



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.

* Fabric Name :

* Fabric Template :

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

Enable MPLS Handoff ?

* Underlay MPLS Loopback Id ? 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 n. Any mismatches will yield unexpected diffs during de

? Note ! All configs should strictly match 'show run' ou

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 :

* Fabric Template :

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 ? *Typically Loopback0 IP Address Range*

* Underlay VTEP Loopback IP Range ? *Typically Loopback1 IP Address Range*

* Underlay RP Loopback IP Range ? *Anycast or Phantom RP IP Address Range*

* Underlay Subnet IP Range ? *Address range to assign Numbered and Peer Link SVI IPs*

* Underlay MPLS Loopback IP Range ? *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 ? *Overlay Network Identifier Range (Min:1, Max:16777214)*

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

* Network VLAN Range ? *Per Switch Overlay Network VLAN Range (Min:2, Max:3967)*

* VRF VLAN Range ? *Per Switch Overlay VRF VLAN Range (Min:2, Max:3967)*

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

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.

* Fabric Name :

* Fabric Template :

General | **Advanced** | Resources | Configuration Backup | Bootstrap

* vPC Peer Link VLAN ? VLAN for vPC Peer Link SVI (Min:2, Max:3967)

* Power Supply Mode ? Default Power Supply Mode For The Fabric

Enable MPLS Handoff ?

* Underlay MPLS Loopback Id ? (Min:0, Max:1023)

Enable AAA IP Authorization ? Enable only, when IP Authorization is enabled in the AAA Server

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.

* Fabric Name :

* Fabric Template :

General | Advanced | **Resources** | Configuration Backup | Bootstrap

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

* Underlay Routing Loopback IP Range ? Typically Loopback0 IP Address Range

* Underlay MPLS Loopback IP Range ? MPLS Loopback IP Address Range

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**.
- Step 2** Choose a VXLAN fabric from which you want to create an underlay inter-fabric connection to MPLS.
- Step 3** Click **Tabular view** in the **Actions** panel that is displayed at the upper left of the window.
- Step 4** 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.

The screenshot shows the Cisco Data Center Network Manager interface. The top navigation bar includes the Cisco logo, 'Data Center Network Manager', and a 'SCOPE: easy101' dropdown. Below this, the 'Fabric Builder: easy101' section is visible, with a 'Save & Deploy' button. The 'Links' tab is active, showing a table of discovered links. The table has columns for 'Update Link', 'Name', 'Name', 'Policy', 'Info', 'Admin State', and 'Oper State'. Link 3 is selected, and the 'Update Link' icon is highlighted.

	Update Link	Name	Name	Policy	Info	Admin State	Oper State
1	<input type="checkbox"/>	easy101	n3k-31-Ethernet1/3---n9k-1-spine-Ethernet2/1		Neighbor Present	Up:-	Up:-
2	<input type="checkbox"/>	easy101	n3k-31-Ethernet1/2---n9k-17-Ethernet2/5		Neighbor Present	Up:-	Up:-
3	<input checked="" type="checkbox"/>	easy101<->external...	n3k-31-Ethernet1/5---n3k-33-Ethernet1/5	ext_vxlan_mpls_underlay_setup	Link Present	Up:Up	Up:Up
4	<input type="checkbox"/>	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

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 / Advanced

* 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)

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	
Advanced	
* IP Address/Mask	10.33.102.1/30 <small>? IP address with mask (e.g. 192.168.10.1/24)</small>
* Neighbor IP	10.33.102.2 <small>? Neighbor IP address</small>
* MPLS Fabric	LDP <small>? MPLS LDP or Segment-Routing</small>
Source SR Index	<small>? Unique SR SID index for the source border</small>
Destination SR Index	<small>? Unique SR SID index for the destination border</small>
SR Global Block Range	<small>? For Segment Routing binding (Min:16000, Max:471804)</small>
* DCI Routing Protocol	is-is <small>? Routing protocol used on the DCI MPLS link</small>
OSPF Area ID	<small>? OSPF Area ID in decimal format or IP address format</small>
* DCI Routing Tag	MPLS_UNDERLAY <small>? Routing Process Tag (Max Size 20)</small>

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**.

- 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
✕

* **Link Type**

* **Link Sub-Type**

* **Link Template**

* **Source Fabric**

* **Destination Fabric**

* **Source Device**

* **Source Interface**

* **Destination Device**

* **Destination Interface**

▼ **Link Profile**

General

* **BGP Local ASN** ? BGP Local Autonomous System Number

* **BGP Neighbor IP** ? Neighbor IP address for eBGP peering

* **BGP Neighbor ASN** ? BGP Neighbor Autonomous System Number

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](#).

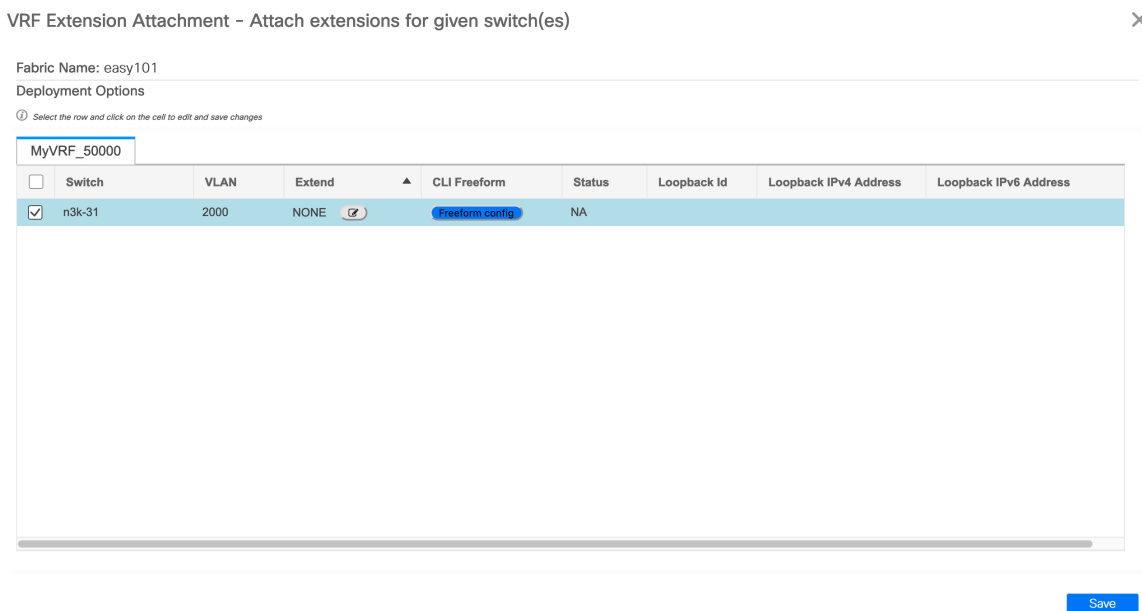
Step 3 Select the newly added VRF and click **Continue**.



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.



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.

