



# 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 2](#)

## 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 are the guidelines and limitations for External Layer-3 Connectivity for VXLAN BGP EVPN Fabrics:

- Support added for Cisco Nexus 3600 platform switches.
- A physical Layer-3 Interface (Parent-Interface) can be used for external Layer-3 connectivity (ie VRF default).
- The Parent-Interface to multiple Sub-Interfaces can not be used for external Layer-3 connectivity (ie Ethernet1/1 for VRF default). A Sub-Interface can be used instead.
- VTEPs do not support VXLAN encapsulated traffic over Parent-Interfaces if Sub-Interfaces are configured. This is regardless of VRF participation.
- VTEPs do not support VXLAN encapsulated traffic over Sub-Interfaces. This is regardless of VRF participation or IEEE 802.1q encapsulation.
- Mixing Sub-Interfaces for VXLAN and non-VXLAN enabled VLANs is not supported.

# 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 External VRF Connectivity and Route Leaking

The following are the guidelines and limitations for External Layer-3 Connectivity for VXLAN BGP EVPN Fabrics:

- Support added for Cisco Nexus 3600 platform switches.
- A physical Layer-3 Interface (Parent-Interface) can be used for external Layer-3 connectivity (ie VRF default).
- The Parent-Interface to multiple Sub-Interfaces can not be used for external Layer-3 connectivity (ie Ethernet1/1 for VRF default). A Sub-Interface can be used instead.
- VTEPs do not support VXLAN encapsulated traffic over Parent-Interfaces if Sub-Interfaces are configured. This is regardless of VRF participation.
- VTEPs do not support VXLAN encapsulated traffic over Sub-Interfaces. This is regardless of VRF participation or IEEE 802.1q encapsulation.
- Mixing Sub-Interfaces for VXLAN and non-VXLAN enabled VLANs is not supported.


## 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 1: Centralized VRF Route-Leaking - Shared Internet 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
<b>Step 1</b>	<b>configure terminal</b>	Enter global configuration mode.
<b>Step 2</b>	<b>vrf context vrf-name</b>	Configure the VRF.
<b>Step 3</b>	<b>vni number</b>	Specify the VNI.  The VNI associated with the VRF is often referred to as Layer-3 VNI, L3VNI or L3VPN.

	Command or Action	Purpose
		The L3VNI is configured as common identifier across the participating VTEPs.
<b>Step 4</b>	<b>rd auto</b>	Specify the VRFs Route Distinguisher (RD). The RD uniquely identifies a VTEP within a L3VNI.
<b>Step 5</b>	<b>address-family ipv4 unicast</b>	Configure the IPv4 Unicast address-family. Required for IPv4 over VXLAN with IPv4 underlay.
<b>Step 6</b>	<b>route-target both {auto   as:vni}</b>	Configure the Route Target (RT) for import/export of IPv4 prefixes within the IPv4 unicast address-family. The Route Target (RT) is used for a per-VRF prefix import/export policy. If <i>as:vni</i> is entered, the value is in the format of ASN:NN, ASN4:NN, or IPv4:NN.
<b>Step 7</b>	<b>route-target both {auto   as:vni } evpn</b>	Configure the Route Target (RT) for import/export of IPv4 prefixes within the IPv4 unicast address-family. The Route Target (RT) is used for a per-VRF prefix import/export policy. If <i>as:vni</i> is entered, the value is in the format of ASN:NN, ASN4:NN, or IPv4:NN.
<b>Step 8</b>	<b>route-target import rt-from-different-vrf</b>	Configure the Route Target (RT) for importing IPv4 prefixes from the leaked-from VRF (ie AS:VNI).
<b>Step 9</b>	<b>route-target import rt-from-different-vrf evpn</b>	Configure the Route Target (RT) for importing IPv4 prefixes from the leaked-from VRF (ie AS:VNI).

## Configuring the BGP VRF instance on the Routing-Block

This procedure applies equally to IPv6.

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b>	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp <i>autonomous-system number</i></b>	Configure BGP.
<b>Step 3</b>	<b>vrf <i>vrf-name</i></b>	Specify the VRF.
<b>Step 4</b>	<b>address-family ipv4 unicast</b>	Configure address family for IPv4
<b>Step 5</b>	<b>advertise l2vpn evpn</b>	Enable the advertisement of EVPN routes within IPv4 address-family.

