

MPLS SR to VxLAN Handoff

- MPLS Segment Routing to VxLAN Handoff, on page 1

MPLS Segment Routing to VxLAN Handoff

MPLS SR to VxLAN handoff enables seamless routing and forwarding between MPLS Segment Routing (SR) domains and VxLAN overlays in data center and WAN edge architectures.

- Interconnects MPLS SR (WAN/core) and VxLAN EVPN (data center) domains.
- Maintains L3VPN segmentation using per-VRF label allocation.
- Handles control plane and data plane translation, including next-hop resolution, label/VNI mapping, and QoS marking.

The MPLS SR to VxLAN handoff is a gateway capability on border leaf or spine devices that enables the transition of routed traffic between an MPLS Segment Routing core and a VxLAN-based overlay network.

How MPLS Segment Routing to VxLAN Handoff Works

The handoff enables communication between a core MPLS SR network and a VxLAN-based data center fabric, typically at the border leaf or spine (DCI node). This is essential for multi-domain connectivity, data center expansion, and migration scenarios.

- The DCI node acts as the gateway, performing protocol translation and encapsulation/decapsulation between MPLS SR and VxLAN overlays.

Summary

This process describes how traffic is handed off between an MPLS Segment Routing (SR) network and a VxLAN overlay at the Data Center Interconnect (DCI) node, enabling seamless L3VPN connectivity between MPLS and VxLAN domains.

Key components in this process are:

- DCI Node (Border Leaf/Border Spine): Performs the handoff and encapsulation functions between MPLS SR and VxLAN overlays.
- MPLS SR Core: Provides L3VPN connectivity using segment routing.

- VxLAN EVPN Fabric: Connects ToRs and other leaf switches using VxLAN overlays.

Workflow

1. Route and Label Advertisement
   - The DCI node receives BGP route updates from both the MPLS SR domain and the VxLAN EVPN domain, including VPN labels and next-hop information.
   - BGP control plane exchanges ensure appropriate import/export of routes between domains using route targets.

When...	And...	Then...	And...
A new host comes online in the VxLAN domain	The DCI node imports the EVPN route	The DCI reoriginates and advertises a VPN label to the MPLS SR core	The route is reachable from both domains

   The DCI node synchronizes the control plane state across both domains using BGP, with appropriate label allocation per VRF.

   The result is that both MPLS SR and VxLAN domains can learn and use routes across the handoff boundary.

2. Data Plane Handoff (Packet Forwarding)
   - Packets arriving from the MPLS SR core are decapsulated and re-encapsulated into VxLAN (or vice versa) by the DCI node.
   - QoS, TTL, and ECN fields are mapped between MPLS and VxLAN headers according to platform-specific rules (e.g., uniform or pipe modes).

When...	And...	Then...	And...
A packet arrives from the MPLS SR core with a VPN label	The DCI node matches the label to a VRF and destination	The DCI strips MPLS headers and applies VxLAN encapsulation with the correct VNI	The packet is forwarded into the VxLAN fabric towards the destination host
A packet arrives from the VxLAN domain with a VNI	The DCI node matches the VNI to a VRF and destination	The DCI strips VxLAN headers and applies MPLS SR encapsulation with the correct label stack	The packet is forwarded into the MPLS SR core towards the remote PE

   The DCI node ensures the correct translation and forwarding of packets between domains, applying platform-specific QoS and statistics handling.

   The result is end-to-end traffic flow between MPLS SR and VxLAN domains, with resiliency provided by fallback to underlay if the VxLAN overlay is unavailable.

3. Resiliency and Fallback Handling
   - If the VxLAN NVE interface is down, the DCI node automatically falls back to using the MPLS SR underlay for next-hop resolution, maintaining reachability.

When...	And...	Then...	And...
NVE1 (VxLAN) is down on the DCI node	The VxLAN overlay is unavailable	The next-hop is resolved via MPLS SR underlay routes	Traffic continues to flow using backup MPLS SR paths until overlay is restored

This stage maintains operational continuity and network resilience by leveraging dual-domain routing.

The result is uninterrupted service during overlay outages, with automatic return to VxLAN overlay when available.

Guidelines and Limitations

Platform and Feature Guidelines

Use only supported hardware and configurations for successful deployment of SR MPLS to VxLAN handoff. Ensure compatibility and adhere to supported operational modes to prevent unexpected results.

- This feature is supported only on Cisco Nexus 9000 Cloudscale platforms, including FX2, FX3, GX, GX2, and select modular platforms.
- Coexistence of VxLAN-EVPN and MPLS Segment Routing (SR-MPLS) or MPLS L3VPN (LDP) features is required on the same device for DCI handoff.
- vPC, VMCT, and pMCT configurations are not supported with SR MPLS to VxLAN handoff.
- Supported only on physical interfaces for core-facing (WAN) ports. SVI and sub-interface handoff for core-facing links is not supported.
- Only per-VRF label allocation is supported for VPN label assignment. Per-prefix label allocation is not supported.

