



Implementing DCI Layer 3 Gateway between MPLS-VPN and EVPN Data Center

This chapter module provides conceptual and configuration information for Data Center Interconnect (DCI) Layer 3 Gateway between MPLS-VPN and EVPN Data Center.


- [Data Center Interconnect between MPLS-VPN and EVPN-MPLS](#) , on page 1
- [Data Center Interconnect between MPLS-VPN and EVPN-VxLAN](#), on page 27

Data Center Interconnect between MPLS-VPN and EVPN-MPLS

This part provides conceptual and configuration information for Data Center Interconnect (DCI) Layer 3 Gateway with EVPN-MPLS on Cisco ASR 9000 Series Router.

DCI Layer 3 Gateway with EVPN-MPLS

You can use SR-EVPN for Data Center on routers for a spine-leaf architecture with edge devices such as border leaf. DCI L3 stitching allows Data Centers that run SR-EVPN to communicate with legacy and existing MPLS VPN (VPNv4) sites.



In this topology,

Leaf (ToR) – Router acts as both access switch and distributed PE. Leaf establishes BGP EVPN neighborhood with Spine route-reflector (RR). This router sends and receives prefixes from the DCI Gateway. Leaf ToR provides the following types of services:

- Regular L3 VRF configuration using subinterfaces to attach some CE devices. Traditional PE-CE scenario without EVPN configuration.
- L3 EVPN VRF using L2VPN configuration to attach multiple Data Centers services.

Leaf sends and receives prefixes from or to the DCI gateway:

- Leaf sends prefixes to DCI: Leaf re-originates local learned VRF subnet route as EVPN Route Type 5 with the EVPN RT (stitching-rt or regular RT), then sends to Spine RR. Spine RR sends prefixes to DCI gateway.
- Leaf receives prefixes from DCI: Leaf receives EVPN Route Type 5 from Spine RR that is re-originated at DCI gateway due to stitching between VPNv4 and EVPN. Leaf imports remote VPNv4 prefixes to local VRF matching VPNv4 RT (stitching-rt or regular RT).

Spine RR: Spine RR establishes BGP EVPN neighborhood with Leaf (ToR) and Edge DCI Gateway serving as Route-Reflector for EVPN prefixes between the devices in the Data Center. Leaf and DCI Gateway must be configured as clients of Spine RR.

Edge (DCI gateway): Edge (DCI gateway) acts as an edge router that allows communication between services connected at Leaf and CEs in legacy MPLS network architecture. The edge DCI gateway establishes BGP EVPN neighborship with Spine RR and remote PEs, or RR depending on legacy MPLS network architecture.

The edge DCI gateway sends and receives prefixes from or to the Data Center:

- DCI gateway receives prefixes from legacy MPLS VPNv4 network and sends prefixes to Leaf: DCI gateway receives L3VPN (VPNv4) routes from remote MPLS VPN (VPNv4) PE or RR depending on legacy MPLS network architecture matching the VPNv4 RT (stitching-rt or regular RT). Then re-originate these prefixes as EVPN Route Type 5 with the EVPN RT (stitching-rt or regular RT) advertising to Spine RR due to BGP EVPN neighbor with the Spine.
- DCI gateway receives prefixes from Leaf and sends prefixes to legacy MPLS VPNv4 network: DCI gateway receives EVPN Route Type 5 originated from Leaf (ToR) by Spine RR due to BGP EVPN neighbor with the Spine. Leaf and DCI gateway does not have a direct BGP neighborship. Then import the routes to local VRF matching the EVPN RT (stitching-rt or regular RT) and re-originate this prefix as VPNv4 router with the VPNv4 RT (stitching-rt or regular RT) and advertise to remote MPLS VPN (VPNv4) PE or RR depending on legacy MPLS network architecture.

Remote PE: Remote PE receives traditional MPLS L3VPN prefixes (VPNv4) by DCI Gateway or RR depending on legacy MPLS network architecture. You must have a unique Route-Distinguisher (RD) between remote PEs and DCI gateway to allow stitching re-originate prefixes from VPNv4 to EVPN at DCI Gateway.

Stitching RTs and Regular RTs can be assigned to any side, EVPN or VPNv4, irrespective of the address-family. Consider the following supported scenarios:


VPNv4-Regular RT and EVPN-Stitching RT

For each VRF on the DCI gateway, there are two sets of manually configured import and export route-targets for VPNv4 as a regular side and EVPN as a stitching side. Consider the following sets:

- Data Center Route-Targets for EVPN associated with EVPN BGP neighbor (Stitching RT).
- MPLS L3VPN Route-Targets for VPNv4 or VPNv6 associated with L3VPN BGP neighbor (Regular RT).

This separation of RTs enables the two sets of RTs to be independently configured. The RTs associated with the EVPN BGP neighbor require **stitching-rt** keyword under VRF configuration. The route-types associated with the L3VPN BGP neighbor do not require the keyword.

The following topology shows regular/normal and stitching side.



Route Targets

The RTs associated with the EVPN BGP neighbor are labelled as stitching RTs. The RTs associated with the L3VPN BGP neighbor are normal RTs.

Route Re-Origination

Consider control plane information propagation by the edge DCI gateway from the L3VPN (regular/normal side) to the Data Center (stitching side). Edge DCI gateway advertises to its BGP EVPN neighbor the routes that are re-originated after importing them from the L3VPN BGP neighbor. For this case of VPNv4 or VPNv6 routes being propagated to the BGP EVPN neighbors (Data Center neighbors), re-originating the routes refers to replacing the normal route-targets with the local route-target values (stitching-rt) associated with the BGP EVPN neighbors.

Route Address-Family and Encoded Address-Family

When an address-family is configured for a BGP neighbor, it means that the specified address-family routes encoded with the NLRI for that address-family are advertised to the neighbor. This does not hold for Data Center BGP neighbors because they use only EVPN address-family. Here, BGP neighbors advertise VPNv4 or VPNv6 unicast routes using the EVPN NLRI encoding. Thus, the encoded address-family and route address family can be possibly different. You can advertise the VPNv4 or VPNv6 address-family using the **advertise vpnv4 unicast** or **advertise vpnv6 unicast** command. For example, an EVPN address-family BGP neighbor configured with the **advertise vpnv4 unicast** command sends VPNv4 unicast routes in an EVPN encoded NLRI.

Local VPNv4 or VPNv6 Route Advertisement

On the edge DCI gateway, the locally sourced VPNv4 or VPNv6 routes (any CE directly connected not using L2VPN with BD/EVI/BVI, using only regular L3 VRF) can be advertised to the BGP EVPN neighbors with the normal route targets (RTs) configured for the VRF or the stitching RTs associated with the BGP EVPN neighbors. By default, these routes are advertised with the normal route targets. You can configure this local VPNv4 or VPNv6 route advertisements to be advertised with stitching RTs to the BGP EVPN neighbors by using the **advertise vpnv4 unicast local stitching-rt** or **advertise vpnv6 unicast local stitching-rt** command as required.

VPNv4 neighbors do not require any additional configuration. By default, these routes are advertised with the normal route-targets to BGP L3VPN neighbors.

