



Configuring Seamless Integration of EVPN with L3VPN SRv6

This chapter contains the following sections:

- [About Seamless Integration of EVPN with L3VPN SRv6 Handoff, on page 1](#)
- [Guidelines and Limitations for EVPN to L3VPN SRv6 Handoff, on page 2](#)
- [Importing L3VPN SRv6 Routes into EVPN VXLAN, on page 3](#)
- [Importing EVPN VXLAN Routes into L3VPN SRv6, on page 5](#)
- [Example Configuration for VXLAN EVPN to L3VPN SRv6 Handoff, on page 6](#)


About Seamless Integration of EVPN with L3VPN SRv6 Handoff

Data Center (DC) deployments have adopted VXLAN EVPN for its benefits such as EVPN control-plane learning, multitenancy, seamless mobility, redundancy, and easier POD additions. Similarly, the CORE is either an IP-based L3VPN SRv6 network or transitioning from the IPv6-based L3VPN underlay to a more sophisticated solution like IPv6 Segment Routing (SRv6) for IPv6. SRv6 is adopted for its benefits such as:

- Simpler traffic engineering (TE) methods
- Easier configuration
- SDN adoption

With two different technologies, one within the data center (DC) and one in the Core, there is traffic handoff from VXLAN to an SRv6 core that becomes a necessity at the DCI nodes, which sit at the edge of the DC domain and interface with the Core edge router.

Figure 1: BGP EVPN VXLAN to L3VPN SRv6 Handoff



For traffic ingressing the EVPN-VxLAN fabric, the BGP EVPN routes get imported into a local VRF which contains the RD of the VRF. The bestpath is calculated and installed in the VRF's RIB, then inserted into the L3VPN SRv6 table. Along with the bestpath, the VRF's RD and per-VRF SRv6 SID are included. The L3VPN SRv6 route target is sent with the route, which is advertised to the L3VPN SRv6 peer.

For traffic egressing the EVPN VxLAN fabric, the BGP L3VPN SRv6 routes get imported into a local VRF which contains the RD of the VRF. The bestpath is calculated and installed in the VRF's RIB, then inserted into the EVPN table. Along with the bestpath, the VRF's RD and VNI are included. The EVPN-VXLAN route target is sent with the route, which is advertised to the EVPN-VxLAN peer.

Guidelines and Limitations for EVPN to L3VPN SRv6 Handoff

This feature has the following guidelines and limitations:

- The same RD import is supported for L3VPN SRv6 fabrics.
- The same RD import is not supported for EVPN VXLAN fabrics.
- On a handoff device, do not use the same RD import on the EVPN VXLAN side.
- Beginning with Cisco NX-OS Release 9.3(3), support is added for the following switches:

- Cisco Nexus C93600CD-GX
 - Cisco Nexus C9364C-GX
 - Cisco Nexus C9316D-GX
- Beginning with Cisco NX-OS Release 10.2(1q)F, SRv6 DCI handoff is supported on Cisco Nexus 9332D-GX2B platform switches.
 - Beginning with Cisco NX-OS Release 10.4(1)F, SRv6 DCI handoff is supported on Cisco Nexus 9332D-H2R switches.
 - Beginning with Cisco NX-OS Release 10.4(2)F, SRv6 DCI handoff is supported on Cisco Nexus 93400LD-H1 switches.
 - Beginning with Cisco NX-OS Release 10.4(3)F, SRv6 DCI handoff is supported on Cisco Nexus 9364C-H1 switches.
 - Beginning with Cisco NX-OS Release 10.2(3)F, EVPN to L3VPN SRv6 Handoff is supported on Cisco Nexus 9300-GX2 platform switches.
 - Beginning with Cisco NX-OS Release 10.4(1)F, EVPN to L3VPN SRv6 Handoff is supported on Cisco Nexus 9332D-H2R switches.
 - Beginning with Cisco NX-OS Release 10.4(2)F, EVPN to L3VPN SRv6 Handoff is supported on Cisco Nexus 93400LD-H1 switches.
 - Beginning with Cisco NX-OS Release 10.4(3)F, EVPN to L3VPN SRv6 Handoff is supported on Cisco Nexus 9364C-H1 switches.

Importing L3VPN SRv6 Routes into EVPN VXLAN

The process of handing off routes from the L3VPN SRv6 domain to the EVPN VXLAN fabric requires configuring the import condition for L3VPN SRv6 routes. Routes can be either IPv4 or IPv6. This task configures unidirectional route advertisement into the EVPN VXLAN fabric. For bidirectional advertisement, you must explicitly configure the import condition for the L3VPN SRv6 domain.

Before you begin

Make sure you have a fully configured L3VPN SRv6 fabric. For more information, see "Configuring Layer 3 VPN over SRv6" in the *Cisco Nexus 9000 Series NX-OS SRv6 Configuration Guide*.

