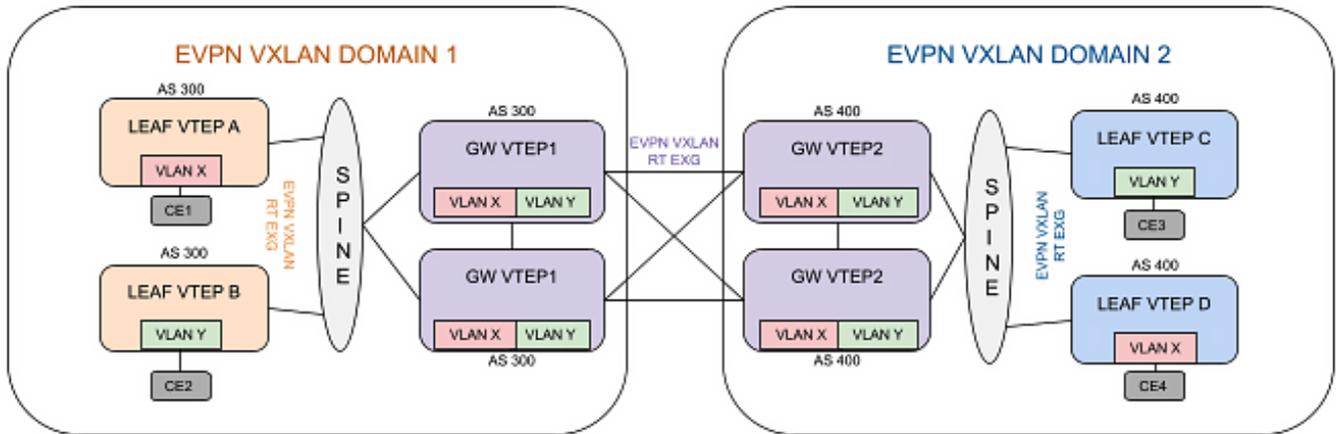


Description

This feature provides the ability to interconnect EVPN VXLAN domains. Domains may or may not be within the same data center network, and the decision to stretch/interconnect a subnet between domains is configurable. The following diagram shows a multi-domain deployment using symmetric IRB. Note that two domains are shown for simplicity, but this solution supports any number of domains.



Within domain #1 and domain #2, VTEPs exchange EVPN reachability as normal. Between domains, gateway nodes advertise intra-domain EVPN routes with the gateway inserting itself as the nexthop. From the perspective of a gateway node, there is the **local EVPN domain**, which contains intra-domain VTEPs and peer gateway nodes, and the **remote EVPN domain**, which contains the gateway nodes of other EVPN VXLAN domains. From the perspective of a leaf VTEP, all routes to remote domains appear to be directly connected to the gateway nodes in that domain. This means that leaf VTEPs are not visible to remote domains.

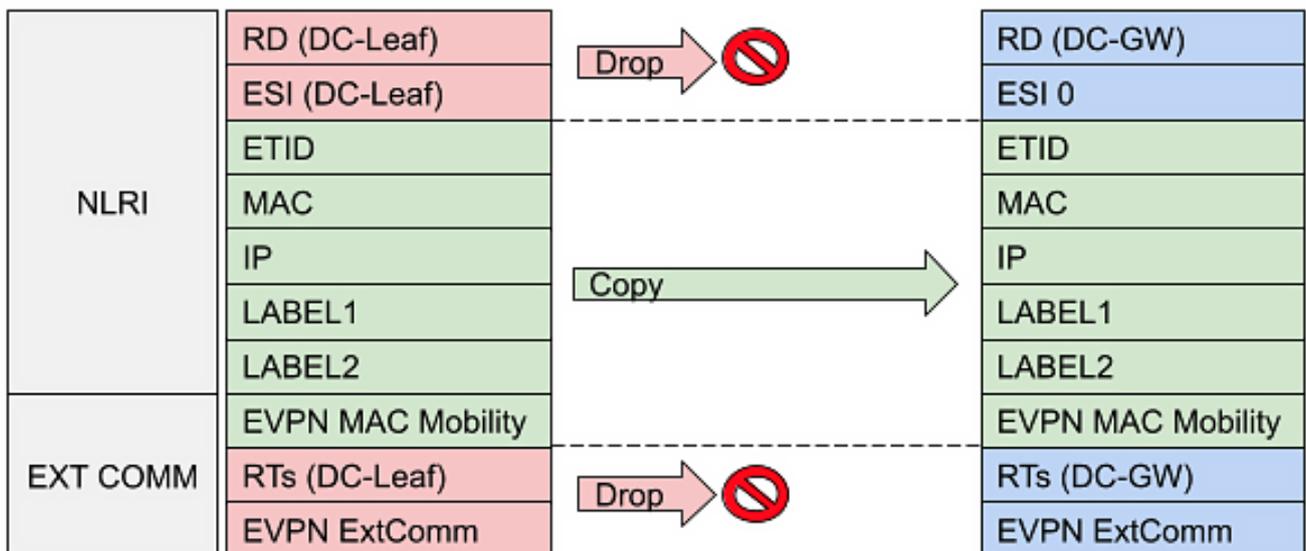
Collapsing all VTEPs in a domain behind the gateway nodes results in significant hardware savings in terms of flood-list entries. The following tables illustrate the per-VLAN flood-list and VXLAN address table for each node from the diagram above:

VTEP	VLAN	Flood-List
LEAF A	X	[GW1]
LEAF B	Y	[GW1]
GW1	X	[LEAF_A, GW2]
GW1	Y	[LEAF_B, GW2]
GW2	X	[LEAF_D, GW1]
GW2	Y	[LEAF_C, GW1]
LEAF C	Y	[GW2]

VTEP	VLAN	Flood-List
LEAF D	X	[GW2]
VTEP	VLAN	VXLAN MAC Entries
LEAF A	X	CE4 ? GW1
LEAF B	Y	CE3 ? GW1
GW1	X	CE1 ? LEAF_A, CE4 ? GW2
GW1	Y	CE2 ? LEAF_B, CE3 ? GW2
GW2	X	CE1 ? GW1, CE4 ? LEAF_D
GW2	Y	CE2 ? GW1, CE3 ? LEAF_C
LEAF C	Y	CE2 ? GW2
LEAF D	X	CE1 ? GW2

Gateway Route Advertisement

Gateway nodes only advertise intra-domain MAC-IP (type-2), IMET (type-3), and IP-PREFIX (type-5) routes to remote domains. IP-PREFIX routes, which are advertised using nexthop-self, are updated according to the diagram in the [EVPN L3 Gateway TOI](#). MAC-IP routes are updated as follows:



Exporting type-2 and type-3 routes into the remote EVPN domain sets the route distinguisher (RD) and route targets (RT) to those configured for the remote EVPN domain. This is illustrated in the configuration section below. The Ethernet Segment identifier is set to zero when inter-

domain multi-homing is not enabled. Imported EVPN extended communities other than MAC Mobility are dropped and re-originated.

MAC Mobility

MAC moves within and between EVPN VXLAN domains is supported. The MAC Mobility sequence number is incremented on each move. This means that intra-domain moves result in an EVPN MAC-IP update being sent to remote domains. The remote domain will not reprogram this route, though, as the next hop has not changed.

VXLAN Routing

The asymmetric, symmetric, and centralized gateway routing models are supported with multi-domain EVPN VXLAN, while the overlay may be IPv4 or IPv6. For an IPv4 overlay with L2 VTEPs, the gateway nodes will proxy reply to ARP requests when possible in order to avoid flooding to remote domains.

Gateway Redundancy

Gateway redundancy is accomplished by provisioning an MLAG pair. This allows for underlay redundancy between the gateway node(s) and intra-domain VTEPs and the gateway node(s) and remote domain VTEPs using L3 ECMP.

Please refer to the [Using Anycast Gateway for Multi-Domain EVPN VXLAN](#) and [EVPN VXLAN All-Active Multihoming Gateway](#) for anycast and multihoming redundancy model.