Route Distinguishers

The Router Distinguisher (RD) associated per VRF must be unique per PE in the network. There are few available options to keep unique RD per device:

- Manual configuration: You must manually assign a unique value per device in the network. For example, in this scenario:
 - Leaf (ToR) = RD 1
 - Edge DCI Gateway = RD 2
 - Remote PE = RD 3
- Use **rd auto** command under VRF. To assign a unique route distinguisher for each router, you must ensure that each router has a unique BGP router-id. If so, the **rd auto** command assigns a Type 1 route distinguisher to the VRF using the following format: *ip-address:number*. The IP address is specified by the BGP router-id statement and the number (which is derived as an unused index in the 0 to 65535 range) is unique across the VRFs.



Note In a DCI deployment, for route re-originate with stitching-rt for a particular VRF, using the same Route Distinguisher (RD) between edge DCI gateway and MPLS-VPN PE or same RD between edge DCI gateway and Leaf (ToR) is not supported.

Configure VPNv4-Regular RT and EVPN-Stitching RT

This section describes tasks to configure VPNv4-Regular RT and EVPN-Stitching RT. Perform the following tasks to complete the configuration:

- Configure Leaf (ToR)
- Configure Spine-RR (Route Reflector)
- Configure Edge DCI Gateway
- Configure EVPN BGP neighbor and route advertisements
- Configure L3VPN BGP neighbor relationship and route advertisements

Configure Leaf (ToR)

Configure VRF in Leaf (ToR) at BGP-EVPN (Stitching Side) with Stitching-RT.

```
vrf data-center1
  address-family ipv4 unicast
    import route-target
      1:2 stitching                               // BGP - EVPN (Stitching Side)
  !
  export route-target
    1:2 stitching                               // BGP - EVPN (Stitching Side)
  !
router bgp 100
  neighbor 10.10.1.1                             // Spine Loopback IP Address
    address-family l2vpn evpn
      advertise vpnv4 unicast
      advertise vpnv6 unicast
  !
```



Note Advertise vpnv4/vpnv6 unicast enables local learned regular L3 VRF prefixes to be advertised as EVPN prefixes to BGP – EVPN neighbor. This means any local prefixes such as PE-CE without L2VPN with BD/EVI/BVI configuration. If all the services are pure EVPN with L2VPN with BD/EVI/BVI configuration these commands are not required.

Configure Spine-RR

Configure Spine RR with Leaf (ToR) and edge DCI gateway as RR client for AFI L2VPN EVPN. VRF configuration is not required.

```
// VRF Config is not required //

router bgp 100
  neighbor 10.10.2.1                             // Leaf (ToR) Loopback IP Address
    address-family l2vpn evpn
      route-reflector-client
  !
  neighbor 10.10.3.1                             // Edge DCI Gateway Loopback IP Address
    address-family l2vpn evpn
      route-reflector-client
  !
```

Configure Edge DCI Gateway

You can configure DCI with the same VRF as Leaf (ToR). Use the same RT as remote PE for L3VPN network or the same VRF if that is possible.

Configure VRF and Route Targets Import and Export rules

Perform the following steps to configure VRF and define route targets to be used for import and export of forwarding information.

```
vrf data-center1
  address-family ipv4 unicast
    import route-target
      1:1                                         // BGP - L3VPN (Regular/normal Side)
```

```

1:2 stitching                // BGP - EVPN (Stitching Side)
!
export route-target
1:1                          // BGP - L3VPN (Regular/normal Side)
1:2 stitching                // BGP - EVPN (Stitching Side)
!

```

Configure EVPN BGP Neighbor and Route Advertisements

Perform this task on the edge DCI gateway to configure BGP neighbor relationship and route advertisements with the EVPN BGP neighbor.

```

router bgp 100
 address-family l2vpn evpn
!
 neighbor 10.10.1.1          // Spine Loopback IP Address
 address-family l2vpn evpn
  import stitching-rt re-originate //Imp EVPN 1:2, reoriginate VPNv4 RT 1:1
  advertise vpnv4 unicast re-originated stitching-rt //Send routes EVPN 1:2
  advertise vpnv6 unicast re-originated stitching-rt //Send routes EVPN 1:2
!

```

Configure L3VPN BGP Neighbor Relationship and Route Advertisements

Perform the following steps to configure BGP neighbor relationship and route advertisements with the L3VPN BGP neighbor.

```

router bgp 100
 address-family vpnv4 unicast
!
 neighbor 10.10.1.1          // Spine Loopback IP Address
 address-family vpnv4 unicast // Same config for VPNv6
  import re-originate stitching-rt // Imp VPNv4 1:1, re-originate EVPN 1:2
  advertise vpnv4 unicast re-originated // Send routes VPNv4 RT 1:1
!

```

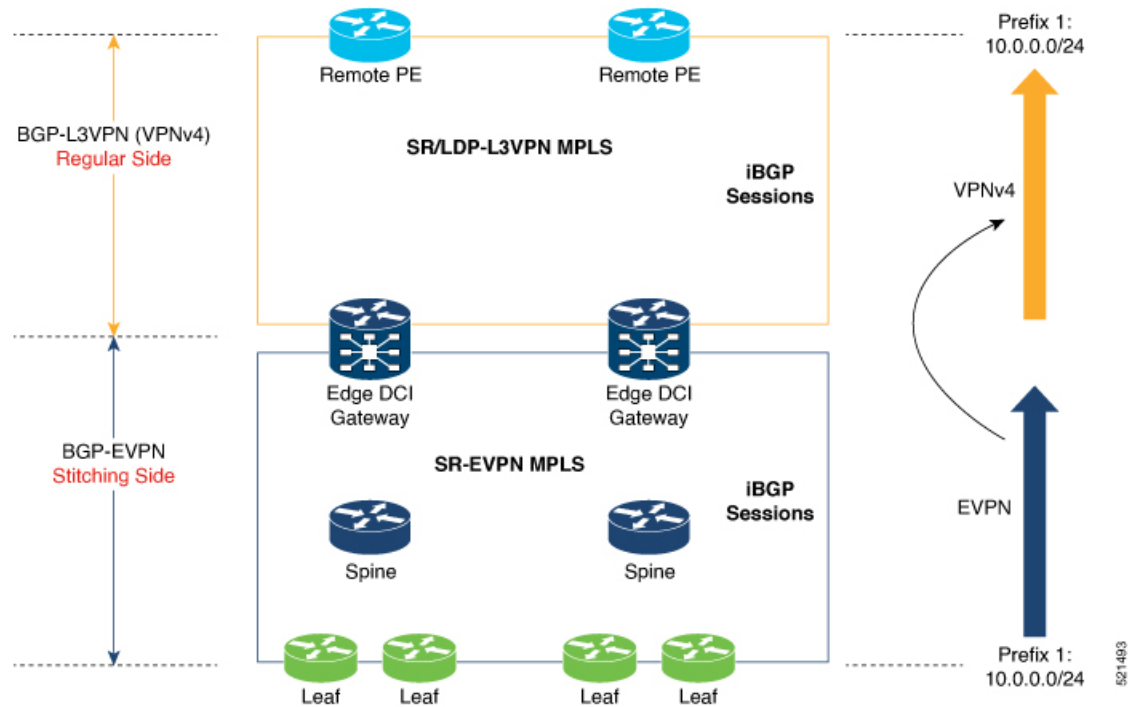
Configuration applies in two directions:

- Stitching from VPNv4 to EVPN routes. Prefixes received from MPLS L3VPN network and re-originated as EVPN prefixes towards Data Center Spine RR and Leaf (ToR).
1. Importing VPNv4 routes with import **re-originate stitching-rt** command under AFI VPNv4 UNICAST. This command imports routes using RT 1:1 and then reoriginate with BGP EVPN 1:2 **stitching-rt**.
 2. Advertising re-originated EVPN routes with VPNv4 RT with advertise **vpn4 unicast re-originated** command under AFI L2VPN EVPN. This command advertises routes from MPLS L3VPN network (VPNv4) to BGP EVPN neighbors inside Data Center (Spine RR and then Leaf (ToR)), re-originating these routes using BGP EVPN 1:2 **stitching-rt**.



- Stitching from EVPN to VPNv4 routes. Prefixes received from BGP-EVPN Data Center and re-originated as MPLS L3VPN prefixes towards VPNv4 RR or remote PE in L3VPN network.

1. Importing EVPN routes with import **stitching-rt re-originate** command under AFI L2VPN EVPN. This command imports routes using RT 1:2 **stitching-rt** and then re-originate with VPNv4 regular/normal VPNv4 RT 1:1.
2. Advertising re-originated EVPN routes with VPNv4 RT with **advertise vpnv4 unicast re-originated** command under AFI VPNv4 UNICAST. This command advertises routes from EVPN Data Center to VPNv4 RR or remote PEs, re-originating these routes using regular/normal VPNv4 RT 1:1.



Verification of Edge DCI Gateway Configuration

```
Router# show bgp l2vpn evpn
```

```
Fri Aug 21 00:24:10.773 PDT
BGP router identifier 30.30.30.30, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 16
BGP NSR Initial initsync version 1 (Reached)
BGP NSR/ISSU Sync-Group versions 16/0
BGP scan interval 60 secs
```

```
Status codes: s suppressed, d damped, h history, * valid, > best
i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
Network Next Hop Metric LocPrf Weight Path
Route Distinguisher: 100:1
*>i [2] [10000] [48] [0226.51bd.c81c] [32] [200::1001]/232
11.0.0.1 100 0 i
*>i [2] [10000] [48] [0226.51bd.c81c] [32] [200:1::1001]/232
11.0.0.1 100 0 i
*>i [2] [10000] [48] [0226.51bd.c81c] [32] [200.1.1.1]/136
11.0.0.1 100 0 i
*>i [2] [10000] [48] [0226.51bd.c81c] [32] [200.1.1.2]/136
11.0.0.1 100 0 i
*>i [5] [4231] [32] [100.1.1.1]/80
11.0.0.1 100 0 i
*>i [5] [4231] [32] [100.1.1.2]/80
11.0.0.1 100 0 i
*>i [5] [4231] [112] [fec0::1001]/176
11.0.0.1 100 0 i
```

```
*>i[5][4232][112][fec0::1:1001]/176
    11.0.0.1                100        0 i
```

Processed 8 prefixes, 8 paths

```
Router# show bgp l2vpn evpn rd 100:1 [5][4231][112][fec0::1001]/176 detail
```

```
Fri Aug 21 00:34:43.747 PDT
BGP routing table entry for [5][4231][112][fec0::1001]/176, Route Distinguisher: 100:1
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          5         5
  Flags: 0x04040001+0x00000000;
Last Modified: Aug 21 00:16:58.000 for 00:17:46
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Flags: 0x4000600025060005, import: 0x3f
  Not advertised to any peer
  Local
    11.0.0.1 (metric 2) from 20.0.0.1 (11.0.0.1)
    Received Label 16001
    Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
reoriginate, not-in-vrf
    Received Path ID 0, Local Path ID 1, version 5
    Extended community: Flags 0x6: RT:1:1
    Originator: 11.0.0.1, Cluster list: 20.20.20.20
    EVPN ESI: ffff.ffff.ffff.ffff.ff01, Gateway Address : fec0::254
```

```
Router# show bgp l2vpn evpn neighbors 20.0.0.1 detail
```

```
Fri Aug 21 00:25:37.383 PDT

BGP neighbor is 20.0.0.1
  Remote AS 100, local AS 100, internal link
  Remote router ID 20.20.20.20
  BGP state = Established, up for 00:08:58
  NSR State: NSR Ready
  Last read 00:00:34, Last read before reset 00:00:00
  Hold time is 180, keepalive interval is 60 seconds
  Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
  Last write 00:00:36, attempted 19, written 19
  Second last write 00:01:36, attempted 143, written 143
  Last write before reset 00:00:00, attempted 0, written 0
  Second last write before reset 00:00:00, attempted 0, written 0
  Last write pulse rcvd Aug 21 00:25:03.667 last full not set pulse count 33
  Last write pulse rcvd before reset 00:00:00
  Socket not armed for io, armed for read, armed for write
  Last write thread event before reset 00:00:00, second last 00:00:00
  Last KA expiry before reset 00:00:00, second last 00:00:00
  Last KA error before reset 00:00:00, KA not sent 00:00:00
  Last KA start before reset 00:00:00, second last 00:00:00
  Precedence: internet
  Non-stop routing is enabled
  Entered Neighbor NSR TCP mode:
    TCP Initial Sync :          Aug 21 00:18:07.291
    TCP Initial Sync Phase Two : Aug 21 00:18:07.319
    TCP Initial Sync Done :      Aug 21 00:18:08.334
  Multi-protocol capability received
  Neighbor capabilities:
    Adv          Rcvd
  Route refresh:      Yes      Yes
  4-byte AS:         Yes      Yes
  Address family VPNv4 Unicast: Yes      No
  Address family VPNv6 Unicast: Yes      No
```

```

Address family L2VPN EVPN:      Yes      Yes
Message stats:
InQ depth: 0, OutQ depth: 0
      Last_Sent      Sent      Last_Rcvd      Rcvd
Open:      Aug 21 00:16:38.087      1      Aug 21 00:16:40.123      1
Notification:  ---      0      ---      0
Update:      Aug 21 00:24:01.421      9      Aug 21 00:24:03.652      13
Keepalive:      Aug 21 00:25:01.434      8      Aug 21 00:25:03.667      9
Route_Refresh: Aug 21 00:24:01.377      3      ---      0
Total:      21      23
Minimum time between advertisement runs is 0 secs
Inbound message logging enabled, 3 messages buffered
Outbound message logging enabled, 3 messages buffered

