



Configuring External VRF Connectivity and Route Leaking

This chapter contains the following sections:

- [Configuring External VRF Connectivity, on page 1](#)
- [Configuring Route Leaking, on page 15](#)

Configuring External VRF Connectivity

About External Layer-3 Connectivity for VXLAN BGP EVPN Fabrics

A VXLAN BGP EVPN fabric can be extended by using per-VRF IP routing to achieve external connectivity. The approach that is used for the Layer-3 extensions is commonly referred to as VRF Lite, while the functionality itself is more accurately defined as Inter-AS Option A or back-to-back VRF connectivity.

Guidelines and Limitations for External VRF Connectivity and Route Leaking

The following guidelines and limitations apply to external Layer 3 connectivity for VXLAN BGP EVPN fabrics:


- A physical Layer 3 interface (parent interface) can be used for external Layer 3 connectivity (that is, VRF default).
- The parent interface to multiple subinterfaces cannot be used for external Layer 3 connectivity (that is, Ethernet1/1 for a VRF default). You can use a subinterface instead.
- VTEPs do not support VXLAN-encapsulated traffic over parent interfaces if subinterfaces are configured, regardless of VRF participation.
- VTEPs do not support VXLAN-encapsulated traffic over subinterfaces, regardless of VRF participation or IEEE 802.1Q encapsulation.
- Mixing subinterfaces for VXLAN and non-VXLAN VLANs is not supported.

VXLAN BGP EVPN - VRF-lite brief

Some pointers are given below:

- The VXLAN BGP EVPN fabrics is depicted on the left in the following figure.
- Routes within the fabric are exchanged between all Edge-Devices (VTEPs) as well as Route-Reflectors; the control-plane used is MP-BGP with EVPN address-family.
- The Edge-Devices (VTEPs) acting as border nodes are configured to pass on prefixes to the external router (ER). This is achieved by exporting prefixes from MP-BGP EVPN to IPv4/IPv6 per-VRF peerings.
- Various routing protocols can be used for the per-VRF peering. While eBGP is the protocol of choice, IGP's like OSPF, IS-IS or EIGRP can be leveraged but require redistribution

Figure 1: External Layer-3 Connectivity - VRF-lite



Configuring VXLAN BGP EVPN with eBGP for VRF-lite

Configuring VRF for VXLAN Routing and External Connectivity using BGP

Configure the VRF on the border node.

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code>	Enter global configuration mode.
Step 2	<code>vrf context vrf-name</code>	Configure the VRF.
Step 3	<code>vni number</code>	Specify the VNI. The VNI associated with the VRF is often referred to as a Layer 3 VNI, L3VNI, or L3VPN. The L3VNI is configured as the common identifier across the participating VTEPs.
Step 4	<code>rd {auto rd}</code>	Specify the VRF's route distinguisher (RD). The RD uniquely identifies a VTEP within an L3VNI. If you enter an RD, the following formats are supported: ASN2:NN, ASN4:NN, or IPV4:NN.

	Command or Action	Purpose
Step 5	address-family {ipv4 ipv6} unicast	Configure the IPv4 or IPv6 unicast address family.
Step 6	route-target both {auto rt}	Configure the route target (RT) for import and export of IPv4 prefixes. The RT is used for a per-VRF prefix import/export policy. If you enter an RT, the following formats are supported: ASN2:NN, ASN4:NN, or IPV4:NN.
Step 7	route-target both {auto rt} evpn	Configure the route target (RT) for import and export of IPv4 prefixes. The RT is used for a per-VRF prefix import/export policy. If you enter an RT, the following formats are supported: ASN2:NN, ASN4:NN, or IPV4:NN.
Step 8	Repeat Step 1 through Step 7 for every L3VNI.	

Configuring the L3VNI's Fabric Facing VLAN and SVI on the Border Node

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enter configuration mode.
Step 2	vlan <i>number</i>	Specify the VLAN id that is used for the L3VNI.
Step 3	vn-segment <i>number</i>	Map the L3VNI to the VLAN for VXLAN EVPN routing.
Step 4	interface <i>vlan-number</i>	Specify the SVI (Switch Virtual Interface) for VXLAN EVPN routing.
Step 5	mtu <i>value</i>	Specify the MTU for the L3VNI.
Step 6	vrf member <i>vrf-name</i>	Map the SVI to the matching VRF context.
Step 7	ip forward	Enable IPv4 forwarding for the L3VNI.
Step 8	no ip redirects	Disable ICMP redirects
Step 9	ipv6 <i>ip-address</i>	Enable IPv6 forwarding for the L3VNI.
Step 10	no ipv6 redirects	Disable ICMPv6 redirects.
Step 11	Repeat Step 2 through Step 10 for every L3VNI.	

Configuring the VTEP on the Border Node

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code>	Enter global configuration mode.
Step 2	<code>interface nve1</code>	Configure the NVE interface.
Step 3	<code>member vni vni associate-vrf</code>	Add Layer-3 VNIs, one per tenant VRF, to the overlay.
Step 4		Repeat Step 3 for every L3VNI.

