

Default Gateway Coexistence of HSRP and Anycast Gateway (VXLAN EVPN)

This chapter contains these sections:

- Default gateways in VXLAN EVPN fabrics, on page 1
- Best practice for migrating from Classic Ethernet or FabricPath to VXLAN, on page 2
- Migrate Classic Ethernet or FabricPath to VXLAN, on page 4
- Configure an external port on a border leaf for migration, on page 5
- Configure external IP address for migration, on page 6

Default gateways in VXLAN EVPN fabrics

A default gateway is a network device that

- enables communication between devices in a local network and external networks
- supports simultaneous operation of HSRP-based and anycast (DAG) gateway configurations during migration in VXLAN EVPN fabrics, and
- provides seamless migration capability between Classic Ethernet/FabricPath and VXLAN EVPN environments with minimal disruption.

The default gateway coexistence feature allows both traditional HSRP gateways and distributed anycast gateways (DAG) to function at the same time within VXLAN EVPN fabrics. By configuring a common default gateway MAC and IP address, you avoid disruptive cut-overs and inefficient routing methods. This feature is enabled on the VXLAN EVPN side, specifically on border nodes that connect to Classic Ethernet or FabricPath networks. As a result, migrations can occur without software or hardware upgrades on the legacy infrastructure, supporting efficient and less disruptive transitions.

After completing the required premigration steps on the Classic Ethernet or FabricPath HSRP gateway, both the DAG (on the VXLAN network) and the HSRP gateway (on the legacy network) can operate simultaneously for the same VLAN. This enables minimal traffic impact and optimal routing during migration.

Layer 2 Interconnection and Migration Steps

Migration can now be performed with minimal traffic impact even when both DAG is functional on VXLAN network and HSRP gateway is functional on Classic Ethernet / FabricPath network for the same VLAN after

the premigration step is performed on the Classic Ethernet / FabricPath HSRP gateway. For more information, see details for premigration step in Migrate Classic Ethernet or FabricPath to VXLAN, on page 4.

Coexistence of both DAG and HSRP gateway was not possible earlier for the same VLAN even after the premigration step was performed. This coexistence will enable optimal routing for the Layer 3 workloads that are migrated to VXLAN network during migration.

Layer 2 Interconnection

- Interconnecting the two networks via Layer 2 is crucial to facilitate seamless workload migration from Classic Ethernet / FabricPath to VXLAN.
- The border leaf on VXLAN network is connected via a Layer 2 interface to the Classic Ethernet / FabricPath network.
- The Layer 2 link can be a port channel trunk or a physical Ethernet trunk.
- The VXLAN border leaf switch can be a vPC or a NX-OS switch and the switch can be a TOR or an EOR. Similarly, the Classic Ethernet / FabricPath border-edge switch can be a vPC or a NX-OS switch. The switch could also host the HSRP gateway for the Classic Ethernet / FabricPath network.

For migration, you must configure the following on the VXLAN border leaf:

- The Layer 2 ports connecting the two network infrastructures must be configured as **port-type external**. These ports are referred as *external* interfaces.
- A unique Burned In Address (BIA) address for IPv4 and IPv6 must be configured on the SVI of each VXLAN border leaf during migration of the VLAN.
- If the VXLAN border leaf is in a vPC configuration, then the BIA address for the SVI must be different on both switches.

The following table provides few Layer 2 interconnection combinations:

Table 1: Layer 2 Interconnection Combinations

VXLAN Border Leaf	Classic Ethernet / FabricPath Border Edge Switch
VPC	VPC
NX-OS switch	NX-OS switch
NX-OS switch	VPC
VPC	NX-OS switch

Best practice for migrating from Classic Ethernet or FabricPath to VXLAN

Follow these best practices when migrating workloads from Classic Ethernet or FabricPath to VXLAN:

- Carve the ingress PACL region and ensure it is available before configuring workload migration for Cisco Nexus 9300-FX/FX2 platforms deployed as VXLAN border leaf nodes.

For Example: Before you configure the **port-type external** command on the ports connecting the VXLAN and Classic Ethernet / FabricPath networks, verify that the PACL region is carved. Use the **show hardware access-list tcam region** command to check. If unavailable, configure the region with **hardware access-list tcam region ing-ifacl 512** and reload the switch.