```

```

For Address Family: VPNv4 Unicast
BGP neighbor version 35
Update group: 0.3 Filter-group: 0.1 No Refresh request being processed
Advertise Reorigination Enabled
Advertise AFI EoR can be sent
Route refresh request: received 0, sent 0
0 accepted prefixes, 0 are bestpaths
Cumulative no. of prefixes denied: 0.
Prefix advertised 4, suppressed 0, withdrawn 0
Maximum prefixes allowed 2097152
Threshold for warning message 75%, restart interval 0 min
AIGP is enabled
An EoR was not received during read-only mode
Last ack version 35, Last synced ack version 35
Outstanding version objects: current 0, max 1
Additional-paths operation: None
Send Multicast Attributes

For Address Family: VPNv6 Unicast
BGP neighbor version 29
Update group: 0.3 Filter-group: 0.1 No Refresh request being processed
Advertise Reorigination Enabled
Advertise AFI EoR can be sent
Route refresh request: received 0, sent 0
0 accepted prefixes, 0 are bestpaths
Cumulative no. of prefixes denied: 0.
Prefix advertised 0, suppressed 0, withdrawn 0
Maximum prefixes allowed 1048576
Threshold for warning message 75%, restart interval 0 min
AIGP is enabled
An EoR was not received during read-only mode
Last ack version 29, Last synced ack version 29
Outstanding version objects: current 0, max 0
Additional-paths operation: None
Send Multicast Attributes
Advertise VPNv4 routes enabled with Reoriginate,Local with stitching-RT option

```

```

For Address Family: L2VPN EVPN
BGP neighbor version 18
Update group: 0.2 Filter-group: 0.1 No Refresh request being processed
Route refresh request: received 0, sent 3
8 accepted prefixes, 8 are bestpaths
Cumulative no. of prefixes denied: 0.
Prefix advertised 4, suppressed 0, withdrawn 6
Maximum prefixes allowed 2097152
Threshold for warning message 75%, restart interval 0 min
AIGP is enabled
An EoR was received during read-only mode
Last ack version 18, Last synced ack version 18
Outstanding version objects: current 0, max 2

```

```

Additional-paths operation: None
Send Multicast Attributes
Advertise VPNv4 routes enabled with Reoriginate, option
Advertise VPNv6 routes is enabled with Reoriginate, option
Import Stitching is enabled for this neighbor address-family
Import Reoriginate is enabled for this neighbor address-family

Connections established 1; dropped 0
Local host: 30.0.0.1, Local port: 59405, IF Handle: 0x00000000
Foreign host: 20.0.0.1, Foreign port: 179
Last reset 00:00:00

```

At the end of each one AFI VPNv4, VPNv6, or L2VPN EVPN, you can see import and advertise information based on the configuration.

```
Router# show bgp sessions
```

```
Fri Aug 21 00:25:57.216 PDT
```

Neighbor	VRF	Spk	AS	InQ	OutQ	NBRState	NSRState
20.0.0.1	default	0	100	0	0	Established	NSR Ready[PP]
32.0.0.2	default	0	200	0	0	Established	NSR Ready

```
Router# show bgp vpnv4 unicast
```

```

Fri Aug 21 00:28:41.253 PDT
BGP router identifier 30.30.30.30, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 39
BGP NSR Initial initsync version 4 (Reached)
BGP NSR/ISSU Sync-Group versions 39/0
BGP scan interval 60 secs

```

```

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete

```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 1:1					
*> 1.1.1.0/24	32.0.0.2			0 200 300	i
*> 1.1.2.0/24	32.0.0.2			0 200 300	i
Route Distinguisher: 30.30.30.30:0 (default for vrf foo)					
*> 1.1.1.0/24	32.0.0.2			0 200 300	i
*> 1.1.2.0/24	32.0.0.2			0 200 300	i
*>i100.1.1.1/32	11.0.0.1		100	0	i
*>i100.1.1.2/32	11.0.0.1		100	0	i
*>i200.1.1.1/32	11.0.0.1		100	0	i
*>i200.1.1.2/32	11.0.0.1		100	0	i

```
Router# show bgp vpnv4 unicast rd 30.30.30.30:0 1.1.1.0/24 detail
```

```

Fri Aug 21 00:28:57.824 PDT
BGP routing table entry for 1.1.1.0/24, Route Distinguisher: 30.30.30.30:0
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          26        26
  Flags: 0x04103001+0x00000000;
Last Modified: Aug 21 00:24:01.000 for 00:04:58
Paths: (1 available, best #1)
  Advertised to peers (in unique update groups):

```

```

20.0.0.1
Path #1: Received by speaker 0
Flags: 0x4000c00005060001, import: 0x80
Advertised to peers (in unique update groups):
  20.0.0.1
200 300
  32.0.0.2 from 32.0.0.2 (40.40.40.40)
    Received Label 24001
    Origin IGP, localpref 100, valid, external, best, group-best, import-candidate,
imported, reoriginated with stitching-rt
    Received Path ID 0, Local Path ID 1, version 26
    Extended community: RT: 1:2
    Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 1:1

```

Router# **show bgp vrf foo**

```

Fri Aug 21 00:24:36.523 PDT
BGP VRF foo, state: Active
BGP Route Distinguisher: 30.30.30.30:0
VRF ID: 0x60000002
BGP router identifier 30.30.30.30, local AS number 100
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000011 RD version: 35
BGP main routing table version 35
BGP NSR Initial initsync version 4 (Reached)
BGP NSR/ISSU Sync-Group versions 31/0

```

```

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete

```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 30.30.30.30:0 (default for vrf foo)					
*> 1.1.1.0/24	32.0.0.2			0 200 300	i
*> 1.1.2.0/24	32.0.0.2			0 200 300	i
*>i100.1.1.1/32	11.0.0.1	100		0	i
*>i100.1.1.2/32	11.0.0.1	100		0	i
*>i200.1.1.1/32	11.0.0.1	100		0	i
*>i200.1.1.2/32	11.0.0.1	100		0	i

Processed 6 prefixes, 6 paths

Router# **show bgp vrf foo ipv4 unicast 100.1.1.1/32 detail**

```

Mon Dec 8 23:24:50.243 PST
BGP routing table entry for 100.1.1.1/32, Route Distinguisher:
30.30.30.30:0
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          43        43
    Local Label: 24001 (with rewrite);
    Flags: 0x05081001+0x00000200;
Last Modified: Dec 8 18:04:21.000 for 05:20:30
Paths: (1 available, best #1)
  Advertised to PE peers (in unique update groups):
    32.0.0.2
    Path #1: Received by speaker 0
    Flags: 0x400061000d060005, import: 0x80
    Advertised to PE peers (in unique update groups):
      32.0.0.2
  Local
    11.0.0.1 (metric 2) from 20.0.0.1 (11.0.0.1)
      Received Label 1234

```

```

Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
imported, reoriginated
Received Path ID 0, Local Path ID 1, version 43
Extended community: RT:1:2
Originator: 11.0.0.1, Cluster list: 20.20.20.20
Source AFI: L2VPN EVPN, Source VRF: default, Source Route Distinguisher: 100:1

