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

This chapter contains these sections:

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

This feature provides coexistence between traditional Default Gateways using First Hop Gateway Protocol (HSRP being the mode supported in this release), and Distributed Anycast Gateway (DAG) for VXLAN EVPN fabrics. Instead of a disruptive cut-over or inefficient hair pinning, Default Gateways with HSRP can now be active at the same time as VXLAN EVPNs DAG, as long as the common Default Gateway MAC and IP is configured. The functionality as part of this feature provides ease for migration and coexistence between Classic Ethernet / FabricPath and VXLAN EVPN fabrics. This functionality is solely enabled on the VXLAN EVPN side, more specifically on the Border nodes neighboring the Classic Ethernet / FabricPath network. This feature allows more efficient routing and less disruptive migrations without the requirement for Software or Hardware upgrades on the Classic Ethernet / FabricPath side.

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 Configuring Classic Ethernet / FabricPath to VXLAN Migration, 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

Guidelines and Limitations for Migrating from Classic Ethernet / FabricPath to VXLAN

- Ingress PACL region must be carved and made available before configuring the migration of workloads for 9300-FX/FX2/FX3/GX/GX2B platforms deployed as VXLAN border leaf nodes.

  For example: You need to verify if the PACL region is carved before configuring the **port-type external** command on the ports connecting the VXLAN and Classic Ethernet / FabricPath networks. You can verify if the ingress PACL region is configured by using the **show hardware access-list tcam region** command. If the region is unavailable, configure the region using the **hardware access-list tcam region ing-ifacl 512** command. Ensure that you reload the switch after the PACL region is configured.
- Verify that there is no ingress PACL policy configured on the external interfaces before migration. If they are configured, you must remove them before configuring the **port-type external** command.
- vPC Fabric Peering, Egress CNTACL, VRRP, and VXLAN Flood and Learn are not supported with this migration. Also, this migration does not support moving workloads that are multicast sources or receivers.
- It is recommend that you configure only up to six external interfaces.

- For migration, ensure that you do not have the *Extended IFACL* feature configured using the **hardware access-list tcam label ing-ifacl 6** command.
- Migration of IPv4 and IPv6 applications must be performed sequentially as mentioned below:
  1. Premigration step must be performed on HSRP gateway for IPv4 gateway-IP for a particular VLAN. For more information, see details for premigration step in Configuring Classic Ethernet / FabricPath to VXLAN Migration, on page 4.
  2. The migration procedure in terms of configuring SVIs with BIA address for IPv4 must be performed 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 the IPv4 hosts for all VLANs are migrated from Classic Ethernet / FabricPath to VXLAN, the premigration step and migration procedure has to be repeated for IPv6.

Note It is recommended that you limit the migration of concurrent host to a maximum of 1000 hosts. Start the next migration only after the previous migration of hosts is complete.

- If we have a vPC VXLAN border leaf configured, Layer 3 peer-router needs to be enabled.
- If the Suppress ARP or Suppress ND feature is enabled on the VXLAN network during Classic Ethernet / FabricPath to VXLAN migration, the host must be learned in the respective ARP or ND tables on the VXLAN border leaf. You can send a GARP/ND before moving the host to VXLAN.

  If adjacency is not learned for the host that is moved to VXLAN, then traffic from the host behind Classic Ethernet / FabricPath network to this host can fail on the Classic Ethernet / FabricPath network.

  For example:

  - When host 10.10.1.8 is being moved to VXLAN, initially, it is not learned as shown:

```
switch# sh ip arp 10.10.1.8 vrf vrf1501

IP ARP Table
Total number of entries: 0
Address         Age       MAC Address     Interface       Flags
switch#

switch(config)# sh ip route 10.10.1.8 vrf vrf1501

10.10.1.0/24, ubest/mbest: 2/0, attached
    *via 10.10.1.1, Vlan1001, [0/0], 22:55:42, direct
    *via 10.10.1.4, Vlan1001, [0/0], 22:55:42, direct
```

  - After sending GARP from host 10.10.1.8, the ARP table output of the border leaf switch is as shown:

```
switch# sh ip arp 10.10.1.8 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
10.10.1.8       00:00:04  0000.8aa9.79d3  Vlan1001

switch(config)# sh ip route 10.10.1.8 vrf vrf1501

10.10.1.8/32, ubest/mbest: 1/0, attached
    *via 10.10.1.8, 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 10.10.1.8 vrf vrf1501

10.10.1.8/32, ubest/mbest: 1/0
    *via 2.2.2.5%default, [200/0], 00:00:23, bgp-200, internal, tag 200, segid:
11501 tunnelid: 0x2020205 encap: VXLAN
```

- Beginning with Cisco NX-OS Release 10.5(2)F, the following features are supported on Cisco Nexus 9500 Series switches with N9K-X9736C-FX3 line card:
  - HSRP
  - VRRP
  - VXLAN Flood and Learn (Static Ingress Replication)
  - FabricPath to VXLAN Migration

Configuring Classic Ethernet / FabricPath to VXLAN Migration

To migrate workloads from Classic Ethernet / FabricPath to VXLAN, perform these steps:

Note Check if PACL region was carved using the **show hardware access-list tcam region** command for EX/FX/FX2 platforms. If not, ensure that PACL region is carved and made available before configuring migration of workloads.

Procedure

Step 1 Ensure that you have a Layer 2 interconnection between the VXLAN and the Classic Ethernet / FabricPath networks. As specified in Table 1: Layer 2 Interconnection Combinations, on page 2, this can be between a VXLAN border leaf (with or without vPC configuration) and the Classic Ethernet / FabricPath edge switch (with or without vPC configuration). This interface can be a physical Ethernet Layer 2 port or a Layer 2 port channel. For more information, see Configuring VXLAN BGP EVPN.

