

Configure VXLAN eBGP Peering with External Node on Nexus 9000

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Topology](#)

[BL Configuration for leaking:](#)

[Verify Route towards External router](#)

[Configure](#)

[Create dedicated loopback](#)

[Configure](#)

[Advertise loopback on BGP](#)

[Configure](#)

[Verify](#)

[Verify route is imported to BGP tenant VRF on LEAF VTEP](#)

[Verify route is imported to BGP L2VPN on LEAF VTEP](#)

[Verify route is imported received on BL tenant VRF](#)

[Verify route is imported received on BL default VRF](#)

[Configure eBGP peering on LEAF](#)

[Configure](#)

[Configure eBGP peering on External router \(EXT-R\)](#)

[Configure](#)

[Summary steps](#)

[LEAF](#)

[EXT-R](#)

[Verify](#)

[Verify eBGP neighbor on LEAF](#)

[Verify eBGP neighbor on External Router](#)

Introduction

This document describes how to configure eBGP peering between VTEP and external node.

Prerequisites

Existing VXLAN environment.

Requirements

CISCO recommends that you have knowledge of these topics:


- NXOS Platform
- QoS
- Elan understanding
- VXLAN VRF Laking

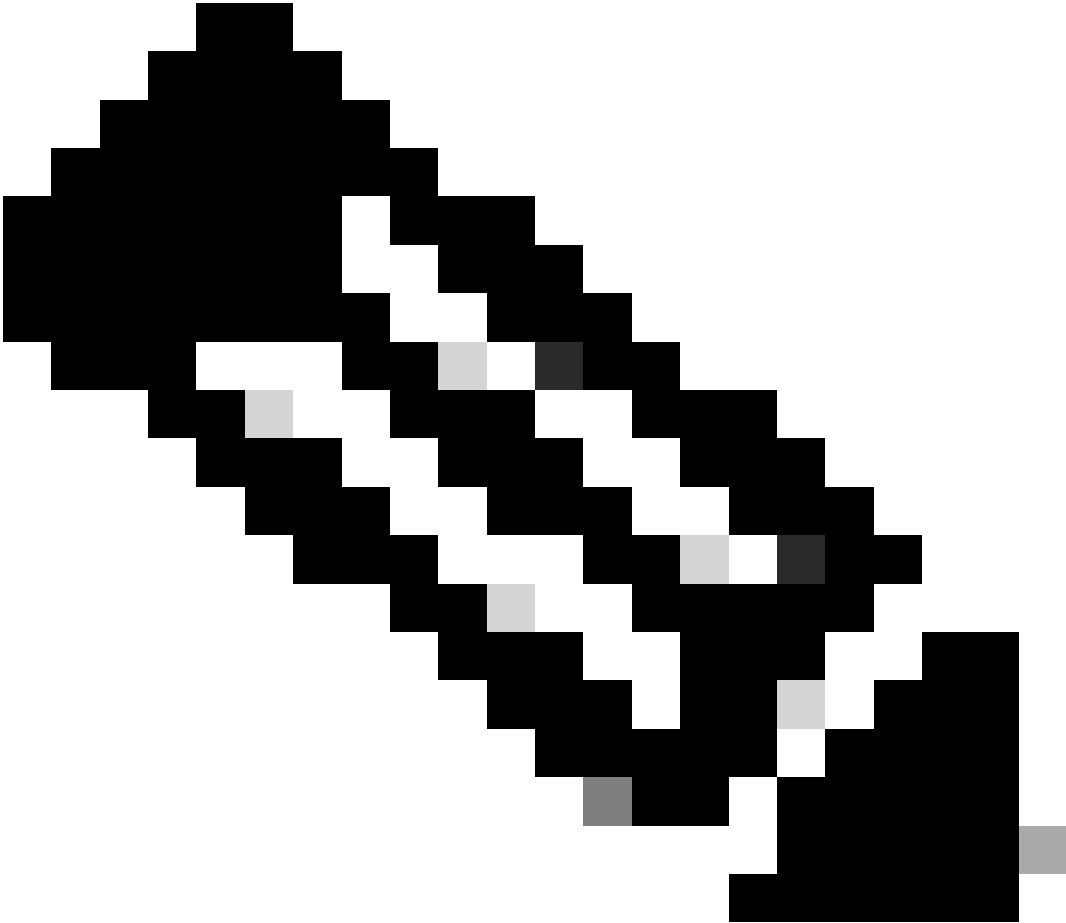
Components Used

Name	Platform	Version
LEAF	N9K-C9332D-GX2B	10.2(6)
BL	N9K-C9348D-GX2A	10.2(5)
EXT-R	N9K-C9348D-GX2A	10.2(3)
SPINE	N9K-C93108TC-FX3P	10.1(1)

"The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command."

Topology





Note: For this example an eBGP peering session is going to be configured between Leaf (vtep) and EXT-R (external device, out of VXLAN fabric.)

For this example BL knows EXT-R IP via OSPF on default vrf and knows LEAF loopback 100 IP via VXLAN and it is redistributed to OSPF in default VRF.

In order to perform this leaking Tech Zone Configure and Verify VXLAN VRF leaking on Nexus Switches can be reviewed: <https://www.cisco.com/c/en/us/support/docs/switches/nexus-9000-series-switches/221709-configure-and-verify-vxlan-vrf-leaking-o.html>

To simplify this document a summary of VRF leaking configuration on BL is written below.

BL Configuration for leaking:

```
ip prefix-list VXLAN-VRF-default-to-Tenant permit 172.16.120.55/32
```

```