Configuring the BGP VRF Instance on the Border Node for IPv4 per-VRF Peering

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code>	Enter global configuration mode.
Step 2	<code>router bgp <i>autonomous-system-number</i></code>	Configure BGP. The range of the <i>autonomous-system-number</i> is from 1 to 4294967295.
Step 3	<code>vrf <i>vrf-name</i></code>	Specify the VRF.
Step 4	<code>address-family ipv4 unicast</code>	Configure address family for IPv4.
Step 5	<code>advertise l2vpn evpn</code>	Enable the advertisement of EVPN routes within IPv4 address-family.
Step 6	<code>maximum-paths ibgp <i>number</i></code>	Enabling equal cost multipathing (ECMP) for iBGP prefixes. The range for <i>number</i> is 1 to 64. The default is 1.
Step 7	<code>maximum-paths <i>number</i></code>	Enabling equal cost multipathing (ECMP) for eBGP prefixes.
Step 8	<code>neighbor <i>address</i> remote-as <i>number</i></code>	Define eBGP neighbor IPv4 address and remote Autonomous-System (AS) number.
Step 9	<code>update-source <i>type/id</i></code>	Define interface for eBGP peering.
Step 10	<code>address-family ipv4 unicast</code>	Activate the IPv4 address family for IPv4 prefix exchange.
Step 11	Repeat Step 3 through Step 10 for every L3VNI that requires external connectivity for IPv4.	

Configuring the BGP VRF Instance on the Border Node for IPv6 per-VRF Peering

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system-number</i>	Configure BGP.
Step 3	vrf <i>vrf-name</i>	Specify the VRF.
Step 4	address-family ipv6 unicast	Configure address family for IPv4.
Step 5	advertise l2vpn evpn	Enable the advertisement of EVPN routes within IPv6 address-family.
Step 6	maximum-paths ibgp <i>number</i>	Enabling equal cost multipathing (ECMP) for iBGP prefixes.
Step 7	maximum-paths <i>number</i>	Enabling equal cost multipathing (ECMP) for eBGP prefixes.
Step 8	neighbor address remote-as <i>number</i>	Define eBGP neighbor IPv6 address and remote Autonomous-System (AS) number.
Step 9	update-source <i>type/id</i>	Define interface for eBGP peering.
Step 10	address-family ipv6 unicast	Configure address family for IPv6.
Step 11	Repeat Step 3 Through Step 10 for every L3VNI that requires external connectivity for IPv6.	

Configuring the Sub-Interface Instance on the Border Node for Per-VRF Peering - Version 1

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.
Step 2	interface <i>type/id</i>	Configure parent interface.
Step 3	no switchport	Disable Layer-2 switching mode on interface.
Step 4	no shutdown	Bring up parent interface.
Step 5	exit	Exit interface configuration mode.
Step 6	interface <i>type/id</i>	Define the Sub-Interface instance.
Step 7	encapsulation dot1q <i>number</i>	Configure the VLAN ID for the sub-interface. The <i>number</i> argument can have a value from 1 to 3967.

	Command or Action	Purpose
Step 8	vrf member <i>vrf-name</i>	Map the Sub-Interface to the matching VRF context.
Step 9	ip address <i>address</i>	Configure the Sub-Interfaces IP address.
Step 10	no shutdown	Bring up Sub-Interface.
Step 11	Repeat Step 5 through Step 9 for every per-VRF peering.	

VXLAN BGP EVPN - Default-Route, Route Filtering on External Connectivity

About Configuring Default Routing for External Connectivity

For default-route advertisement into a VXLAN BGP EVPN fabric, we have to ensure that the default-route advertised into the fabric is at the same time not advertised outside of the fabric. For this case, it is necessary to have route filtering in place that prevents this eventuality.

Configuring the Default Route in the Border Nodes VRF

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.
Step 2	vrf context <i>vrf-name</i>	Configure the VRF.
Step 3	ip route 0.0.0.0/0 <i>next-hop</i>	Configure the IPv4 default-route.
Step 4	ipv6 route 0::/0 <i>next-hop</i>	Configure the IPv6 default-route.

Configuring the BGP VRF Instance on the Border Node for IPv4/IPv6 Default-Route Advertisement

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.
Step 2	router bgp <i>autonomous-system-number</i>	Configure BGP.
Step 3	vrf <i>vrf-name</i>	Specify the VRF.
Step 4	address-family ipv4 unicast	Configure the IPv4 Unicast address-family. Required for IPv6 over VXLAN with IPv4 underlay.
Step 5	network 0.0.0.0/0	Creating IPv4 default-route network statement.
Step 6	address-family ipv6 unicast	Configure the IPv6 unicast address-family.

	Command or Action	Purpose
Step 7	network 0::/0	Creating IPv6 default-route network statement.
Step 8	neighbor addressremote-as number	Define eBGP neighbor IPv4 address and remote Autonomous-System (AS) number.
Step 9	update-source type/id	Define interface for eBGP peering
Step 10	address-family {ipv4 ipv6} unicast	Activate the IPv4 or IPv6 address family for IPv4/IPv6 prefix exchange.
Step 11	route-map name out	Attach route-map for egress route filtering.
Step 12	Repeat Step 3 through Step 11 for every L3VNI that requires external connectivity with default-route filtering.	

Configuring Route Filtering for IPv4 Default-Route Advertisement

You can configure route filtering for IPv4 default-route advertisement.

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.
Step 2	ip prefix-list name seq 5 permit 0.0.0.0/0	Configure IPv4 prefix-list for default-route filtering.
Step 3	route-map name deny 10	Create route-map with leading deny statement to prevent the default-route of being advertised via External Connectivity.
Step 4	match ip address prefix-list name	Match against the IPv4 prefix-list that contains the default-route.
Step 5	route-map name permit 1000	Create route-map with trailing allow statement to advertise non-matching routes via External Connectivity.

Configuring Route Filtering for IPv6 Default-Route Advertisement

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.
Step 2	ipv6 prefix-list name seq 5 permit 0::/0	Configure IPv6 prefix-list for default-route filtering.

	Command or Action	Purpose
Step 3	route-map <i>name</i> deny 10	Create route-map with leading deny statement to prevent the default-route of being advertised via External Connectivity.
Step 4	match ipv6 address prefix-list <i>name</i>	Match against the IPv6 prefix-list that contains the default-route.
Step 5	route-map <i>name</i> permit 1000	Create route-map with trailing allow statement to advertise non-matching routes via External Connectivity.