	Command or Action	Purpose
<b>Step 6</b>	<b>aggregate-address</b> <i>prefix/mask</i>	Create less specific prefix aggregate into the destination VRF.
<b>Step 7</b>	<b>maximum-paths</b> <i>ibgp number</i>	Enabling equal cost multipathing (ECMP) for iBGP prefixes.
<b>Step 8</b>	<b>maximum-paths</b> <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 2: 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
<b>Step 1</b>	<b>configure terminal</b>	Enter global configuration mode.
<b>Step 2</b>	<b>vrf context <i>vrf-name</i></b>	Configure the VRF.
<b>Step 3</b>	<b>vni <i>number</i></b>	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 common identifier across the participating VTEPs.
<b>Step 4</b>	<b>ip route 0.0.0.0/0 <i>next-hop</i></b>	Configure default-route in shared internet VRF to external router (example).
<b>Step 5</b>	<b>rd auto</b>	Specify the VRFs Route Distinguisher (RD).

	Command or Action	Purpose
		The RD uniquely identifies a VTEP within a L3VNI.
<b>Step 6</b>	<b>address-family ipv4 unicast</b>	Configure the IPv4 Unicast address-family. Required for IPv4 over VXLAN with IPv4 underlay.
<b>Step 7</b>	<b>route-target both {auto   <i>as:vni</i>}</b>	Configure the Route Target (RT) for import/export of EVPN and IPv4 prefixes within the IPv4 unicast address-family.
<b>Step 8</b>	<b>route-target both <i>shared-vrf-rt evpn</i></b>	Configure a special Route Target (RT) for the import/export of the shared IPv4 prefixes. 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
<b>Step 1</b>	<b>configure terminal</b>	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp <i>autonomous-system number</i></b>	Configure BGP.
<b>Step 3</b>	<b>vrf <i>vrf-name</i></b>	Specify the VRF.
<b>Step 4</b>	<b>address-family ipv4 unicast</b>	Configure address family for IPv4
<b>Step 5</b>	<b>advertise l2vpn evpn</b>	Enable the advertisement of EVPN routes within IPv4 address-family.
<b>Step 6</b>	<b>aggregate-address <i>prefix/mask</i></b>	Create less specific prefix aggregate into the destination VRF.
<b>Step 7</b>	<b>maximum-paths ibgp <i>number</i></b>	Enabling equal cost multipathing (ECMP) for iBGP prefixes.
<b>Step 8</b>	<b>maximum-paths <i>number</i></b>	Enabling equal cost multipathing (ECMP) for eBGP prefixes.

## Configuring Custom VRF Context on the Border Node - 1

This procedure applies equally to IPv6.



**Procedure**

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b>	Enters global configuration mode.
<b>Step 2</b>	<b>vrf context</b> <i>vrf-name</i>	Configure the VRF.
<b>Step 3</b>	<b>vni</b> <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 the common identifier across the participating VTEPs.
<b>Step 4</b>	<b>rd</b> <i>auto</i>	Specify the VRFs Route Distinguisher (RD).  The Route Distinguisher (RD) uniquely identifies a VTEP within a L3VNI.
<b>Step 5</b>	<b>ip route</b> <i>0.0.0.0/0 Null0</i>	Configure default-route in common VRF to attract traffic towards Border Node with Shared Internet VRF.
<b>Step 6</b>	<b>address-family</b> <i>ipv4 unicast</i>	Configure the IPv4 Unicast address-family.  Required for IPv4 over VXLAN with IPv4 underlay.
<b>Step 7</b>	<b>route-target</b> <i>both {auto   as:vni}</i>	Configure the Route Target (RT) for import/export of IPv4 prefixes within the IPv4 unicast address-family The Route Target (RT) is used for a per-VRF prefix import/export policy. If <i>as:vni</i> is entered, the value is in the format of ASN:NN, ASN4:NN, or IPv4:NN.
<b>Step 8</b>	<b>route-target</b> <i>both {auto   as:vni} evpn</i>	Configure the Route Target (RT) for import/export of IPv4 prefixes within the IPv4 unicast address-family The Route Target (RT) is used for a per-VRF prefix import/export policy. If <i>as:vni</i> is entered, the value is in the format of ASN:NN, ASN4:NN, or IPv4:NN.
<b>Step 9</b>	<b>import map</b> <i>name</i>	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
<b>Step 1</b>	<b>configure terminal</b>	Enters global configuration mode.