- Perform migration of IPv4 and IPv6 applications sequentially, as follows:

  1. Complete the premigration step for the IPv4 HSRP gateway IP for each VLAN. For more information, see details for premigration step in Migrate Classic Ethernet or FabricPath to VXLAN, on page 4.
  2. Configure SVIs with a BIA address for IPv4 on each VXLAN border leaf node connecting to the Classic Ethernet / FabricPath network.
  3. Migrate all the IPv4 hosts from Classic Ethernet / FabricPath to VXLAN side.
  4. After all IPv4 hosts are migrated, repeat the premigration and migration procedure for IPv6.



Note Limit the migration to 1000 concurrent hosts. Begin new migrations only after the previous batch completes.

Follow these requirements during migration:

- Ensure that no ingress PACL policies are configured on the external interfaces before migration. Remove any existing policies before you configure the **port-type external** command.
- Before migration, ensure that the *Extended IFACL* feature is not configured using the **hardware access-list tcam label ing-ifacl 6** command.
- If a vPC VXLAN border leaf is configured, enable Layer 3 peer-router functionality.
- If the Suppress ARP or Suppress ND feature is enabled on the VXLAN network during migration, ensure the host is learned in the corresponding ARP or ND table of the VXLAN border leaf. Send a GARP or ND message before moving the host to VXLAN.

  For example: If a host is not learned in the ARP table after migration, traffic from Classic Ethernet / FabricPath to the host may fail. Send a GARP from the host to update the ARP table.

  - When host 192.0.2.8 is being moved to VXLAN, initially, it is not learned as shown:

```
switch# sh ip arp 192.0.2.8 vrf vrf1501

Flags: * - Adjacencies learnt on non-active FHRP router
       + - Adjacencies synced via CFSoE
       # - Adjacencies Throttled for Glean
       CP - Added via L2RIB, Control plane Adjacencies
       PS - Added via L2RIB, Peer Sync
       RO - Re-Originated Peer Sync Entry
       D - Static Adjacencies attached to down interface

IP ARP Table
Total number of entries: 1
Address         Age       MAC Address     Interface       Flags
192.0.2.8       00:00:04  0000.8aa9.79d3  Vlan1001

switch(config)# sh ip route 192.0.2.8 vrf vrf1501
```

```
192.0.2.8/32, ubest/mbest: 1/0, attached
    *via 192.0.2.8, Vlan1001, [190/0], 00:00:14, hmm
```

- After sending GARP from host 192.0.2.8, the ARP table output of the border leaf switch is as shown:

```
switch# show ip arp 192.0.2.8 vrf vrf1501

Flags: * - Adjacencies learnt on non-active FHRP router
       + - Adjacencies synced via CFSoE
       # - Adjacencies Throttled for Glean
       CP - Added via L2RIB, Control plane Adjacencies
       PS - Added via L2RIB, Peer Sync
       RO - Re-Originated Peer Sync Entry
       D - Static Adjacencies attached to down interface

IP ARP Table
Total number of entries: 1
Address         Age       MAC Address     Interface       Flags
192.0.2.8       00:00:04  0000.8aa9.79d3  Vlan1001

switch(config)# sh ip route 192.0.2.8 vrf vrf1501

192.0.2.8/32, ubest/mbest: 1/0, attached
    *via 192.0.2.8, Vlan1001, [190/0], 00:00:14, hmm
```

- After GARP, the host is moved to leaf in the VXLAN network as shown:

```
switch(config)# sh ip route 192.0.2.8 vrf vrf1501

192.0.2.8/32, ubest/mbest: 1/0
    *via 192.0.2.5%default, [200/0], 00:00:23, bgp-200, internal, tag 200, segid:
11501 tunnelid: 0x2020205 encap: VXLAN
```

Migrate Classic Ethernet or FabricPath to VXLAN

Use this task when upgrading your network infrastructure from legacy Classic Ethernet or FabricPath environments to a VXLAN-based fabric, minimizing downtime and configuration errors during migration.

Before you begin

- Verify that the PACL region on FX/FX2 platforms is carved and enabled using the **show hardware access-list tcam region** command.
- If the PACL region is not carved, configure and enable it before proceeding.

Procedure

Step 1 Establish a Layer 2 interconnection between the Classic Ethernet / FabricPath network and the VXLAN fabric using a physical Ethernet port or port channel between the VXLAN border leaf and Classic Ethernet / FabricPath edge switch. vPC is supported. For more information, see Table 1: Layer 2 Interconnection Combinations, on page 2.

Step 2 If using a vPC VXLAN border leaf, configure the **peer-gateway** and **layer3 peer-router** commands.