About Configuring Default-Route Distribution and Host-Route Filter

Per-default, a VXLAN BGP EVPN fabric always advertises all known routes via the External Connectivity. As not in all circumstances it is beneficial to advertise IPv4 /32 or IPv6 /128 Host-Routes, a respective route filtering approach can become necessary.

Configuring the BGP VRF Instance on the Border Node for IPv4/IPv6 Host-Route Filtering

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.
Step 2	router bgp <i>autonomous-system-number</i>	Configure BGP.
Step 3	vrf <i>vrf-name</i>	Specify the VRF.
Step 4	neighbor <i>address</i> remote-as <i>number</i>	Define eBGP neighbor IPv4/IPv6 address and remote Autonomous-System (AS) number.
Step 5	update-source <i>type/id</i>	Define interface for eBGP peering.
Step 6	address-family {<i>ipv4</i> <i>ipv6</i>} unicast	Activate the IPv4 or IPv6 address family for IPv4/IPv6 prefix exchange.
Step 7	route-map <i>name</i> out	Attach route-map for egress route filtering.
Step 8	Repeat Step 3 through Step 7 for every L3VNI that requires external connectivity with host-route filtering.	

Configuring Route Filtering for IPv4 Host-Route Advertisement

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.

	Command or Action	Purpose
Step 2	ip prefix-list <i>name</i> seq 5 permit 0.0.0.0/0 eq 32	Configure IPv4 prefix-list for host-route filtering.
Step 3	route-map <i>name</i> deny 10	Create route-map with leading deny statement to prevent the default-route of being advertised via External Connectivity.
Step 4	match ip address prefix-list <i>name</i>	Match against the IPv4 prefix-list that contains the host-route.
Step 5	route-map <i>name</i> permit 1000	Create route-map with trailing allow statement to advertise non-matching routes via external connectivity.

Configuring Route Filtering for IPv6 Host-Route Advertisement

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.
Step 2	ipv6 prefix-list <i>name</i> seq 5 permit 0::/0 eq 128	Configure IPv4 prefix-list for host-route filtering.
Step 3	route-map <i>name</i> deny 10	Create route-map with leading deny statement to prevent the default-route of being advertised via External Connectivity.
Step 4	match ipv6 address prefix-list <i>name</i>	Match against the IPv4 prefix-list that contains the host-route.
Step 5	route-map <i>name</i> permit 1000	Create route-map with trailing allow statement to advertise non-matching routes via External Connectivity.

Example - Configuring VXLAN BGP EVPN with eBGP for VRF-lite

An example of external connectivity from VXLAN BGP EVPN to an external router using VRF-lite.

Configuring VXLAN BGP EVPN Border Node

The VXLAN BGP EVPN Border Node acts as neighbor device to the External Router. The VRF Name is purely localized and can be different to the VRF Name on the External Router, only significance is the L3VNI must be consistent across the VXLAN BGP EVPN fabric. For the ease of reading, the VRF and interface enumeration will be consistently used.

The configuration examples represents a IPv4 and IPv6 dual-stack approach; IPv4 or IPv6 can be substituted of each other.

```
vrf context myvrf_50001
  vni 50001
  rd auto
```

```

    address-family ipv4 unicast
      route-target both auto
      route-target both auto evpn
    address-family ipv6 unicast
      route-target both auto
      route-target both auto evpn
  !
vlan 2000
  vn-segment 50001
!
interface Vlan2000
  no shutdown
  mtu 9216
  vrf member myvrf_50001
  no ip redirects
  ip forward
  ipv6 address use-link-local-only
  no ipv6 redirects
!
interface nve1
  no shutdown
  host-reachability protocol bgp
  source-interface loopback1
  member vni 50001 associate-vrf
!
router bgp 65002
  vrf myvrf_50001
    router-id 10.2.0.6
    address-family ipv4 unicast
      advertise l2vpn evpn
      maximum-paths ibgp 2
      maximum-paths 2
    address-family ipv6 unicast
      advertise l2vpn evpn
      maximum-paths ibgp 2
      maximum-paths 2
    neighbor 10.31.95.95
      remote-as 65099
      address-family ipv4 unicast
    neighbor 2001::95/64
      remote-as 65099
      address-family ipv4 unicast
  !
interface Ethernet1/3
  no switchport
  no shutdown
interface Ethernet1/3.2
  encapsulation dot1q 2
  vrf member myvrf_50001
  ip address 10.31.95.31/24
  ipv6 address 2001::31/64
  no shutdown

```

Configuring Default-Route, Route Filtering on External Connectivity

The VXLAN BGP EVPN Border Node has the ability to advertise IPv4 and IPv6 default-route within the fabric. In cases where it is not beneficial to advertise the Host Routes from the VXLAN BGP EVPN fabric to the External Router, these IPv4 /32 and IPv6 /128 can be filtered at the External Connectivity peering configuration.

```

ip prefix-list default-route seq 5 permit 0.0.0.0/0 le 1
ipv6 prefix-list default-route-v6 seq 5 permit 0::/0
!
ip prefix-list host-route seq 5 permit 0.0.0.0/0 eq 32

```

```

ipv6 prefix-list host-route-v6 seq 5 permit 0::/0 eq 128
!
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-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
!
vrf context myvrf_50001
  ip route 0.0.0.0/0 10.31.95.95
  ipv6 route 0::/0 2001::95/64
!
router bgp 65002
  vrf myvrf_50001
    address-family ipv4 unicast
      network 0.0.0.0/0
    address-family ipv6 unicast
      network 0::/0

  neighbor 10.31.95.95
    remote-as 65099
    address-family ipv4 unicast
      route-map extcon-rmap-filter out
  neighbor 2001::95/64
    remote-as 65099
    address-family ipv4 unicast
      route-map extcon-rmap-filter-v6 out

```

Configuring External Router

The External Router performs as a neighbor device to the VXLAN BGP EVPN border node. The VRF Name is purely localized and can be different to the VRF Name on the VXLAN BGP EVPN Fabric. For the ease of reading, the VRF and interface enumeration will be consistently used.