SUMMARY STEPS

1. **config terminal**
2. **router bgp** *as-number*
3. **neighbor bgp** *ipv6-address remote-as as-number*
4. **address family vpnv4 unicast** or **address family vpnv6 unicast**
5. **import l2vpn evpn route-map** *name* [reoriginate]

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	config terminal Example: <pre>switch-1# config terminal</pre> Enter configuration commands, one per line. End with CNTL/Z. <pre>switch-1(config)#</pre>	Enter configuration mode.
Step 2	router bgp as-number Example: <pre>switch-1(config)# router bgp 100</pre> <pre>switch-1(config-router)#</pre>	Enter BGP router configuration mode.
Step 3	neighbor bgp ipv6-address remote-as as-number Example: <pre>switch-1(config-router)# neighbor 1234::1 remote-as 200</pre> <pre>switch-1(config-router-neighbor)#</pre>	Enter BGP router configuration mode.
Step 4	address family vpnv4 unicast or address family vpnv6 unicast Example: <pre>switch-1(config-router-neighbor)# address-family vpnv4 unicast</pre> <pre>switch-1(config-router-neighbor-af)#</pre> Example: <pre>switch-1(config-router-neighbor)# address-family vpnv6 unicast</pre> <pre>switch-1(config-router-neighbor-af)#</pre>	Configure the IPv4 or IPv6 address family for unicast traffic that the EVPN VXLAN will handoff to L3VPN SRv6.
Step 5	import l2vpn evpn route-map name [reoriginate] Example: <pre>switch-1(config-router-neighbor-af)# import l2vpn evpn route-map test reoriginate</pre> <pre>switch-1(config-router-neighbor-af)#</pre>	Configure the IPv4 or IPv6 address family for unicast traffic that EVPN VXLAN will handoff to L3VPN SRv6. This command enables routes learned from L3VPN SRv6 domain to be advertised to the EVPN VXLAN domain. Using the optional reoriginate keyword advertises only domain-specific RTs.

What to do next

For bidirectional route advertisement, configure importing EVPN VXLAN routes into the L3VPN SRv6 domain.

Importing EVPN VXLAN Routes into L3VPN SRv6

The process of handing off routes from the EVPN VXLAN fabric to the L3VPN SRv6 domain requires configuring the import condition for EVPN VXLAN routes. Routes can be either IPv4 or IPv6. This task configures unidirectional route advertisement into the L3VPN SRv6 fabric. For bidirectional advertisement, you must explicitly configure the import condition for the EVPN VXLAN domain.

Before you begin

Make sure you have a fully configured L3VPN SRv6 fabric. For more information, see "Configuring Layer 3 VPN over SRv6" in the *Cisco Nexus 9000 Series NX-OS SRv6 Configuration Guide*.

SUMMARY STEPS

1. **config terminal**
2. **router bgp *as-number***
3. **neighbor *ipv6-address* remote-as *as-number***
4. **address-family *l2vpn evpn***
5. **import vpn unicast route-map *name* [reoriginate]**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	config terminal Example: <pre>switch-1# config terminal</pre> Enter configuration commands, one per line. End with CNTL/Z. <pre>switch-1(config)#</pre>	Enter configuration mode.
Step 2	router bgp <i>as-number</i> Example: <pre>switch-1(config)# router bgp 200</pre> <pre>switch-1(config-router)#</pre>	Enter BGP router configuration mode.
Step 3	neighbor <i>ipv6-address</i> remote-as <i>as-number</i> Example: <pre>switch-1(config-router)# neighbor 1234::1 remote-as</pre> <pre>100</pre> <pre>switch-1(config-router-neighbor)#</pre>	Enter BGP router configuration mode.
Step 4	address-family <i>l2vpn evpn</i> Example: <pre>switch(config-router-neighbor)# address-family</pre> <pre>l2vpn evpn</pre> <pre>switch(config-router-neighbor-af)#</pre>	Configure the address family for unicast traffic that EVPN VXLAN will handoff to L3VPN SRv6.

	Command or Action	Purpose
Step 5	import vpn unicast route-map <i>name</i> [reoriginate] Example: <pre>switch-1(config-router-neighbor-af)# import vpn unicast route-map test reoriginate switch-1(config-router-neighbor-af)#</pre>	Configure the IPv4 or IPv6 address family for unicast traffic that EVPN VXLAN will handoff to L3VPN SRv6. This command enables routes learned from the EVPN VXLAN domain to be advertised to the L3VPN SRv6 domain. Using the optional reoriginate keyword advertises only domain-specific RTs.

What to do next

For bidirectional route advertisement, configure importing L3VPN SRv6 routes into the EVPN VXLAN fabric.

Example Configuration for VXLAN EVPN to L3VPN SRv6 Handoff

```
feature vn-segment-vlan-based
feature nv overlay
feature interface-vlan
nv overlay evpn
feature srv6

vrf context customer1
vni 10000
rd auto
address-family ipv4 unicast
route-target both 1:1
route-target both auto evpn
address-family ipv6 unicast
route-target both 1:1
route-target both auto evpn

segment-routing
srv6
encapsulation
source-address loopback1
locators
locator DCI_1
prefix café:1234::/64

interface loopback0
ip address 1.1.1.0/32

interface loopback1
ip address 1.1.1.1/32
ipv6 address 4567::1/128

interface nve1
source-interface loopback0
member vni 10000 associate-vrf
host-reachability protocol bgp

vlan 100
vn-segment 10000

interface vlan 100
ip forward
```

```
ipv6 address use-link-local-only
vrf member customer1

router bgp 65000
  segment-routing srv6
  locator DCI_1
  neighbor 2.2.2.2 remote-as 200
    remote-as 75000
  address-family l2vpn evpn
    import vpn route-map | reoriginate
  neighbor 1234::1 remote-as 100
    remote-as 65000
  address-family vpv4 unicast
  import l2vpn evpn route-map | reoriginate
  address-family vpv6 unicast
  import l2vpn evpn route-map | reoriginate

vrf customer
  segment-routing srv6
  alloc-mode per-vrf
  address-family ipv4 unicast
  address-family ipv6 unicast
```



Note In the **vni number** command, do not use the **L3** keyword during configuration of VNI under VRF, as the new L3 VNI configuration is not supported on VLAN-BD for VNIs which are dynamically allocated.