```

Router# **show bgp vpnv4 unicast update-group**

Fri Aug 21 00:27:57.910 PDT

Update group for VPNv4 Unicast, index 0.1:

```

Attributes:
  Outbound policy: pass
  First neighbor AS: 200
  Send communities
  Send GSHUT community if originated
  Send extended communities
  4-byte AS capable
  Send Re-originated VPN routes
  Send multicast attributes
  Minimum advertisement interval: 30 secs
Update group desynchronized: 0
Sub-groups merged: 0
Number of refresh subgroups: 0
Messages formatted: 8, replicated: 8
All neighbors are assigned to sub-group(s)
  Neighbors in sub-group: 0.2, Filter-Groups num:1
  Neighbors in filter-group: 0.2(RT num: 0)
  32.0.0.2

```

Update group for VPNv4 Unicast, index 0.3:

```

Attributes:
  Neighbor sessions are IPv4
  Internal
  Common admin
  First neighbor AS: 100
  Send communities
  Send GSHUT community if originated
  Send extended communities
  4-byte AS capable
  Send AIGP
  Send Re-originated VPN routes
  Send multicast attributes
  Minimum advertisement interval: 0 secs
Update group desynchronized: 0
Sub-groups merged: 0
Number of refresh subgroups: 0
Messages formatted: 2, replicated: 2
All neighbors are assigned to sub-group(s)
  Neighbors in sub-group: 0.1, Filter-Groups num:1
  Neighbors in filter-group: 0.1(RT num: 0)
  20.0.0.1

```

Router# **show bgp l2vpn evpn update-group**

Fri Aug 21 00:27:42.786 PDT

Update group for L2VPN EVPN, index 0.2:

```

Attributes:
  Neighbor sessions are IPv4
  Internal
  Common admin

```

```
First neighbor AS: 100
Send communities
Send GSHUT community if originated
Send extended communities
4-byte AS capable
Send AIGP
Send multicast attributes
Minimum advertisement interval: 0 secs
Update group desynchronized: 0
Sub-groups merged: 0
Number of refresh subgroups: 0
Messages formatted: 4, replicated: 4
All neighbors are assigned to sub-group(s)
  Neighbors in sub-group: 0.1, Filter-Groups num:1
    Neighbors in filter-group: 0.1(RT num: 0)
      20.0.0.1
```


EVPN-Regular RT and VPNv4-Stitching RT

For each VRF on the DCI gateway, there are two sets of manually configured import and export route-targets for EVPN as regular side and VPNv4 as stitching side. Consider the following sets:

- Data Center Route-Targets for EVPN associated with EVPN BGP neighbor (Regular RT)
- MPLS L3VPN Route-Targets for VPNv4 or VPNv6 associated with L3VPN BGP neighbor (Stitching RT)

This separation of RTs enables the two sets of RTs to be independently configured. The RTs associated with the EVPN BGP neighbor does not require the keyword, it remains a normal configuration. The RTs associated with the L3VPN BGP neighbor require **stitching-rt** keyword under VRF configuration.

The following topology shows regular or normal and stitching side.



Route Targets

The RTs associated with the L3VPN BGP neighbor are labelled as stitching RTs. The RTs associated with the EVPN BGP neighbor are normal RTs.

Route Re-Origination

Consider control plane information propagation by the edge DCI gateway from the L3VPN (stitching side) to the Data Center (regular/normal side). Edge DCI gateway advertises to its BGP EVPN neighbor the routes that are re-originated after importing them from the L3VPN BGP neighbor. For this case of VPNv4 or VPNv6 routes being propagated to the BGP EVPN neighbors (Data Center neighbors), re-originating the routes refers to replacing the stitching route-targets with the local route-target values (regular/normal) associated with the BGP EVPN neighbors.

Local VPNv4 or VPNv6 Route Advertisement

On the edge DCI gateway, the locally sourced VPNv4 or VPNv6 routes (any CE directly connected not using L2VPN with BD/EVI/BVI, using only regular L3 VRF) can be advertised to the BGP EVPN neighbors with the normal route targets (RTs) configured for the VRF or the stitching RTs associated with the BGP EVPN neighbors. By default, these routes are advertised with the normal route targets to the BGP EVPN Neighbors (regular/normal side)

VPNv4 neighbors require an additional configuration on the existing legacy VRF to allow these routes to be advertised to VPNv4 RR or remote PEs. Configure **stitching-rt** keyword on existing VRF under import/export RT.

Route Distinguishers

The Router Distinguisher (RD) associated per VRF must be unique per PE in the network. There are few available options to keep unique RD per device:

- Manual configuration: You must manually assign a unique value per device in the network. For example, in this scenario:
 - Leaf (ToR) = RD 1
 - Edge DCI Gateway = RD 2
 - Remote PE = RD 3
- Use **rd auto** command under VRF. To assign a unique route distinguisher for each router, you must ensure that each router has a unique BGP router-id. If so, the **rd auto** command assigns a Type 1 route distinguisher to the VRF using the following format: *ip-address:number*. The IP address is specified by the BGP router-id statement and the number (which is derived as an unused index in the 0 to 65535 range) is unique across the VRFs.



Note In a DCI deployment, for route re-originate with stitching-rt for a particular VRF, using the same Route Distinguisher (RD) between edge DCI gateway and MPLS-VPN PE or same RD between edge DCI gateway and Leaf (ToR) is not supported.

Configure EVPN-Regular RT and VPNv4-Stitching RT

This section describes tasks to configure EVPN-Regular RT and VPNv4-Stitching RT. Perform the following tasks to complete the configuration:

- Configure Leaf (ToR)
- Configure Spine-RR (Route Reflector)
- Configure Edge DCI Gateway
- Configure EVPN BGP neighbor and route advertisements
- Configure L3VPN BGP neighbor relationship and route advertisements

Configure Leaf (ToR)

Configure VRF in Leaf (ToR) at BGP-EVPN (regular/normal side). Note that the **stitching-rt** keyword is not required.

```
vrf data-center1
  address-family ipv4 unicast
    import route-target
      1:2                               // BGP - EVPN (Regular/Normal Side)
  !
```

```

export route-target
  1:2 // BGP - EVPN (Regular/Normal Side)
!
router bgp 100
  neighbor 10.10.1.1 // Spine Loopback IP Address
    address-family l2vpn evpn
      advertise vpnv4 unicast
      advertise vpnv6 unicast
!

```



Note Advertise vpnv4/vpnv6 unicast enables local learned regular L3 VRF prefixes to be advertised as EVPN prefixes to BGP-EVPN neighbor. This means any local prefixes such as PE-CE without L2VPN with BD/EVI/BVI configuration. If all the services are pure EVPN with L2VPN with BD/EVI/BVI configuration these commands are not required.

Configure Spine-RR

Configure Spine RR with Leaf (ToR) and edge DCI gateway as RR client for AFI L2VPN EVPN.

```

// VRF Config is not required //
router bgp 100
  neighbor 10.10.2.1 // Leaf (ToR) Loopback IP Address
    address-family l2vpn evpn
      route-reflector-client
  !
  neighbor 10.10.3.1 // Edge DCI Gateway Loopback IP Address
    address-family l2vpn evpn
      route-reflector-client
!

```

Configure Edge DCI Gateway

You can configure DCI with the same VRF as Leaf (ToR). Use the same RT as remote PE for L3VPN network or the same VRF if that is possible.