The configuration examples represents a IPv4 and IPv6 dual-stack approach; IPv4 or IPv6 can be substituted of each other.

```

vrf context myvrf_50001
!
router bgp 65099
  vrf myvrf_50001
    address-family ipv4 unicast
      maximum-paths 2
    address-family ipv6 unicast
      maximum-paths 2
  neighbor 10.31.95.31
    remote-as 65002
    address-family ipv4 unicast
  neighbor 2001::31/64
    remote-as 65002
    address-family ipv4 unicast
!
interface Ethernet1/3
  no switchport
  no shutdown
interface Ethernet1/3.2
  encapsulation dot1q 2

```

```
vrf member myvrf_50001
ip address 10.31.95.95/24
Ipv6 address 2001::95/64
no shutdown
```

Configuring VXLAN BGP EVPN with OSPF for VRF-lite

Configuring VRF for VXLAN Routing and External Connectivity using OSPF

Configure the BGP VRF instance on the border node for OSPF per-VRF peering.

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system-number</i>	Configure BGP.
Step 3	vrf <i>vrf-name</i>	Specify the VRF.
Step 4	address-family ipv4 unicast	Configure the IPv4 address family.
Step 5	advertise l2vpn evpn	Enable the advertisement of EVPN routes within the address family.
Step 6	maximum-paths ibgp <i>number</i>	Enabling equal-cost multipathing (ECMP) for iBGP prefixes.
Step 7	redistribute ospf <i>name</i> route-map <i>name</i>	Define redistribution from OSPF into BGP.
Step 8	Repeat Step 3 through Step 7 for every per-VRF peering.	

Configuring the Route-Map for BGP to OSPF Redistribution

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	route-map <i>name</i> permit 10	Create route-map for BGP to OSPF redistribution
Step 3	match route-type internal	Redistribution route-map must allow the matching of BGP internal route-types if iBGP is used in the VXLAN BGP EVPN fabric.

Configuring the OSPF on the Border Node for Per-VRF Peering

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code>	Enter global configuration mode.
Step 2	<code>router ospf instance</code>	Configure OSPF.
Step 3	<code>vrf vrf-name</code>	Specify the VRF.
Step 4	<code>redistribute bgp autonomous-system-number</code> <code>route-map name</code>	Define redistribution from BGP to OSPF.
Step 5	Repeat Step 3 through Step 4 for every per-VRF peering.	

Configuring the Sub-Interface Instance on the Border Node for Per-VRF Peering - Version 2

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code>	Enters global configuration mode.
Step 2	<code>interface type/id</code>	Configure parent interface.
Step 3	<code>no switchport</code>	Disable Layer-2 switching mode on interface.
Step 4	<code>no shutdown</code>	Bring up parent interface.
Step 5	<code>exit</code>	Exit interface configuration mode.
Step 6	<code>interface type/id</code>	Define the Sub-Interface instance.
Step 7	<code>encapsulation dot1q number</code>	Configure the VLAN ID for the sub-interface. The range is from 2 to 4093.
Step 8	<code>vrf member vrf-name</code>	Map the Sub-Interface to the matching VRF context.
Step 9	<code>ip address address</code>	Configure the Sub-Interfaces IP address.
Step 10	<code>ip ospf network point-to-point</code>	Define OSPF network-type for sub-interface.
Step 11	<code>ip router ospf name area area-id</code>	Configure the OSPF instance.
Step 12	<code>no shutdown</code>	Bring up Sub-Interface.
Step 13	Repeat Step 5 through Step 12 for every per-VRF peering.	

Example - Configuration VXLAN BGP EVPN with OSPF for VRF-lite

An example of external connectivity from VXLAN BGP EVPN to an External Router using VRF-lite.

Configuring VXLAN BGP EVPN Border Node with OSPF

The VXLAN BGP EVPN Border Node acts as neighbor device to the External Router. The VRF Name is purely localized and can be different to the VRF Name on the External Router, only significance is the L3VNI must be consistent across the VXLAN BGP EVPN fabric. For the ease of reading, the VRF and interface enumeration will be consistently used.

The configuration examples represents a IPv4 approach with OSPFv2.

```

route-map extcon-rmap-BGP-to-OSPF permit 10
  match route-type internal
route-map extcon-rmap-OSPF-to-BGP permit 10
!
vrf context myvrf_50001
  vni 50001
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
!
vlan 2000
  vn-segment 50001
!
interface Vlan2000
  no shutdown
  mtu 9216
  vrf member myvrf_50001
  no ip redirects
  ip forward
!
interface nve1
  no shutdown
  host-reachability protocol bgp
  source-interface loopback1
  member vni 50001 associate-vrf
!
router bgp 65002
  vrf myvrf_50001
  router-id 10.2.0.6
  address-family ipv4 unicast
    advertise l2vpn evpn
    maximum-paths ibgp 2
    maximum-paths 2
  redistribute ospf EXT route-map extcon-rmap-OSPF-to-BGP
!
router ospf EXT
  vrf myvrf_50001
  redistribute bgp 65002 route-map extcon-rmap-BGP-to-OSPF
!
interface Ethernet1/3
  no switchport
  no shutdown
interface Ethernet1/3.2
  encapsulation dot1q 2
  vrf member myvrf_50001
  ip address 10.31.95.31/24
  ip ospf network point-to-point
  ip router ospf EXT area 0.0.0.0
  no shutdown

```

Configuring Route Leaking

About Centralized VRF Route-Leaking for VXLAN BGP EVPN Fabrics

VXLAN BGP EVPN uses MP-BGP and its route-policy concept to import and export prefixes. The ability of this very extensive route-policy model allows to leak routes from one VRF to another VRF and vice-versa; any combination of custom VRF or VRF default can be used. VRF route-leaking is a switch-local function at specific to a location in the network, the location where the cross-VRF route-target import/export configuration takes place (leaking point). The forwarding between the different VRFs follows the control-plane, the location of where the configuration for the route-leaking is performed - hence Centralized VRF route-leaking. With the addition of VXLAN BGP EVPN, the leaking point requires to advertise the cross-VRF imported/exported route and advertise them towards the remote VTEPs or External Routers.