Step 2 If there is a vPC VXLAN border leaf, ensure that **peer-gateway** and **layer3 peer-router** commands are configured.

Step 3 As part of the premigration step, configure the Anycast gateway MAC address (value present on VXLAN fabric) in HSRP for a particular VLAN on the Classic Ethernet / FabricPath network using the **mac-address** *address* {**ipv4** | **ipv6**} under HSRP.

With this premigration step configured, a GARP is triggered and it will update all hosts in that VLAN with the Anycast gateway MAC address.

Step 4 Configure a port on VXLAN border leaf as an external port using the **port-type external** for the Layer 2 port connecting the two fabrics.

Step 5 Ensure that the SVI for the VLAN that is to be migrated is configured on all the VXLAN leafs including border leaf. This step is required if there is a routed traffic for the VLAN. Ensure that you keep the SVI in the shutdown state.

Step 6 On the VXLAN border leaf, ensure that the SVI is configured with the IPv4 and/or IPv6 BIA address.

This configuration is required so that a proxy-ARP or ND request can be sent using this BIA IP address as the source-IP address and VDC-MAC as source-MAC over the external interfaces to Classic Ethernet / FabricPath network. This configuration ensures that you do not use the regular gateway-IP and the Anycast gateway MAC. This configuration will prevent collision of MACs after the premigration step.

Step 7 The IPv4 or IPv6 BIA address must be in the same subnet as the source address on the SVI of the VXLAN border leaf.

Step 8 Bring up the SVI using the **no shutdown** command under the SVI on all Leafs of VXLAN.

With this configuration, when a workload on a VLAN is moved from Classic Ethernet / FabricPath to VXLAN, it can route on the source VXLAN leaf following the VXLAN Distributed Anycast Gateway (DAG) paradigm.

Step 9 Hosts for the VLAN that continue to exist on Classic Ethernet / FabricPath side will route at the HSRP gateway. With this, both DAG and HSRP are coexisting and functional for the VLAN.

Step 10 Move all hosts from Classic Ethernet / FabricPath to VXLAN for a given VLAN.

Step 11 Ensure that all the hosts in one address family (IPv4 or IPv6) are migrated completely before migrating the other address family.

Step 12 After all the hosts for a VLAN are moved from Classic Ethernet / FabricPath to VXLAN, the HSRP gateway SVI can be removed from the Classic Ethernet / FabricPath side for the VLAN.

Step 13 After all the VLANs have been migrated from Classic Ethernet / FabricPath to VXLAN for both address families (IPv4 and IPv6), run the **no port-type external** command on the Layer 2 interfaces connecting the two fabrics. The BIA address are no longer required and can be removed from the SVI of border leafs.

The migration will now be complete.

Configuring an External Port on Border Leaf for Migration

For migrating applications or workloads from Classic Ethernet / FabricPath to VXLAN, you must configure ports on border leaf as an external port for Layer 2 interconnection.

Before you begin

For migrating hosts in a VLAN from Classic Ethernet / FabricPath to VXLAN, ensure that you complete the premigration step for the VLAN on the Classic Ethernet / FabricPath side. For this, configure an Anycast gateway MAC address in HSRP for Classic Ethernet / FabricPath network for the VLAN.

SUMMARY STEPS

1. **configure terminal**
2. **interface port-channel** *number*

3. **port-type external**

DETAILED STEPS

Procedure

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters configuration mode.

Step 2	**interface port-channel** *number*
	Example:
	`switch(config)# interface port-channel 40`
	`switch(config-if)#`
		Enters configuration mode and configures a port channel interface.

Step 3	**port-type external**
	Example:
	`switch(config-if)# port-type external`
	`switch(config-if)#`
		Configures the interface to be the external interface that connects to a Classic Ethernet / FabricPath network.

What to do next

As mentioned in the steps, we need to configure a BIA address for IPv4 or IPv6 on the SVI where VLAN-hosts are being moved from Classic Ethernet / FabricPath to VXLAN. For configuration this, see Configuring External IP Address for Migration, on page 6.

Configuring External IP Address for Migration

SUMMARY STEPS

1. **configure terminal**
2. **interface vlan** *vlan-id*
3. **vrf member** *vrf-name*
4. **ip address** *address netmask*
5. **ip address** *address netmask* **secondary use-bia**
6. **ipv6 address** *address netmask*
7. **ipv6 address** *address netmask* **use-bia**

DETAILED STEPS

Procedure

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters configuration mode.

Step 2	**interface vlan** *vlan-id*
	Example:
	`switch(config)# interface vlan 1100`
	`switch(config-if)#`
		Creates a VLAN interface and enters the interface configuration mode.

Step 3	**vrf member** *vrf-name* **Example:** `switch(config-if)# vrf member vrf50`	Adds this interface to VRF.
Step 4	**ip address** *address netmask* **Example:** `switch(config-if)# ip address 192.168.1.1/24`	Assigns an IPv4 address to the interface.
Step 5	**ip address** *address netmask* **secondary use-bia** **Example:** `switch(config-if)# ip address 192.168.1.10/24 secondary use-bia`	Configures external IPv4 address.
Step 6	**ipv6 address** *address netmask* **Example:** `switch(config-if)# ipv6 address 2001:DB8:1::1/64`	Assigns an IPv6 address to the interface.
Step 7	**ipv6 address** *address netmask* **use-bia** **Example:** `switch(config-if)# ipv6 address 2001:DB8:1::10/64 use-bia`	Configures external IPv6 address.