route-map VXLAN-VRF-default-to-Tenant  
  match ip address prefix-list VXLAN-VRF-default-to-Tenant
```

```

ip prefix-list VXLAN-VRF-Tenant-to-default permit 192.168.150.150/32

route-map VXLAN-VRF-Tenant-to-default
 match ip address prefix-list VXLAN-VRF-Tenant-to-default

route-map BGP-OSPF
 match ip address prefix-list VXLAN-VRF-Tenant-to-default
 match route-type internal

router bgp 65000
 address-family ipv4 unicast
 redistribute ospf 1 route-map VXLAN-VRF-default-to-Tenant

vrf context tenant-a
 address-family ipv4 unicast
 import vrf default map VXLAN-VRF-default-to-Tenant advertise-vpn

vrf context tenant-a
 address-family ipv4 unicast
 export vrf default map VXLAN-VRF-Tenant-to-default allow-vpn

router ospf 1
 redistribute bgp 65000 route-map BGP-OSPF

```

Verify Route towards External router

The first step is to verify if in VTEP exists a route towards External router.

In this example BL is leaking route 172.16.120.55 from default VRF to tenant VRF and advertising it to LEAF.

```

LEAF# show ip route 172.16.120.55 vrf tenant-a
IP Route Table for VRF "tenant-a"
'*' denotes best ucast next-hop
 '**' denotes best mcast next-hop
 '[x/y]' denotes [preference/metric]
 '%<string>' in via output denotes VRF <string>

172.16.120.55/32, ubest/mbest: 1/0
 *via 172.16.0.5%default, [200/2], 00:16:01, bgp-65000, internal, tag 65000, segid: 303030 tunnelid: 0xa

```

Configure

On nexus 9000 eBGP peering is supported between a VTEP and External router.

For this peering to be possible a dedicated loopback on the VTEP needs to be configured.

Create dedicated loopback

Configure


	Command or Action	Purpose
Step 1	LEAF# configure terminal Enter configuration commands, one per line. End with CNTL/Z.	Enters configuration mode.
Step 2	LEAF(config)# interface lo 100	Create dedicated loopback.
Step 3	LEAF(config-if)# vrf member tenant-a Warning: Deleted all L3 config on interface loopback100	Attach tenant VRF to loopback.
Step 4	LEAF(config-if)# ip address 192.168.150.150/32	Assign an IP to loopback.

Advertise loopback on BGP

Created dedicated loopback needs to be advertised on BGP in order to have reachability to external router.