The advantage of Centralized VRF route-leaking is that only the VTEP acting as leaking point requires the special capabilities needed, while all other VTEPs in the network are neutral to this function.

Guidelines and Limitations for Centralized VRF Route-Leaking

The following are the guidelines and limitations for Centralized VRF Route-Leaking:


- Each prefix must be imported into each VRF for full cross-VRF reachability.
- The **feature bgp** command is required for the **export vrf default** command.
- If a VTEP has a less specific local prefix in its VRF, the VTEP might not be able to reach a more specific prefix in a different VRF.
- VXLAN routing in hardware and packet reencapsulation at VTEP is required for Centralized VRF Route-Leaking with BGP EVPN.

Centralized VRF Route-Leaking Brief - Specific Prefixes Between Custom VRF

Some pointers are given below:

- The Centralized VRF route-leaking for VXLAN BGP EVPN fabrics is depicted within Figure 2.
- BGP EVPN prefixes are cross-VRF leaked by exporting them from VRF Blue with an import into VRF Red and vice-versa. The Centralized VRF route-leaking is performed on the centralized Routing-Block (RBL) and could be any or multiple VTEPs.
- Configured less specific prefixes (aggregates) are advertised from the Routing-Block to the remaining VTEPs in the respective destination VRF.
- BGP EVPN does not export prefixes that were previously imported to prevent the occurrence of routing loops.

Figure 2: Centralized VRF Route-Leaking - Specific Prefixes with Custom VRF



Configuring Centralized VRF Route-Leaking - Specific Prefixes between Custom VRF

Configuring VRF Context on the Routing-Block VTEP

This procedure applies equally to IPv6.

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code>	Enter global configuration mode.
Step 2	<code>vrf context vrf-name</code>	Configure the VRF.
Step 3	<code>vni number</code>	Specify the VNI. The VNI associated with the VRF is often referred to as Layer 3 VNI, L3VNI, or L3VPN. The L3VNI is configured as a common identifier across the participating VTEPs.
Step 4	<code>rd auto</code>	Specify the VRF's route distinguisher (RD). The RD uniquely identifies a VTEP within an L3VNI.

	Command or Action	Purpose
Step 5	address-family ipv4 unicast	Configure the IPv4 unicast address family.
Step 6	route-target both {auto rt}	Configure the route target (RT) for import and export of IPv4 prefixes. The RT is used for a per-VRF prefix import/export policy. If you enter an RT, the following formats are supported: ASN2:NN, ASN4:NN, or IPV4:NN.
Step 7	route-target both {auto rt} evpn	Configure the route target (RT) for import and export of IPv4 prefixes. The RT is used for a per-VRF prefix import/export policy. If you enter an RT, the following formats are supported: ASN2:NN, ASN4:NN, or IPV4:NN.
Step 8	route-target import rt-from-different-vrf	Configure the RT for importing IPv4 prefixes from the leaked-from VRF. The following formats are supported: ASN2:NN, ASN4:NN, or IPV4:NN.
Step 9	route-target import rt-from-different-vrf evpn	Configure the RT for importing IPv4 prefixes from the leaked-from VRF. The following formats are supported: ASN2:NN, ASN4:NN, or IPV4:NN.

Configuring the BGP VRF instance on the Routing-Block

This procedure applies equally to IPv6.

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.
Step 2	router bgp <i>autonomous-system number</i>	Configure BGP.
Step 3	vrf <i>vrf-name</i>	Specify the VRF.
Step 4	address-family ipv4 unicast	Configure address family for IPv4
Step 5	advertise l2vpn evpn	Enable the advertisement of EVPN routes within IPv4 address-family.
Step 6	aggregate-address <i>prefix/mask</i>	Create less specific prefix aggregate into the destination VRF.
Step 7	maximum-paths ibgp <i>number</i>	Enabling equal cost multipathing (ECMP) for iBGP prefixes.
Step 8	maximum-paths <i>number</i>	Enabling equal cost multipathing (ECMP) for eBGP prefixes

Example - Configuration Centralized VRF Route-Leaking - Specific Prefixes Between Custom VRF

Configuring VXLAN BGP EVPN Routing-Block

The VXLAN BGP EVPN Routing-Block acts as centralized route-leaking point. The leaking configuration is localized such that control-plane leaking and data-path forwarding follow the same path. Most significantly is the VRF configuration of the Routing-Block and the advertisement of the less specific prefixes (aggregates) into the respective destination VRFs.

```
vrf context Blue
  vni 51010
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
    route-target import 65002:51020
    route-target import 65002:51020 evpn
!
vlan 2110
  vn-segment 51010
!
interface Vlan2110
  no shutdown
  mtu 9216
  vrf member Blue
  no ip redirects
  ip forward
!
vrf context Red
  vni 51020
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
    route-target import 65002:51010
    route-target import 65002:51010 evpn
!
vlan 2120
  vn-segment 51020
!
interface Vlan2120
  no shutdown
  mtu 9216
  vrf member Blue
  no ip redirects
  ip forward
!
interface nve1
  no shutdown
  host-reachability protocol bgp
  source-interface loopback1
  member vni 51010 associate-vrf
  member vni 51020 associate-vrf
!
router bgp 65002
  vrf Blue
    address-family ipv4 unicast
      advertise l2vpn evpn
      aggregate-address 10.20.0.0/16
      maximum-paths ibgp 2
      Maximum-paths 2
  vrf Red
    address-family ipv4 unicast
```

```

advertise l2vpn evpn
aggregate-address 10.10.0.0/16
maximum-paths ibgp 2
Maximum-paths 2