Configure VRF and Route Targets Import and Export rules

Perform the following steps to configure VRF and define route targets to be used for import and export of forwarding information.

```

vrf data-center1
  address-family ipv4 unicast
    import route-target
      1:1 stitching // BGP - L3VPN (Stitching Side)
      1:2 // BGP - EVPN (Regular/normal Side)
  !
  export route-target
    1:1 stitching // BGP - L3VPN (Stitching Side)
    1:2 // BGP - EVPN (Regular/normal Side)
!

```

Configure EVPN BGP Neighbor and Route Advertisements

Perform this task on the edge DCI gateway to configure BGP neighbor relationship and route advertisements with the EVPN BGP neighbor.

```

router bgp 100
 address-family l2vpn evpn
 !
 neighbor 10.10.1.1          // Spine Loopback IP Address
  address-family l2vpn evpn
   import re-originate stitching-rt //Imp EVPN RT 1:2, re-originate VPNv4 1:1
   advertise vpnv4 unicast re-originated //Send routes VPNv4 RT 1:1
 !

```

Configure L3VPN BGP Neighbor Relationship and Route Advertisements

Perform the following steps to configure BGP neighbor relationship and route advertisements with the L3VPN BGP neighbor.

```

router bgp 100
 address-family vpnv4 unicast
 !
 neighbor 10.10.1.1          // Spine Loopback IP Address
  address-family vpnv4 unicast // Same config for VPNv6
   import stitching-rt re-originate // Imp VPNv4 1:1, reoriginate EVPN 1:2
   advertise vpnv4 unicast re-originated stitching-rt //Send Routes EVPN 1:2
   advertise vpnv6 unicast re-originated stitching-rt //Send Routes EVPN 1:2
 !

```




Note The **stitching-rt** applies for L3VPN RT and EVPN RT does not require the **stitching-rt** for this use case.

If there are existing regular local L3 VRF without L2VPN with BD/EVI/BVI in these devices, configure import/export Stitching-RT for existing VRFs to advertise to L3VPN RR or remote PEs.


Configuration applies in two directions:

- Stitching from VPNv4 to EVPN routes. Prefixes received from MPLS L3VPN network and re-originated as EVPN prefixes towards Data Center Spine RR and Leaf (ToR)
 1. Importing VPNv4 routes with **import stitching-rt re-originate** command under AFI VPNv4 UNICAST. This command imports routes using RT 1:1 stitching-rt and then re-originate with BGP EVPN 1:2
 2. Advertising re-originated EVPN routes with VPNv4 RT with **advertise vpnv4 unicast re-originated** command under AFI L2VPN EVPN. This command advertises routes from MPLS L3VPN network (VPNv4) to BGP EVPN neighbors inside Data Center (Spine RR and then Leaf (ToR)), re-originating these routes using BGP EVPN 1:2.



- Stitching from EVPN to VPNv4 routes. Prefixes received from BGP-EVPN Data Center and re-originated as MPLS L3VPN prefixes towards VPNv4 RR or remote PE in L3VPN network.

1. Importing EVPN routes with **import re-originate stitching-rt** command under AFI L2VPN EVPN. This command imports routes using RT 1:2 and then re-originate with VPNv4 RT 1:1 **stitching-rt**.
2. Advertising re-originated EVPN routes with VPNv4 RT with **advertise vpnv4 unicast re-originated stitching-rt** command under AFI VPNv4 UNICAST. This command advertises routes from EVPN Data Center to VPNv4 RR or remote PEs, re-originating these routes using VPNv4 RT 1:1 **stitching-rt**



Verification of Edge DCI Gateway Configuration

Router# show bgp l2vpn evpn

```

Fri Aug 21 00:24:10.773 PDT
BGP router identifier 30.30.30.30, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 16
BGP NSR Initial initsync version 1 (Reached)
BGP NSR/ISSU Sync-Group versions 16/0
BGP scan interval 60 secs
    
```

Status codes: s suppressed, d damped, h history, * valid, > best
 i - internal, r RIB-failure, S stale, N Nexthop-discard
 Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 100:1					
*>i [2] [10000] [48] [0226.51bd.c81c] [32] [200::1001]/232	11.0.0.1	100	0	i	
*>i [2] [10000] [48] [0226.51bd.c81c] [32] [200:1::1001]/232	11.0.0.1	100	0	i	
*>i [2] [10000] [48] [0226.51bd.c81c] [32] [200.1.1.1]/136	11.0.0.1	100	0	i	
*>i [2] [10000] [48] [0226.51bd.c81c] [32] [200.1.1.2]/136	11.0.0.1	100	0	i	
*>i [5] [4231] [32] [100.1.1.1]/80	11.0.0.1	100	0	i	
*>i [5] [4231] [32] [100.1.1.2]/80	11.0.0.1	100	0	i	
*>i [5] [4231] [112] [fec0::1001]/176	11.0.0.1	100	0	i	
*>i [5] [4232] [112] [fec0::1:1001]/176	11.0.0.1	100	0	i	

```

11.0.0.1          100      0 i

Processed 8 prefixes, 8 paths

Router# show bgp l2vpn evpn rd 100:1 [5][4231][112][fec0::1001]/176 detail

Fri Aug 21 00:34:43.747 PDT
BGP routing table entry for [5][4231][112][fec0::1001]/176, Route Distinguisher: 100:1
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          5         5
  Flags: 0x04040001+0x00000000;
Last Modified: Aug 21 00:16:58.000 for 00:17:46
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Flags: 0x4000600025060005, import: 0x3f
  Not advertised to any peer
Local
  11.0.0.1 (metric 2) from 20.0.0.1 (11.0.0.1)
  Received Label 16001
  Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
reoriginate stitching-rt, not-in-vrf
  Received Path ID 0, Local Path ID 1, version 5
  Extended community: Flags 0x6: RT:1:1
  Originator: 11.0.0.1, Cluster list: 20.20.20.20
  EVPN ESI: ffff.ffff.ffff.ffff.ff01, Gateway Address : fec0::254

```

The main difference with scenario 1 is that the prefixes have a **reoriginate stitching-rt** keyword on the output versus scenario 1 having just reoriginate.

```

Router# show bgp l2vpn evpn neighbors 20.0.0.1 detail

Fri Aug 21 00:25:37.383 PDT

BGP neighbor is 20.0.0.1
Remote AS 100, local AS 100, internal link
Remote router ID 20.20.20.20
BGP state = Established, up for 00:08:58
NSR State: NSR Ready
Last read 00:00:34, Last read before reset 00:00:00
Hold time is 180, keepalive interval is 60 seconds
Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
Last write 00:00:36, attempted 19, written 19
Second last write 00:01:36, attempted 143, written 143
Last write before reset 00:00:00, attempted 0, written 0
Second last write before reset 00:00:00, attempted 0, written 0
Last write pulse rcvd Aug 21 00:25:03.667 last full not set pulse count 33
Last write pulse rcvd before reset 00:00:00
Socket not armed for io, armed for read, armed for write
Last write thread event before reset 00:00:00, second last 00:00:00
Last KA expiry before reset 00:00:00, second last 00:00:00
Last KA error before reset 00:00:00, KA not sent 00:00:00
Last KA start before reset 00:00:00, second last 00:00:00
Precedence: internet
Non-stop routing is enabled
Entered Neighbor NSR TCP mode:
  TCP Initial Sync :          Aug 21 00:18:07.291
  TCP Initial Sync Phase Two : Aug 21 00:18:07.319
  TCP Initial Sync Done :      Aug 21 00:18:08.334
Multi-protocol capability received
Neighbor capabilities:
Route refresh:          Yes      Rcvd      Yes
4-byte AS:              Yes      Rcvd      Yes
Address family VPNv4 Unicast: Yes      Rcvd      No

