BGWs do not form a vPC pair. Such a topology is called a Box topology. An IBGP session

is configured on each physical connection. One connection is between the Eth1/21 interfaces, and the other is between the Eth1/22 interfaces.

IBGP Configuration for the Box Topology in the Easy7200 Fabric

The following configuration is generated on each of the nodes if the fabric has numbered interfaces. In case the fabric interfaces are unnumbered, then the IBGP session is formed via the loopback0 address.

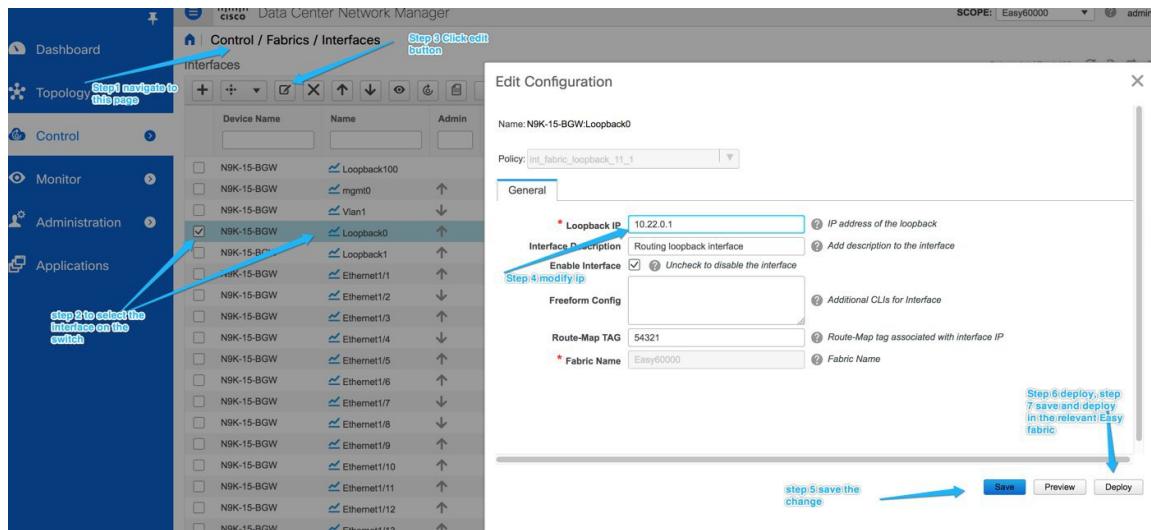


N9K-BGW-3	N9K-BGW-4
<pre>router bgp 7200 neighbor 10.4.0.17 remote-as 7200 update-source ethernet1/22 address-family ipv4 unicast next-hop-self</pre>	<pre>router bgp 7200 neighbor 10.4.0.18 remote-as 7200 update-source Ethernet1/22 address-family ipv4 unicast next-hop-self</pre>
<pre>router bgp 7200 neighbor 10.4.0.13 remote-as 7200 update-source ethernet1/21 address-family ipv4 unicast next-hop-self</pre>	<pre>router bgp 7200 neighbor 10.4.0.14 remote-as 7200 update-source Ethernet1/21 address-family ipv4 unicast next-hop-self</pre>
<pre>interface ethernet1/22 evpn multisite dci-tracking no switchport ip address 10.4.0.18/30 description connected-to-N9K-4-BGW--Ethernet1/22</pre>	<pre>interface Ethernet1/22 evpn multisite dci-tracking no switchport ip address 10.4.0.17/30 description connected-to-N9K-3-BGW-Ethernet1/22</pre>

Route Server Configuration

N9K-BGW-3	N9K-BGW-4
<pre>interface ethernet1/21 evpn multisite dci-tracking no switchport ip address 10.4.0.14/30 description connected-to-N9K-4-BGW-Ethernet1/21</pre>	<pre>interface Ethernet1/21 evpn multisite dci-tracking no switchport ip address 10.4.0.13/30 description connected-to-N9K-3-BGW-Ethernet1/21</pre>

Changing loopback0 Policy to Modify IP Address



Route Server Configuration

The route server overlay and base configurations are only deployed if the external fabric is not in Monitor mode.



Note When an external fabric is set to **Fabric Monitor Mode Only**, you cannot deploy configurations on its switches. Refer the *Creating an External Fabric* topic in the *Control* chapter for details.

Route Server Base Configuration - These are one time deployed on the route server and may be edited or deleted via the corresponding policy. The router server overlay and base configurations are only deployed if the external fabric is not in Monitor mode.

Configuration	Description
<pre>route-map unchanged permit 10 set ip next-hop unchanged</pre>	—

Configuration	Description
<pre>router bgp 65000 address-family ipv4 unicast network /32</pre>	The network command to redistribute the BGP peering address of RS1 to the eBGP underlay sessions so that BGWs know how to reach RS. If operator is using a different method to distribute the route server peering address to BGW, then this is not needed
<pre>interface ethernet1/22 evpn multisite dci-tracking no switchport ip address 10.4.0.18/30 description connected-to-N9K-4-BGW--Ethernet1/22</pre>	<pre>interface Ethernet1/22 evpn multisite dci-tracking no switchport ip address 10.4.0.17/30 description connected-to-N9K-3-BGW-Ethernet1/22</pre>
<pre>template peer OVERLAY-PEERING update-source loopback0 ebgp-multipath 5 address-family l2vpn evpn route-map unchanged out address-family l2vpn evpn retain route-target all send-community send-community extended</pre>	The knob in the external fabric controls if send community is sent in the form shown here, or as send-community both. If this form causes a persistent CC difference, then edit the policy on the device in the external fabric as shown in the Deploying the Send-Community Both Attribute section below.

Multi-Site Overlay IFC Configuration

In the reference topology, there are two BGWs in the Easy7200 fabric. Each BGW forms a BGP overlay connection with the route server.

BGW	Route Server
<pre>router bgp 7200 neighbor remote-as 65000 update-source loopback0 ebgp-multipath 5 peer-type fabric-external address-family l2vpn evpn send-community send-community extended rewrite-evpn-rt-asn</pre>	<pre>router bgp 65000 neighbor 10.2.0.1 remote-as 7200 inherit peer OVERLAY-PEERING address-family l2vpn evpn rewrite-evpn-rt-asn router bgp 65000 neighbor 10.2.0.2 remote-as 7200 inherit peer OVERLAY-PEERING address-family l2vpn evpn rewrite-evpn-rt-asn</pre>

See below for the configurations generated on the BGW and the route server.

Multi-Site Underlay IFC Configuration – Out-of-Box Profiles

The following table shows the Multi-Site IFC configuration deployed by DCNM with the out-of-the box profiles. If the IFC is between two VXLAN fabrics, then both sides have the BGW configurations shown below.

Multi-Site Underlay IFC Configuration – Out-of-Box Profiles

BGW Configuration	Core Router Configuration
<pre>router bgp 7200 neighbor 10.10.1.6 remote-as 65000 update-source ethernet1/47 address-family ipv4 unicast next-hop-self</pre>	<pre>router bgp 65000 neighbor 10.10.1.5 remote-as 7200 update-source ethernet7/4/1 address-family ipv4 unicast next-hop-self</pre>
<pre>interface ethernet1/47 mtu 9216 no shutdown no switchport ip address 10.10.1.5/30 tag 54321 evpn multisite dci-tracking</pre>	<pre>interface ethernet7/4/1 mtu 9216 no shutdown no switchport ip address 10.10.1.6/30 tag 54321</pre>

The tag 54321 attached to the IP address is not required for correct functioning and will be removed in subsequent releases. It is benign.