```

Centralized VRF Route-Leaking Brief - Shared Internet with Custom VRF

Some pointers follow:

- The Shared Internet with VRF route-leaking for VXLAN BGP EVPN fabrics is depicted in the following figure.
- The default-route is made exported from the Shared Internet VRF and re-advertisement within VRF Blue and VRF Red on the Border Node.
- Ensure the default-route in VRF Blue and VRF Red is not leaked to the Shared Internet VRF.
- The less specific prefixes for VRF Blue and VRF Red are exported for the Shared Internet VRF and re-advertised as necessary.
- Configured less specific prefixes (aggregates) that are advertised from the Border Node to the remaining VTEPs to the destination VRF (Blue or Red).
- BGP EVPN does not export prefixes that were previously imported to prevent the occurrence of routing loops.

Figure 3: Centralized VRF Route-Leaking - Shared Internet with Custom VRF



Configuring Centralized VRF Route-Leaking - Shared Internet with Custom VRF

Configuring Internet VRF on Border Node

This procedure applies equally to IPv6.

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vrf context <i>vrf-name</i>	Configure the VRF.
Step 3	vni <i>number</i>	Specify the VNI. The VNI associated with the VRF is often referred to as Layer 3 VNI, L3VNI, or L3VPN. The L3VNI is configured as a common identifier across the participating VTEPs.
Step 4	ip route 0.0.0.0/0 <i>next-hop</i>	Configure the default route in the shared internet VRF to the external router.
Step 5	rd auto	Specify the VRF's route distinguisher (RD). The RD uniquely identifies a VTEP within an L3VNI.
Step 6	address-family ipv4 unicast	Configure the IPv4 unicast address family. This configuration is required for IPv4 over VXLAN with IPv4 underlay.
Step 7	route-target both { <i>auto</i> <i>rt</i> }	Configure the route target (RT) for the import and export of EVPN and IPv4 prefixes. If you enter an RT, the following formats are supported: ASN2:NN, ASN4:NN, or IPV4:NN.
Step 8	route-target both <i>shared-vrf-rt evpn</i>	Configure a special route target (RT) for the import and export of the shared IPv4 prefixes. An additional import/export map for further qualification is supported.

Configuring Shared Internet BGP Instance on the Border Node

This procedure applies equally to IPv6.

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.
Step 2	router bgp <i>autonomous-system number</i>	Configure BGP.

	Command or Action	Purpose
Step 3	<code>vrf vrf-name</code>	Specify the VRF.
Step 4	<code>address-family ipv4 unicast</code>	Configure address family for IPv4
Step 5	<code>advertise l2vpn evpn</code>	Enable the advertisement of EVPN routes within IPv4 address-family.
Step 6	<code>aggregate-address prefix/mask</code>	Create less specific prefix aggregate into the destination VRF.
Step 7	<code>maximum-paths ibgp number</code>	Enabling equal cost multipathing (ECMP) for iBGP prefixes.
Step 8	<code>maximum-paths number</code>	Enabling equal cost multipathing (ECMP) for eBGP prefixes.

Configuring Custom VRF on Border Node

This procedure applies equally to IPv6

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code>	Enters global configuration mode.
Step 2	<code>ip prefix-list name seq 5 permit 0.0.0.0/0</code>	Configure IPv4 prefix-list for default-route filtering.
Step 3	<code>route-map name deny 10</code>	Create route-map with leading deny statement to prevent the default-route of being leaked.
Step 4	<code>match ip address prefix-list name</code>	Match against the IPv4 prefix-list that contains the default-route.
Step 5	<code>route-map name permit 20</code>	Create route-map with trailing allow statement to advertise non-matching routes via route-leaking.

Configuring Custom VRF Context on the Border Node - 1

This procedure applies equally to IPv6.

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code>	Enters global configuration mode.
Step 2	<code>vrf context vrf-name</code>	Configure the VRF.

	Command or Action	Purpose
Step 3	<code>vni number</code>	Specify the VNI. The VNI associated with the VRF is often referred to as Layer 3 VNI, L3VNI, or L3VPN. The L3VNI is configured as the common identifier across the participating VTEPs.
Step 4	<code>rd auto</code>	Specify the VRF's route distinguisher (RD). The RD uniquely identifies a VTEP within an L3VNI.
Step 5	<code>ip route 0.0.0.0/0 Null0</code>	Configure default-route in common VRF to attract traffic towards Border Node with Shared Internet VRF.
Step 6	<code>address-family ipv4 unicast</code>	Configure the IPv4 address family. This configuration is required for IPv4 over VXLAN with IPv4 underlay.
Step 7	<code>route-target both {auto rt}</code>	Configure the route target (RT) for the import and export of IPv4 prefixes within the IPv4 address family. The RT is used for a per-VRF prefix import/export policy. If you enter an RT, the following formats are supported: ASN2:NN, ASN4:NN, or IPV4:NN.
Step 8	<code>route-target both {auto rt} evpn</code>	Configure the route target (RT) for the import and export of IPv4 prefixes within the IPv4 address family. The RT is used for a per-VRF prefix import/export policy. If you enter an RT, the following formats are supported: ASN2:NN, ASN4:NN, or IPV4:NN.
Step 9	<code>import map name</code>	Apply a route-map on routes being imported into this routing table.