```

```

Address family VPNv6 Unicast:  Yes      No
Address family L2VPN EVPN:    Yes      Yes
Message stats:
InQ depth: 0, OutQ depth: 0
      Last_Sent      Sent  Last_Rcvd      Rcvd
Open:      Aug 21 00:16:38.087      1  Aug 21 00:16:40.123      1
Notification:  ---      0  ---      0
Update:      Aug 21 00:24:01.421      9  Aug 21 00:24:03.652      13
Keepalive:    Aug 21 00:25:01.434      8  Aug 21 00:25:03.667      9
Route_Refresh: Aug 21 00:24:01.377      3  ---      0
Total:                               21      23
Minimum time between advertisement runs is 0 secs
Inbound message logging enabled, 3 messages buffered
Outbound message logging enabled, 3 messages buffered

```

```

For Address Family: VPNv4 Unicast
BGP neighbor version 35
Update group: 0.3 Filter-group: 0.1 No Refresh request being processed
Advertise Reorigination Enabled
Advertise AFI EoR can be sent
Route refresh request: received 0, sent 0
0 accepted prefixes, 0 are bestpaths
Cumulative no. of prefixes denied: 0.
Prefix advertised 4, suppressed 0, withdrawn 0
Maximum prefixes allowed 2097152
Threshold for warning message 75%, restart interval 0 min
AIGP is enabled
An EoR was not received during read-only mode
Last ack version 35, Last synced ack version 35
Outstanding version objects: current 0, max 1
Additional-paths operation: None
Send Multicast Attributes

```

```

For Address Family: VPNv6 Unicast
BGP neighbor version 29
Update group: 0.3 Filter-group: 0.1 No Refresh request being processed
Advertise Reorigination Enabled
Advertise AFI EoR can be sent
Route refresh request: received 0, sent 0
0 accepted prefixes, 0 are bestpaths
Cumulative no. of prefixes denied: 0.
Prefix advertised 0, suppressed 0, withdrawn 0
Maximum prefixes allowed 1048576
Threshold for warning message 75%, restart interval 0 min
AIGP is enabled
An EoR was not received during read-only mode
Last ack version 29, Last synced ack version 29
Outstanding version objects: current 0, max 0
Additional-paths operation: None
Send Multicast Attributes
Advertise VPNv4 routes enabled with Reoriginate,Local with stitching-RT option

```

```

For Address Family: L2VPN EVPN
BGP neighbor version 18
Update group: 0.2 Filter-group: 0.1 No Refresh request being processed
Route refresh request: received 0, sent 3
8 accepted prefixes, 8 are bestpaths
Cumulative no. of prefixes denied: 0.
Prefix advertised 4, suppressed 0, withdrawn 6
Maximum prefixes allowed 2097152
Threshold for warning message 75%, restart interval 0 min
AIGP is enabled
An EoR was received during read-only mode
Last ack version 18, Last synced ack version 18

```

```

Outstanding version objects: current 0, max 2
Additional-paths operation: None
Send Multicast Attributes
Advertise VPNv4 routes enabled with Reoriginate, option
Advertise VPNv6 routes is enabled with Reoriginate, option
Import Reoriginate is enabled for this neighbor address-family

Connections established 1; dropped 0
Local host: 30.0.0.1, Local port: 59405, IF Handle: 0x00000000
Foreign host: 20.0.0.1, Foreign port: 179
Last reset 00:00:00

```

At the end of each one AFI VPNv4, VPNv6, or L2VPN EVPN, you can see import and advertise information based on the configuration.

Based on whether stitching-side or regular side, import stitching applies on VPNv4 AFI. In Scenario 1 you can see import stitching under L2VPN EVPN.

```
Router# show bgp sessions
```

```
Fri Aug 21 00:25:57.216 PDT
```

Neighbor	VRF	Spk	AS	InQ	OutQ	NBRState	NSRState
20.0.0.1	default	0	100	0	0	Established	NSR Ready[PP]
32.0.0.2	default	0	200	0	0	Established	NSR Ready

```
Router# show bgp vpnv4 unicast
```

```

Fri Aug 21 00:28:41.253 PDT
BGP router identifier 30.30.30.30, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 39
BGP NSR Initial initsync version 4 (Reached)
BGP NSR/ISSU Sync-Group versions 39/0
BGP scan interval 60 secs

```

```

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete

```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 1:1					
*> 1.1.1.0/24	32.0.0.2		0	200	300 i
*> 1.1.2.0/24	32.0.0.2		0	200	300 i
Route Distinguisher: 30.30.30.30:0 (default for vrf foo)					
*> 1.1.1.0/24	32.0.0.2		0	200	300 i
*> 1.1.2.0/24	32.0.0.2		0	200	300 i
*>i100.1.1.1/32	11.0.0.1		100	0	i
*>i100.1.1.2/32	11.0.0.1		100	0	i
*>i200.1.1.1/32	11.0.0.1		100	0	i
*>i200.1.1.2/32	11.0.0.1		100	0	i

In origin IGP line, you can see that the prefix was reoriginated with regular-RT.

```
Router# show bgp vpnv4 unicast rd 30.30.30.30:0 1.1.1.0/24 detail
```

```

Fri Aug 21 00:28:57.824 PDT
BGP routing table entry for 1.1.1.0/24, Route Distinguisher: 30.30.30.30:0
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          26        26
  Flags: 0x04103001+0x00000000;
Last Modified: Aug 21 00:24:01.000 for 00:04:58

```



```

Paths: (1 available, best #1)
  Advertised to peers (in unique update groups):
    20.0.0.1
  Path #1: Received by speaker 0
  Flags: 0x4000c00005060001, import: 0x80
  Advertised to peers (in unique update groups):
    20.0.0.1
  200 300
    32.0.0.2 from 32.0.0.2 (40.40.40.40)
      Received Label 24001
      Origin IGP, localpref 100, valid, external, best, group-best, import-candidate,
imported, reoriginated
      Received Path ID 0, Local Path ID 1, version 26
      Extended community: RT: 1:2
      Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 1:1

```

Router# **show bgp vrf foo**

```

Fri Aug 21 00:24:36.523 PDT
BGP VRF foo, state: Active
BGP Route Distinguisher: 30.30.30.30:0
VRF ID: 0x60000002
BGP router identifier 30.30.30.30, local AS number 100
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000011 RD version: 35
BGP main routing table version 35
BGP NSR Initial initsync version 4 (Reached)
BGP NSR/ISSU Sync-Group versions 31/0