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

	Command or Action	Purpose
<b>Step 2</b>	<b>router bgp</b> <i>autonomous-system-number</i>	Configure BGP.
<b>Step 3</b>	<b>vrf</b> <i>vrf-name</i>	Specify the VRF.
<b>Step 4</b>	<b>address-family ipv4 unicast</b>	Configure address family for IPv4.
<b>Step 5</b>	<b>advertise l2vpn evpn</b>	Enable the advertisement of EVPN routes within IPv4 address-family.
<b>Step 6</b>	<b>network 0.0.0.0/0</b>	Creating IPv4 default-route network statement.
<b>Step 7</b>	<b>maximum-paths ibgp</b> <i>number</i>	Enabling equal cost multipathing (ECMP) for iBGP prefixes.
<b>Step 8</b>	<b>maximum-paths</b> <i>number</i>	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
!
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 3: 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
<b>Step 1</b>	<b>configure terminal</b>	Enters global configuration mode.
<b>Step 2</b>	<b>ip route 0.0.0.0/0 <i>next-hop</i></b>	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
<b>Step 1</b>	<b>configure terminal</b>	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp <i>autonomous-system number</i></b>	Configure BGP.
<b>Step 3</b>	<b>address-family ipv4 unicast</b>	Configure address family for IPv4.
<b>Step 4</b>	<b>aggregate-address <i>prefix/mask</i></b>	Create less specific prefix aggregate in VRF default.
<b>Step 5</b>	<b>maximum-paths <i>number</i></b>	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
<b>Step 1</b>	<b>configure terminal</b>	Enters global configuration mode.
<b>Step 2</b>	<b>ip prefix-list <i>name</i> seq 5 permit 0.0.0.0/0</b>	Configure IPv4 prefix-list for default-route filtering.
<b>Step 3</b>	<b>route-map <i>name</i> deny 10</b>	Create route-map with leading deny statement to prevent the default-route of being leaked.
<b>Step 4</b>	<b>match ip address prefix-list <i>name</i></b>	Match against the IPv4 prefix-list that contains the default-route.

	Command or Action	Purpose
<b>Step 5</b>	<b>route-map <i>name</i> permit 20</b>	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
<b>Step 1</b>	<b>configure terminal</b>	Enters global configuration mode.
<b>Step 2</b>	<b>route-map <i>name</i> permit 10</b>	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
<b>Step 1</b>	<b>configure terminal</b>	Enters global configuration mode.
<b>Step 2</b>	<b>vrf context <i>vrf-name</i></b>	Configure the VRF.
<b>Step 3</b>	<b>vni <i>number</i></b>	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 common identifier across the participating VTEPs.
<b>Step 4</b>	<b>rd auto</b>	Specify the VRFs Route Distinguisher (RD).  The Route Distinguisher (RD) uniquely identifies a VTEP within a L3VNI.
<b>Step 5</b>	<b>ip route 0.0.0.0/0 Null0</b>	Configure default-route in common VRF to attract traffic towards Border Node with Shared Internet VRF.
<b>Step 6</b>	<b>address-family ipv4 unicast</b>	Configure the IPv4 Unicast address-family.  Required for IPv4 over VXLAN with IPv4 underlay.

	Command or Action	Purpose
<b>Step 7</b>	<b>route-target both auto</b>   <i>AS:VNI</i>	Configure the Route Target (RT) for import/export of EVPN and IPv4 prefixes within the IPv4 unicast address-family.
<b>Step 8</b>	<b>route-target both auto</b>   <i>AS:VNI evpn</i>	Configure the Route Target (RT) for import/export of EVPN and IPv4 prefixes within the IPv4 unicast address-family.
<b>Step 9</b>	<b>route-target both</b> <i>shared-vrf-rt</i>	Configure a special Route Target (RT) for the import/export of the Shared IPv4 prefixes. Additional import/export map for further qualification is supported
<b>Step 10</b>	<b>route-target both</b> <i>shared-vrf-rt evpn</i>	Configure a special Route Target (RT) for the import/export of the Shared IPv4 prefixes. Additional import/export map for further qualification is supported
<b>Step 11</b>	<b>import vrf default map</b> <i>name</i>	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
<b>Step 1</b>	<b>configure terminal</b>	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>autonomous-system-number</i>	Configure BGP.
<b>Step 3</b>	<b>vrf</b> <i>vrf-name</i>	Specify the VRF.
<b>Step 4</b>	<b>address-family ipv4 unicast</b>	Configure address family for IPv4.
<b>Step 5</b>	<b>advertise l2vpn evpn</b>	Enable the advertisement of EVPN routes within IPv4 address-family.
<b>Step 6</b>	<b>network 0.0.0.0/0</b>	Creating IPv4 default-route network statement.
<b>Step 7</b>	<b>maximum-paths ibgp</b> <i>number</i>	Enabling equal cost multipathing (ECMP) for iBGP prefixes.
<b>Step 8</b>	<b>maximum-paths</b> <i>number</i>	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
```



```
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
```