Configuring Custom VRF Instance in BGP on the Border Node

This procedure applies equally to IPv6.

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code>	Enters global configuration mode.
Step 2	<code>router bgp autonomous-system-number</code>	Configure BGP.
Step 3	<code>vrf vrf-name</code>	Specify the VRF.
Step 4	<code>address-family ipv4 unicast</code>	Configure address family for IPv4.

	Command or Action	Purpose
Step 5	<code>advertise l2vpn evpn</code>	Enable the advertisement of EVPN routes within IPv4 address-family.
Step 6	<code>network 0.0.0.0/0</code>	Creating IPv4 default-route network statement.
Step 7	<code>maximum-paths ibgp number</code>	Enabling equal cost multipathing (ECMP) for iBGP prefixes.
Step 8	<code>maximum-paths number</code>	Enabling equal cost multipathing (ECMP) for eBGP prefixes.

Example - Configuration Centralized VRF Route-Leaking - Shared Internet with Custom VRF

An example of Centralized VRF route-leaking with Shared Internet VRF

Configuring VXLAN BGP EVPN Border Node for Shared Internet VRF

The VXLAN BGP EVPN Border Node provides a centralized Shared Internet VRF. The leaking configuration is localized such that control-plane leaking and data-path forwarding following the same path. Most significantly is the VRF configuration of the Border Node and the advertisement of the default-route and less specific prefixes (aggregates) into the respective destination VRFs.

```
vrf context Shared
  vni 51099
  ip route 0.0.0.0/0 10.9.9.1
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
    route-target both 99:99
    route-target both 99:99 evpn
  !
vlan 2199
  vn-segment 51099
  !
interface Vlan2199
  no shutdown
  mtu 9216
  vrf member Shared
  no ip redirects
  ip forward
  !
ip prefix-list PL_DENY_EXPORT seq 5 permit 0.0.0.0/0
  !
route-map RM_DENY_IMPORT deny 10
  match ip address prefix-list PL_DENY_EXPORT
route-map RM_DENY_IMPORT permit 20
  !
vrf context Blue
  vni 51010
  ip route 0.0.0.0/0 Null0
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
    route-target both 99:99
    route-target both 99:99 evpn
  import map RM_DENY_IMPORT
```

Example - Configuration Centralized VRF Route-Leaking - Shared Internet with Custom VRF

```

!
vlan 2110
  vn-segment 51010
!
interface Vlan2110
  no shutdown
  mtu 9216
  vrf member Blue
  no ip redirects
  ip forward
!
vrf context Red
  vni 51020
  ip route 0.0.0.0/0 Null0
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
    route-target both 99:99
    route-target both 99:99 evpn
  import map RM_DENY_IMPORT
!
vlan 2120
  vn-segment 51020
!
interface Vlan2120
  no shutdown
  mtu 9216
  vrf member Blue
  no ip redirects
  ip forward
!
interface nve1
  no shutdown
  host-reachability protocol bgp
  source-interface loopback1
  member vni 51099 associate-vrf
  member vni 51010 associate-vrf
  member vni 51020 associate-vrf
!
router bgp 65002
  vrf Shared
    address-family ipv4 unicast
      advertise l2vpn evpn
      aggregate-address 10.10.0.0/16
      aggregate-address 10.20.0.0/16
      maximum-paths ibgp 2
      maximum-paths 2
    vrf Blue
      address-family ipv4 unicast
        advertise l2vpn evpn
        network 0.0.0.0/0
        maximum-paths ibgp 2
        maximum-paths 2
    vrf Red
      address-family ipv4 unicast
        advertise l2vpn evpn
        network 0.0.0.0/0
        maximum-paths ibgp 2
        maximum-paths 2


```


Centralized VRF Route-Leaking Brief - Shared Internet with VRF Default

Some pointers are given below:

- The Shared Internet with VRF route-leaking for VXLAN BGP EVPN fabrics is depicted within Figure 4.
- The default-route is made exported from VRF default and re-advertisement within VRF Blue and VRF Red on the Border Node.
- Ensure the default-route in VRF Blue and VRF Red is not leaked to the Shared Internet VRF
- The less specific prefixes for VRF Blue and VRF Red are exported to VRF default and re-advertised as necessary.
- Configured less specific prefixes (aggregates) that are advertised from the Border Node to the remaining VTEPs to the destination VRF (Blue or Red).
- BGP EVPN does not export prefixes that were previously imported to prevent the occurrence of routing loops.

Figure 4: Centralized VRF Route-Leaking - Shared Internet with VRF Default



Configuring Centralized VRF Route-Leaking - Shared Internet with VRF Default

Configuring VRF Default on Border Node

This procedure applies equally to IPv6.

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.
Step 2	ip route 0.0.0.0/0 next-hop	Configure default-route in VRF default to external router (example)

Configuring BGP Instance for VRF Default on the Border Node

This procedure applies equally to IPv6.

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.
Step 2	router bgp autonomous-system number	Configure BGP.
Step 3	address-family ipv4 unicast	Configure address family for IPv4.
Step 4	aggregate-address prefix/mask	Create less specific prefix aggregate in VRF default.
Step 5	maximum-paths number	Enabling equal cost multipathing (ECMP) for eBGP prefixes.

Configuring Custom VRF on Border Node

This procedure applies equally to IPv6

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.
Step 2	ip prefix-list name seq 5 permit 0.0.0.0/0	Configure IPv4 prefix-list for default-route filtering.
Step 3	route-map name deny 10	Create route-map with leading deny statement to prevent the default-route of being leaked.
Step 4	match ip address prefix-list name	Match against the IPv4 prefix-list that contains the default-route.