Step 3 In HSRP, configure the Anycast gateway MAC address from the VXLAN fabric for the target VLAN on Classic Ethernet / FabricPath using **mac-address** *address* {**ipv4** | **ipv6**}.

This triggers a GARP, updating VLAN hosts with the Anycast gateway MAC address.

Step 4 On the VXLAN border leaf, configure the connecting port as external using **port-type external**.

Step 5 Create the SVI for the migrating VLAN on all VXLAN leafs (including border leaf); keep the SVI in a shutdown state if the VLAN carries routed traffic.

Step 6 On the VXLAN border leaf, configure the SVI with the appropriate IPv4 or IPv6 BIA address, ensuring the BIA is in the same subnet as the source SVI.

This avoids MAC collisions during migration and enables proxy-ARP/ND requests using the BIA IP and VDC MAC.

Step 7 Enable (**no shutdown**) the SVI for the VLAN on all VXLAN leafs when ready to move workloads.

Workloads now route using VXLAN Distributed Anycast Gateway (DAG).

Step 8 Move all hosts for the VLAN from Classic Ethernet / FabricPath to VXLAN, migrating one address family (IPv4 or IPv6) at a time.

Step 9 Once all hosts for a VLAN are moved, remove the Classic Ethernet / FabricPath HSRP gateway SVI and BIA address from the VXLAN border leaf. After all VLANs for both IPv4 and IPv6 are migrated, remove the external port configuration.

The VLANs and workloads now operate entirely within the VXLAN fabric. Gateway configurations are cleaned up from the Classic Ethernet / FabricPath side and traffic routing is enabled via VXLAN DAG.

Configure an external port on a border leaf for migration

Before you begin

- Complete premigration steps for the VLAN on the Classic Ethernet / FabricPath network.
- Configure an Anycast gateway MAC address in HSRP for the VLAN on the Classic Ethernet / FabricPath side.

Use this task when migrating hosts or workloads in a VLAN from Classic Ethernet / FabricPath to VXLAN. The configuration allows seamless Layer 2 connectivity between networks during transition.

SUMMARY STEPS

1. Enter configuration mode.
2. Select and enter configuration for the port channel interface.
3. Designate the port channel as an external port to connect to the Classic Ethernet / FabricPath network.

DETAILED STEPS

Procedure

Step 1 Enter configuration mode.

Example:

```
configure terminal
```

Step 2 Select and enter configuration for the port channel interface.

Example:

```
switch(config)# interface port-channel 40
switch(config-if)#
```

Step 3 Designate the port channel as an external port to connect to the Classic Ethernet / FabricPath network.

Example:

```
switch(config-if)# port-type external
switch(config-if)#
```

The port channel is configured as an external port on the border leaf, enabling migration of VLAN hosts from Classic Ethernet / FabricPath to VXLAN.

What to do next

Configure a BIA address for IPv4 or IPv6 on the SVI where VLAN hosts are being migrated. For details, see Configure external IP address for migration, on page 6.

Configure external IP address for migration

Use these steps to configure the necessary external IP addresses and VRF membership for a VLAN interface when migrating network connectivity.

Follow these steps to configure an external IP address for migration:

Before you begin

- Ensure you have access to the device with appropriate privileges.
- Identify the VLAN ID, VRF name, and IP addresses to be configured.

Procedure

Step 1 Enter configuration mode.

Example:

```
switch# configure terminal
switch(config)#
```

Step 2 Enter interface configuration mode for the target VLAN.

Example:

```
switch(config)# interface vlan 1100
switch(config-if)#
```

Step 3 Assign the interface to the required VRF.

Example:

```
switch(config-if)# vrf member vrf50
```

Step 4 Configure the primary IPv4 address on the interface.

Example:

```
switch(config-if)# ip address 192.168.1.1/24
```

Assigns an IPv4 address to the interface.

Step 5 Configure a secondary external IPv4 address using the use-bia option.

Example:

```
switch(config-if)# ip address 192.168.1.10/24 secondary use-bia
```

Step 6 Assign the primary IPv6 address to the interface.

Example:

```
switch(config-if)# ipv6 address 2001:DB8:1::1/64
```

Step 7 Configure a secondary external IPv6 address using the use-bia option.

Example:

```
switch(config-if)# ipv6 address 2001:DB8:1::10/64 use-bia
```

The VLAN interface is configured with the designated external IPv4 and IPv6 addresses for migration.