Operational Limitations and Restrictions

Be aware of all known operational limitations when deploying SR MPLS to VxLAN handoff. Avoid unsupported configurations and understand the impact on failover, statistics, and scale.

- Only EVPN Type 5 (IP prefix routes) are supported for handoff to MPLS/SR core. Subnet (Type 2) handoff and L2 extension are not supported in the current release.
- Beginning with Cisco NX-OS Release 10.5(3)F, MPLS Segment Routing can be enabled on Border Gateway (BGW for VXLAN Multi-Site). In the earlier releases, MPLS Segment Routing and Border Gateway (BGW for VXLAN Multi-Site) could not be enabled at the same time.
- On some platforms, MPLS and VxLAN statistics are not supported; on FX2, only VPN label statistics are available (no LSR or adjacency stats).
- End-to-end TTL and ECN propagation is not fully supported due to ASIC limitations. Only pipe mode TTL is supported at the handoff.
- FX2 Platform supports a maximum of 256 VxLAN peers, 900 VRFs (of which up to 100 can be extended to MPLS), 48,000 adjacencies, and 500 MPLS labels.
- Priority Flow Control (PFC) is not supported in DCI handoff mode.

- Route leaking or VRF import/export between VxLAN and MPLS domains is not supported; only same-VRF handoff is allowed.
- During a failure or shutdown of the NVE (VxLAN) interface, next-hop resolution falls back to the MPLS underlay, which is expected behavior for resiliency.

Configure MPLS SR to VXLAN Handoff

This procedure enables seamless routing and forwarding between an MPLS SR core and a VXLAN EVPN-based data center fabric at the DCI border.

This configuration is required when connecting a VXLAN EVPN data center fabric to an MPLS SR or LDP-based WAN/core using a Nexus 9000 as the DCI/border device.

Before you begin

Ensure that the device has appropriate hardware resources for both VXLAN and MPLS SR features.

- Licensing for both VXLAN and MPLS features is applied.
- Required VLANs, VRFs, and interfaces are provisioned.

Follow these steps to configure MPLS SR to VXLAN handoff on the DCI node.

Procedure

Step 1 Enable required features and global configurations.

Example:

```
switch# configure terminal
switch(config)# feature-set mpls
switch(config)# feature ospf
switch(config)# feature bgp
switch(config)# feature mpls l3vpn
switch(config)# feature mpls segment-routing
switch(config)# feature nv overlay
switch(config)# feature vn-segment-vlan-based
```

Ensure all features are enabled before proceeding with interface and protocol configuration.

Note
Not all features are available on every platform or NX-OS release.

Required features are enabled, and the device is prepared for further configuration.

Step 2 Configure VRFs and VXLAN-to-MPLS mapping

Example:

```
switch(config)# vrf context Tenant-A         <<< Create VRF
switch(config-vrf)# vni 10010                        <<< Map VNI to VRF
switch(config-vrf)# rd auto
switch(config-vrf)# address-family ipv4 unicast
switch(config-vrf-af)# route-target both 65000:10010
switch(config-vrf-af)# exit

switch(config-vrf)# address-family l2vpn evpn
```

```
switch(config-vrf-af)# route-target import 65000:10010 evpn
switch(config-vrf-af)# route-target export 65000:10010 evpn
switch(config-vrf-af)# exit

switch(config-vrf)# exit

switch(config)# interface nve1                          <<< Configure NVE Interface
switch(config-if-nve)# no shutdown
switch(config-if-nve)# host-reachability protocol bgp <<< Enable BGP for EVPN
switch(config-if-nve)# source-interface loopback1      <<< Set NVE source
switch(config-if-nve)# member vni 10010 associate-vrf  <<< Associate VNI with VRF
switch(config-if-nve)# exit

switch(config)#
```

Route-targets must match between DC and WAN to import/export correct L3VPN routes.

Note
Only per-VRF label allocation is supported for VXLAN–MPLS interworking.

VRFs and VNIs are mapped for inter-domain handoff.

Step 3 Configure interfaces for both MPLS and VXLAN connectivity

Example:

```
switch(config)# interface Ethernet1/21
switch(config-if)# ip address 6.2.0.1/24
switch(config-if)# mpls ip forwarding
switch(config-if)# no shutdown
switch(config-if)# exit

switch(config)# interface Ethernet1/21.1
switch(config-if)# encapsulation dot1q 1211
switch(config-if)# vrf member evpn
switch(config-if)# ip address 6.22.0.1/24
switch(config-if)# no shutdown
switch(config-if)# exit

switch(config)# end
switch#
```

Enable MPLS on WAN/core-facing interfaces. Assign VRFs and IP addresses per design.