Configure

	Command or Action	Purpose
Step 1	LEAF# configure terminal Enter configuration commands, one per line. End with CNTL/Z.	Enters configuration mode.
Step 2	LEAF(config)# router bgp 65000	Enters BGP configuration.
Step 3	LEAF(config-router)# vrf tenant-a	Enters BGP VRF.
Step 4	LEAF(config-router-vrf)# address-family ipv4 unicast	Enters BGP VRF Address-family ipv4
Step 5	LEAF(config-router-vrf-af)# network 192.168.150.150/32	Advertise loopback IP.



Tip: BGP advertisement can also be done using route-maps and advertising it using redistribute direct.

Verify

Verify route is imported to BGP tenant VRF on LEAF VTEP

```
LEAF# show ip bgp 192.168.150.150 vrf tenant-a
BGP routing table information for VRF tenant-a, address family IPv4 Unicast
BGP routing table entry for 192.168.150.150/32, version 20
Paths: (1 available, best #1)
Flags: (0x80c0002) (high32 0x000020) on xmit-list, is not in urib, exported
vpn: version 25, (0x00000000100002) on xmit-list
```

```
Advertised path-id 1, VPN AF advertised path-id 1
```

Path type: local, path is valid, is best path, no labeled nexthop, is extd
Imported to 1 destination(s)
Imported paths list: tenant-b
AS-Path: NONE, path locally originated
0.0.0.0 (metric 0) from 0.0.0.0 (192.168.150.150)
Origin IGP, MED not set, localpref 100, weight 32768
Extcommunity: RT:1:1 RT:65000:303030

VRF advertise information:
Path-id 1 not advertised to any peer

VPN AF advertise information:
Path-id 1 not advertised to any peer

Verify route is imported to BGP L2VPN on LEAF VTEP

```
LEAF# sh bgp l2vpn evpn 192.168.150.150
BGP routing table information for VRF default, address family L2VPN EVPN
Route Distinguisher: 172.16.0.2:3 (L3VNI 303030)
BGP routing table entry for [5]:[0]:[0]:[32]:[192.168.150.150]/224, version 59
Paths: (1 available, best #1)
Flags: (0x000002) (high32 00000000) on xmit-list, is not in l2rib/evpn
```

```
Advertised path-id 1
Path type: local, path is valid, is best path, no labeled nexthop
Gateway IP: 0.0.0.0
AS-Path: NONE, path locally originated
172.16.0.2 (metric 0) from 0.0.0.0 (172.16.0.2)
Origin IGP, MED not set, localpref 100, weight 32768
Received label 303030
Extcommunity: RT:65000:303030 ENCAP:8 Router MAC:9c54.1651.5cd7
```

```
Path-id 1 advertised to peers:
10.101.11.1
```

Verify route is imported received on BL tenant VRF

```
BL# show ip route 192.168.150.150 vrf tenant-a
IP Route Table for VRF "tenant-a"
'*' denotes best ucast next-hop
'***' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>
```

```
192.168.150.150/32, ubest/mbest: 1/0
```

```
*via 172.16.0.2%default, [200/0], 00:01:13, bgp-65000, internal, tag 65000, segid: 303030 tunnelid: 0xa
```

Verify route is imported received on BL default VRF

Since External router is known via default VRF, VTEP loopback IP must be imported to default VRF on

BL.

```
BL(config-router-vrf-neighbor)# show ip route 192.168.150.150
```

```
IP Route Table for VRF "default"
```

```
'*' denotes best ucast next-hop
```

```
'**' denotes best mcast next-hop
```

```
'[x/y]' denotes [preference/metric]
```

```
'%<string>' in via output denotes VRF <string>
```

```
192.168.150.150/32, ubest/mbest: 1/0
```

```
*via 172.16.0.2, [200/0], 00:12:16, bgp-65000, internal, tag 65000, segid: 303030 tunnelid: 0xac100002
```

Configure eBGP peering on LEAF

Once dedicated loopback is created eBGP peering can be performed on BGP VRF.