	Command or Action	Purpose
Step 5	<code>route-map <i>name</i> permit 20</code>	Create route-map with trailing allow statement to advertise non-matching routes via route-leaking.

Configuring Filter for Permitted Prefixes from VRF Default on the Border Node

This procedure applies equally to IPv6.

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code>	Enters global configuration mode.
Step 2	<code>route-map <i>name</i> permit 10</code>	Create route-map with allow statement to advertise routes via route-leaking to the customer VRF and subsequently remote VTEPs.

Configuring Custom VRF Context on the Border Node - 2

This procedure applies equally to IPv6.

Procedure

	Command or Action	Purpose
Step 1	<code>configure terminal</code>	Enter global configuration mode.
Step 2	<code>vrf context <i>vrf-name</i></code>	Configure the VRF.
Step 3	<code>vni <i>number</i></code>	Specify the VNI. The VNI associated with the VRF is often referred to as Layer 3 VNI, L3VNI, or L3VPN. The L3VNI is configured as the common identifier across the participating VTEPs.
Step 4	<code>rd auto</code>	Specify the VRF's route distinguisher (RD). The RD uniquely identifies a VTEP within an L3VNI.
Step 5	<code>ip route 0.0.0.0/0 Null0</code>	Configure default-route in common VRF to attract traffic towards Border Node with Shared Internet VRF.
Step 6	<code>address-family ipv4 unicast</code>	Configure the IPv4 address family. This configuration is required for IPv4 over VXLAN with IPv4 underlay.
Step 7	<code>route-target both {<i>auto</i> <i>rt</i>}</code>	Configure the route target (RT) for the import and export of EVPN and IPv4 prefixes within the IPv4 address family. If you enter an RT,

	Command or Action	Purpose
		the following formats are supported: ASN2:NN, ASN4:NN, or IPV4:NN.
Step 8	route-target both {auto rt} evpn	Configure the route target (RT) for the import and export of EVPN and IPv4 prefixes within the IPv4 address family. If you enter an RT, the following formats are supported: ASN2:NN, ASN4:NN, or IPV4:NN.
Step 9	route-target both shared-vrf-rt	Configure a special route target (RT) for the import/export of the shared IPv4 prefixes. An additional import/export map for further qualification is supported.
Step 10	route-target both shared-vrf-rt evpn	Configure a special route target (RT) for the import/export of the shared IPv4 prefixes. An additional import/export map for further qualification is supported.
Step 11	import vrf default map name	Permits all routes, from VRF default, from being imported into the custom VRF according to the specific route-map.

Configuring Custom VRF Instance in BGP on the Border Node

This procedure applies equally to IPv6.

Procedure

	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.
Step 2	router bgp autonomous-system-number	Configure BGP.
Step 3	vrf vrf-name	Specify the VRF.
Step 4	address-family ipv4 unicast	Configure address family for IPv4.
Step 5	advertise l2vpn evpn	Enable the advertisement of EVPN routes within IPv4 address-family.
Step 6	network 0.0.0.0/0	Creating IPv4 default-route network statement.
Step 7	maximum-paths ibgp number	Enabling equal cost multipathing (ECMP) for iBGP prefixes.
Step 8	maximum-paths number	Enabling equal cost multipathing (ECMP) for eBGP prefixes.

Example - Configuration Centralized VRF Route-Leaking - VRF Default with Custom VRF

An example of Centralized VRF route-leaking with VRF default

Configuring VXLAN BGP EVPN Border Node for VRF Default

The VXLAN BGP EVPN Border Node provides centralized access to VRF default. The leaking configuration is localized such that control-plane leaking and data-path forwarding following the same path. Most significantly is the VRF configuration of the Border Node and the advertisement of the default-route and less specific prefixes (aggregates) into the respective destination VRFs.

```
ip route 0.0.0.0/0 10.9.9.1
!
ip prefix-list PL_DENY_EXPORT seq 5 permit 0.0.0.0/0
!
route-map permit 10
match ip address prefix-list PL_DENY_EXPORT
route-map RM_DENY_EXPORT permit 20
route-map RM_PERMIT_IMPORT permit 10
!
vrf context Blue
  vni 51010
  ip route 0.0.0.0/0 Null0
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
  import vrf default map RM_PERMIT_IMPORT
  export vrf default 100 map RM_DENY_EXPORT allow-vpn
!
vlan 2110
  vn-segment 51010
!
interface Vlan2110
  no shutdown
  mtu 9216
  vrf member Blue
  no ip redirects
  ip forward
!
vrf context Red
  vni 51020
  ip route 0.0.0.0/0 Null0
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
  import vrf default map RM_PERMIT_IMPORT
  export vrf default 100 map RM_DENY_EXPORT allow-vpn
!
vlan 2120
  vn-segment 51020
!
interface Vlan2120
  no shutdown
  mtu 9216
  vrf member Blue
  no ip redirects
  ip forward
!
interface nve1
  no shutdown
  host-reachability protocol bgp
```

Example - Configuration Centralized VRF Route-Leaking - VRF Default with Custom VRF

```
source-interface loopback1
member vni 51010 associate-vrf
member vni 51020 associate-vrf
!
router bgp 65002
address-family ipv4 unicast
  aggregate-address 10.10.0.0/16
  aggregate-address 10.20.0.0/16
  maximum-paths 2
  maximum-paths ibgp 2
vrf Blue
  address-family ipv4 unicast
  advertise l2vpn evpn
  network 0.0.0.0/0
  maximum-paths ibgp 2
  maximum-paths 2
vrf Red
  address-family ipv4 unicast
  advertise l2vpn evpn
  network 0.0.0.0/0
  maximum-paths ibgp 2
  maximum-paths 2
```