Note
Only L3 physical interfaces are supported for core-facing MPLS links.

All physical and logical interfaces for VXLAN and MPLS are configured and active.

Step 4 Set up BGP with appropriate address families and route re-origination

Example:

```
switch(config)# router bgp 600
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)# redistribute direct route-map passall
switch(config-router-af)# allocate-label all              <<< Enables per-prefix label
allocation for MPLS VPNv4/vpnv6 routes (DCI/SR-MPLS handoff)
switch(config-router-af)# exit

switch(config-router)# neighbor 6.6.6.3
switch(config-router-neighbor)# remote-as 300
switch(config-router-neighbor)# update-source loopback6
switch(config-router-neighbor)# ebgp-multihop 255
```

```
switch(config-router-neighbor)# address-family vpnv4 unicast
switch(config-router-neighbor-af)# send-community
switch(config-router-neighbor-af)# send-community extended
switch(config-router-neighbor-af)# next-hop-self
switch(config-router-neighbor-af)# import l2vpn evpn reoriginate    <<< Enables import and
 re-origination of EVPN routes into MPLS VPNv4 for DCI handoff
switch(config-router-neighbor-af)# exit

switch(config-router-neighbor)# exit

switch(config-router)# neighbor 21.21.21.21
switch(config-router-neighbor)# remote-as 600
switch(config-router-neighbor)# update-source loopback1
switch(config-router-neighbor)# address-family l2vpn evpn
switch(config-router-neighbor-af)# send-community
switch(config-router-neighbor-af)# send-community extended
switch(config-router-neighbor-af)# import vpn unicast reoriginate   <<< Enables import and
 re-origination of VPNv4/vpnv6 routes into EVPN for DCI handoff
switch(config-router-neighbor-af)# exit

switch(config-router-neighbor)# exit

switch(config-router)# exit

switch(config)#
```

Configure BGP neighbors for both MPLS (WAN/core) and VXLAN (fabric) sides, enabling import and reorigination for cross-domain route exchange.

Note
Use **import l2vpn evpn reoriginate** and **import vpn unicast reoriginate** for bidirectional handoff.

BGP sessions are established and routes are exchanged between domains.

MPLS SR to VXLAN handoff is successfully configured, enabling seamless L3VPN connectivity between data center and core network domains.

Verify the DCI VxLAN-MPLS Handoff

Perform this verification to confirm that the DCI device correctly forwards traffic between VxLAN and MPLS domains.

- Confirm that overlay and underlay routing tables are populated as expected.
- Check that interface and protocol states are up and operational.
- Validate correct data plane handoff by simulating end-to-end traffic.

Before you begin

Before starting this verification, ensure that all relevant VxLAN, EVPN, MPLS, and BGP configurations have been applied and the devices have completed initial convergence.

- All physical and logical interfaces involved in the handoff are up and configured.
- Control-plane protocols (BGP, OSPF/ISIS, etc.) are established and stable.

Follow these steps to verify the DCI VxLAN-MPLS handoff functionality.

Procedure

Step 1 Check the status of NVE and MPLS interfaces on the DCI node.

Example:

```
switch# show nve interface
switch# show interface nve1
switch# show mpls interfaces
```

These commands display the operational status of the VxLAN and MPLS interfaces. Both must be **up** for successful handoff.

If either interface is down, troubleshoot physical connectivity, configuration, or protocol states before proceeding.

Step 2 Verify route propagation and label allocation between VxLAN and MPLS domains.

Example:

```
switch# show bgp l2vpn evpn summary
switch# show bgp vpnv4 unicast summary
switch# show mpls forwarding-table
```

These commands confirm that route and label exchanges are occurring correctly between the overlay (EVPN) and underlay (MPLS) domains.

- Check for expected prefixes and label bindings for all relevant VRFs.
- Absence of routes or labels may indicate misconfiguration or protocol issues.

Step 3 Test end-to-end traffic forwarding across the DCI node.

Example:

```
switch# ping [destination-ip] vrf [vrf-name] source [source-ip]
switch# traceroute [destination-ip] vrf [vrf-name]
```

These tests verify data plane connectivity and the proper functioning of the handoff between VxLAN and MPLS domains.

- Successful pings and traceroutes indicate a working handoff.
- Failures may indicate issues with routing, label allocation, or interface states.

At the end of this procedure, you have verified that the DCI node correctly hands off traffic between VxLAN and MPLS domains, with all routing, label, and interface states operational. Traffic forwarding is confirmed by end-to-end connectivity tests.