Configure

	Command or Action	Purpose
Step 1	LEAF# configure terminal Enter configuration commands, one per line. End with CNTL/Z.	Enters configuration mode.
Step 2	LEAF(config)# router bgp 65000	Enters BGP configuration.
Step 3	LEAF(config-router)# vrf tenant-a	Enters BGP VRF.
Step 4	LEAF(config-router-vrf)# address-family ipv4 unicast	Enters BGP VRF Address-family ipv4
Step 5	LEAF(config-router-vrf-af)# neighbor 172.16.120.55 remote-as 64577	Creates eBGP neighbor.
Step 6	LEAF(config-router-vrf-neighbor)# update-source loopback 100	Use loopback 100 as update source.
Step 7	LEAF(config-router-vrf-neighbor)# address-family ipv4 unicast	Enters neighbor address-family ipv4 configuration.
Step 8	LEAF(config-router-vrf-neighbor-af)# ebgp-multihop 10	Assign ttl for packets for eBGP peering.



Warning: The use of update source of a dedicated loopback is mandatory.

Configure eBGP peering on External router (EXT-R)

Once dedicated loopback is created eBGP peering can be performed on BGP VRF.

Configure

	Command or Action	Purpose
Step 1	EXT-R# configure terminal Enter configuration commands, one per line. End with CNTL/Z.	Enters configuration mode.
Step 2	EXT-R(config)# router bgp 64577	Enters BGP configuration.
Step 3	EXT-R(config-router)# vrf blue	Enters BGP VRF.
Step 4	EXT-R(config-router-vrf)# address-family ipv4 unicast	Enters BGP VRF Address-family ipv4

Step 5	EXT-R(config-router-vrf-af)#neighbor 192.168.150.150 remote-as 65000	Creates eBGP neighbor.
Step 6	EXT-R(config-router-vrf-neighbor)# update-source loopback 10	Use loopback 100 as update source.
Step 7	EXT-R(config-router-vrf-neighbor)# address-family ipv4 unicast	Enters neighbor address-family ipv4 configuration.
Step 8	EXT-R(config-router-vrf-neighbor-af)# ebgp-multihop 10	Assign ttl for packets for eBGP peering.

Summary steps

LEAF

1. Configure terminal
2. interface lo 100
3. vrf member tenant-a
4. ip address 192.168.150.150/32
5. router bgp 65000
6. vrf tenant-a
7. address-family ipv4 unicast
8. network 192.168.150.150/32
9. neighbor 172.16.120.55
10. remote-as 64577
11. update-source loopback 100
12. address-family ipv4 unicast
13. ebgp-multihop 10

EXT-R

1. Configure terminal
2. router bgp 64577
3. vrf blue
4. address-family ipv4 unicast
5. neighbor 172.16.120.55
6. remote-as 64577
7. update-source loopback 100
8. address-family ipv4 unicast
9. ebgp-multihop 10

Verify

Verify eBGP neighbor on LEAF

```
LEAF# show ip bgp summary vrf tenant-a
BGP summary information for VRF tenant-a, address family IPv4 Unicast
BGP router identifier 192.168.150.150, local AS number 65000
BGP table version is 23, IPv4 Unicast config peers 1, capable peers 1
13 network entries and 14 paths using 2436 bytes of memory
BGP attribute entries [8/2816], BGP AS path entries [0/0]
BGP community entries [0/0], BGP clusterlist entries [3/12]
```

```
Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
172.16.120.55 4 64577 6 6 23 0 0 00:00:27 0
```

Verify eBGP neighbor on External Router

```
EXT-R(config-router-vrf-neighbor)# show ip bgp summary vrf blue
BGP summary information for VRF blue, address family IPv4 Unicast
BGP router identifier 172.16.120.55, local AS number 64577
BGP table version is 15, IPv4 Unicast config peers 1, capable peers 1
13 network entries and 13 paths using 3588 bytes of memory
BGP attribute entries [2/704], BGP AS path entries [1/6]
BGP community entries [0/0], BGP clusterlist entries [0/0]
```

```
Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
192.168.150.150 4 65000 13 6 15 0 0 00:00:48 13
```