Platform Compatibility

- DCS-7020 series
- DCS-7280 series
- DCS-7500 series
- DCS-7800 series
- vEOS-lab

Feature History

Release	Update
4.26.1F	Initial introduction
4.33.1F	Support for deploying the L3 DCI gateway as an intra-DC route reflector

Configuration

The following sample shows BGP configuration on GW1. Note that the routing model between VTEPs in domain #1 is asymmetric IRB.

```
GW1(config-router-bgp)# show active
router bgp 300
...
!
address-family evpn
  neighbor LEAF_A activate
  neighbor LEAF_B activate
  neighbor GW2 activate
  neighbor GW2 domain remote
!
vlan 10
  rd evpn domain local 10.255.1.1:10
  rd evpn domain remote 10.255.1.1:1010
  route-target import export 64500:10
  route-target import export evpn domain remote 64501:10
  redistribute learned
!
vlan 20
  rd evpn domain local 10.255.1.1:20
  rd evpn domain remote 10.255.1.1:1020
  route-target import export 64500:20
  route-target import export evpn domain remote 64501:20
  redistribute learned
!
vrf red
  rd 10.255.1.1:0
  route-target import evpn 64500:20000
  route-target export evpn 64500:20000
  router-id 10.255.1.1
```

Note that GW2 is explicitly configured to be in the remote EVPN domain, while the leaf VTEPs in domain #1 are in the local EVPN domain by default. Under each MAC-VRF, both the local domain and the remote domain route distinguisher and route target are configured. When the gateway re-exports MAC-IP routes received from the remote domain to the local domain, it will remove the route distinguisher and the route target in the received routes and attach the local domain route distinguisher and route target. Likewise, when the gateway re-exports MAC-IP routes received from the local domain to the remote domain, it will remove the route distinguisher and the route target in the received routes and attach the remote domain route distinguisher and route target. The MAC-IP routes originated for the attached hosts that will be sent to the local domain will contain the local domain route distinguisher and route target.

Likewise, the MAC-IP routes originated for the attached hosts that will be sent to the remote domain will contain the remote domain route distinguisher and route target.

Alternatively, a route distinguisher can be configured to be used in both domains with *rd evpn domain all 10.255.1.1:10*. And a route target can be configured to be used in both domains with *route-target import export evpn domain all 64500:10*.

For IP-VRFs, the configured route-targets apply to both domains. IP-PREFIX routes crossing domains are covered in the EVPN L3 Multi-Domain section below.

On the leaf devices, there is no additional EVPN configuration and no knowledge of the remote EVPN domain(s).

The relevant VXLAN config on one of the peers forming GW1 is as follows. Note that the SVI configuration is required on gateway nodes for VLAN-VNIs with VXLAN routing.

```
vlan 10,20
!
interface Loopback1
  ip address 10.255.1.1/32
!
interface Vlan10
  vrf red
  ip address virtual 10.10.0.254/24
!
interface Vlan20
  vrf red
  ip address virtual 10.20.0.254/24
!
interface Vxlan1
  vxlan source-interface Loopback1
  vxlan virtual-router encapsulation mac-address mlag-system-id
  vxlan udp-port 4789
  vxlan vlan 10 vni 10010
  vxlan vlan 20 vni 10020
  vxlan vrf red vni 20000
!
ip virtual-router mac-address 00:00:80:00:00:00
```

EVPN L3 Multi-Domain

The existing EVPN L3 Gateway functionality interops with multi-domain EVPN VXLAN through a CLI extension. Without the *inter-domain* keyword, received IP-PREFIX routes **do not** cross domain boundaries. With the keyword, IP-PREFIX routes received from the local EVPN domain

are advertised to the remote EVPN domain (with next-hop self) and vice versa.

The following sample shows the BGP configuration on GW1 to enable IP-PREFIX routes to cross EVPN domains:

```
GW1(config-router-bgp)# show active
router bgp 300
...
!
address-family evpn
  neighbor LEAF_A activate
  neighbor LEAF_B activate
  neighbor GW2 activate
  neighbor GW2 domain remote
  neighbor default next-hop-self received-evpn-routes route-type ip-prefix inter-
domain
!
vrf red
  rd 10.255.1.1:0
  route-target import evpn 64500:20000
  route-target export evpn 64500:20000
  router-id 10.255.1.1
  redistribute connected
```

Note that this CLI applies **only** to IP-PREFIX routes crossing EVPN domains. GW1 will advertise IP-PREFIX routes to BGP peers in the same EVPN domain using next-hop unchanged. To advertise both intra-domain and inter-domain routes using next-hop self, configure the following:

```
address-family evpn
  neighbor LEAF_A activate
  neighbor LEAF_B activate
  neighbor GW2 activate
  neighbor GW2 domain remote
  neighbor default next-hop-self received-evpn-routes route-type ip-prefix
```

L3 Route Distinguishers and Route Targets

Note that IP-PREFIX routes do not use separate route distinguishers or route targets for local vs remote domain routes. The GW will not change the route distinguishers or the route targets in the IP-PREFIX routes when forwarding them across domain boundaries.

Use the L3 DCI Gateway as an Intra-DC Route Reflector

The L3-only DCI gateway can be configured as an intra-DC route reflector as well to reflect EVPN routes between the leaf VTEPs within the DC.

TCAM Profiles

When VXLAN routing is required, the `vxlan-routing` TCAM profile must be used.

Show Commands

The configuration of a remote EVPN domain virtually partitions the BGP EVPN RIB. As a result, the `domain (remote | local)` filter was added to BGP show commands. The following example filters for IMET routes in the remote domain. Present is the GW1 IMET route for VNI 10010 that is advertised to remote domain peers, and the IMET route for VNI 10010 that is received from GW2.

```
GW1#show bgp evpn route-type imet vni 10010 domain remote
      Network                Next Hop                Metric  LocPref Weight  Path
* >   RD: 10.255.1.1:10 imet 10.255.1.1 remote
      -                      -                      -       -       0       i
* >   RD: 10.255.2.1:10 imet 10.255.2.1 remote
      10.255.2.1           -                      100     0       400 i
```

Inspecting a MAC-IP route illustrates how the route distinguisher and export route-targets differ based on domain:

```
GW1#show bgp evpn route-type mac-ip 10.10.0.1 detail
BGP routing table information for VRF default
Router identifier 0.0.0.1, local AS number 300
BGP routing table entry
for mac-ip 4aed.4645.70de 10.10.0.1, Route Distinguisher: 10.255.1.1:10
Paths: 1 available
Local
- from - (0.0.0.0)
Origin IGP, metric -, localpref -, weight 0, valid, local, best
Extended Community: Route-Target-AS:64500:10 Route-Target-AS:64500:20000
TunnelEncap:tunnelTypeVxlan EvpnRouterMac:00:00:78:01:00:00
VNI: 10010 L3 VNI: 20000 ESI: 0000:0000:0000:0000:0000
BGP routing table entry f
or mac-ip 4aed.4645.70de 10.10.0.1 remote, Route Distinguisher: 10.255.1.1:1010
Paths: 1 available
Local
```

```
- from - (0.0.0.0)
Origin IGP, metric -, localpref -, weight 0, valid, local, best
Extended Community: Route-Target-AS:64500:20000 Route-Target-AS:64501:10
                    TunnelEncap:tunnelTypeVxlan EvpnRouterMac:00:00:78:01:00:00
VNI: 10010 L3 VNI: 20000 ESI: 0000:0000:0000:0000:0000
```

Note that in the local domain, the L2 route target is 64500:10 and the route distinguisher is 10.255.1.1:10, while in the remote domain the L2 route target is 64501:10 and the route distinguisher is 10.255.1.1:1010. The L3 route-target is the same in both domains.

Troubleshooting

- Verify that gateways in remote domains are configured as remote domain neighbors.
- Verify that routes are advertised to remote domain peers using:

```
GW1#show bgp neighbors PEER evpn advertised-routes domain remote
```

- Verify that MAC-VRFs stretched to remote domains contain remote domain import/export route-targets.
 - Whether to use the same route distinguisher and route-targets between domains is up to the network administrator.
- Verify VXLAN interface status with:

```
GW1#show interfaces vxlan 1
```

- Ensure that VLAN to VNI mappings are correct and that the flood-list for each VLAN contains all expected VTEPs.
 - Ensure that VTEP to VTEP bridging shows as enabled.
 - Ensure that VRF to VNI mapping is correct.
- Verify VXLAN flood-lists with:

```
GW1#show vxlan flood vtep
```

- Verify underlay redundancy between VTEPs in the same domain and between gateway nodes in different domains.
- If using *ip address virtual* and *ip virtual-router mac-address*, verify addresses are consistent across nodes.

Limitations

- Import route-maps do not apply to routes received from remote domain peers.
- Multi-VTEP MLAG is not supported as a gateway.
- IPv6 underlay is not supported.

- Gateways do not support VARP VTEP IP

References

- [EVPN L3 Gateway TOI](#)
- [RFC 9014: Interconnect Solution for Ethernet VPN \(EVPN\) Overlay Networks](#)