```

```

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete

```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 30.30.30.30:0 (default for vrf foo)					
*> 1.1.1.0/24	32.0.0.2			0 200 300	i
*> 1.1.2.0/24	32.0.0.2			0 200 300	i
*>i100.1.1.1/32	11.0.0.1		100	0	i
*>i100.1.1.2/32	11.0.0.1		100	0	i
*>i200.1.1.1/32	11.0.0.1		100	0	i
*>i200.1.1.2/32	11.0.0.1		100	0	i

Processed 6 prefixes, 6 paths

Router# **show bgp vrf foo ipv4 unicast 100.1.1.1/32 detail**

```

Mon Dec 8 23:24:50.243 PST
BGP routing table entry for 100.1.1.1/32, Route Distinguisher:
30.30.30.30:0

```

Versions:

Process	bRIB/RIB	SendTblVer
Speaker	43	43

Local Label: 24001 (with rewrite);

Flags: 0x05081001+0x00000200;

Last Modified: Dec 8 18:04:21.000 for 05:20:30

Paths: (1 available, best #1)

```

  Advertised to PE peers (in unique update groups):
    32.0.0.2

```

Path #1: Received by speaker 0

Flags: 0x400061000d060005, import: 0x80

```

  Advertised to PE peers (in unique update groups):
    32.0.0.2

```

Local

11.0.0.1 (metric 2) from 20.0.0.1 (11.0.0.1)

Received Label 1234

```

Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
imported, reoriginated with stitching-rt
Received Path ID 0, Local Path ID 1, version 43
Extended community: RT:1:2
Originator: 11.0.0.1, Cluster list: 20.20.20.20
Source AFI: L2VPN EVPN, Source VRF: default, Source Route Distinguisher: 100:1v

```

```
Router# show bgp vpnv4 unicast update-group
```

```
Fri Aug 21 00:27:57.910 PDT
```

```
Update group for VPNv4 Unicast, index 0.1:
```

```

Attributes:
  Outbound policy: pass
  First neighbor AS: 200
  Send communities
  Send GSHUT community if originated
  Send extended communities
  4-byte AS capable
  Send Re-originated VPN routes
  Send multicast attributes
  Minimum advertisement interval: 30 secs
Update group desynchronized: 0
Sub-groups merged: 0
Number of refresh subgroups: 0
Messages formatted: 8, replicated: 8
All neighbors are assigned to sub-group(s)
  Neighbors in sub-group: 0.2, Filter-Groups num:1
  Neighbors in filter-group: 0.2(RT num: 0)
  32.0.0.2

```

```
Update group for VPNv4 Unicast, index 0.3:
```

```

Attributes:
  Neighbor sessions are IPv4
  Internal
  Common admin
  First neighbor AS: 100
  Send communities
  Send GSHUT community if originated
  Send extended communities
  4-byte AS capable
  Send AIGP
  Send Re-originated VPN routes
  Send multicast attributes
  Minimum advertisement interval: 0 secs
Update group desynchronized: 0
Sub-groups merged: 0
Number of refresh subgroups: 0
Messages formatted: 2, replicated: 2
All neighbors are assigned to sub-group(s)
  Neighbors in sub-group: 0.1, Filter-Groups num:1
  Neighbors in filter-group: 0.1(RT num: 0)
  20.0.0.1

```

```
Router# show bgp l2vpn evpn update-group
```

```
Fri Aug 21 00:27:42.786 PDT
```

```
Update group for L2VPN EVPN, index 0.2:
```

```

Attributes:
  Neighbor sessions are IPv4
  Internal
  Common admin
  First neighbor AS: 100
  Send communities

```

```

Send GSHUT community if originated
Send extended communities
4-byte AS capable
Send AIGP
Send multicast attributes
Minimum advertisement interval: 0 secs
Update group desynchronized: 0
Sub-groups merged: 0
Number of refresh subgroups: 0
Messages formatted: 4, replicated: 4
All neighbors are assigned to sub-group(s)
Neighbors in sub-group: 0.1, Filter-Groups num:1
Neighbors in filter-group: 0.1(RT num: 0)
20.0.0.1
    
```

Data Center Interconnect between MPLS-VPN and EVPN-VxLAN

This part provides conceptual and configuration information for Data Center Interconnect (DCI) VXLAN Layer 3 Gateway on the router.


Release	Modification
Release 5.3.2	This feature was introduced.
Release 6.1.x	OpFlex
Release 6.6.x	EVPN VxLAN VRF leaking

Data Center Interconnect VXLAN Layer 3 Gateway

Router can serve as a Data Center Interconnect (DCI) L3 Gateway using stitching technology between VPNv4/v6 and EVPN-VXLAN. The DCI provides a solution for a new EVPN-VXLAN Data Center that needs to communicate with legacy and existing traditional MPLS VPN networks (VPNv4) having PE-CE architecture.


The DCI L3 gateway provides the following functions:

- IP connectivity between multi-tenant remote Data Center sites:** Consider the following network topology that has two Data Center sites connected through the intermediate service provider network. The multi-tenant Data Centers use VXLAN encapsulation to carry separate tenant IP traffic. The VXLAN-enabled Data Center sites use the MP-BGP EVPN control plane for distributing both Layer-2 and Layer-3 forwarding information within the site. The router uses MPLS L3VPN application service over the service provider network to provide L3 connectivity between the two Data Center sites. Making this translation between EVPN-VXLAN to VPNv4 overlay.



521487

- **IP Connectivity between Data Center and remote PEs in a legacy network:** Consider the following network topology that has one new Data Center site connected through the intermediate service provider network. The multi-tenant Data Center uses VXLAN encapsulation to carry separate tenant IP traffic. The VXLAN-enabled Data Center site uses the MP-BGP EVPN control plane for distributing both Layer-2 and Layer-3 forwarding information within the site. The router uses MPLS L3VPN application service over the service provider network to provide L3 connectivity between the Data Center services and the legacy CEs using VPNv4 to communicate with services placed inside the Data Center. Making this translation between EVPN-VXLAN to VPNv4 overlay.



521488

**Note**

- DCI gateway does not provide layer 2 inter-connectivity across Data Centers.
- In a DCI deployment, for route reoriginate with stitching-rt for a particular VRF, using the same Route Distinguisher (RD) between DCI and MPLS-VPN PE or same RD between DCI and VxLAN Top of Rack (ToR) is not supported.

Route Targets

For each VRF on the DCI router, there are two sets of manually configured import and export route-targets. One set of import and export route-targets is associated with the Data Center BGP neighbor that uses EVPN address-family to exchange L3 information; the other set of import and export route-targets is associated with the L3VPN BGP neighbor that use VPNv4 or VPNv6 unicast address-family to exchange L3 information. This separation of route targets (RTs) enables the two sets of RTs to be independently configured. The DCI router effectively stitches the two set of RTs. The RTs associated with the EVPN BGP neighbor are labelled as stitching RTs. The RTs associated with the L3VPN BGP neighbor are normal RTs.

Route Re-origination

Consider the case of control plane information propagation by the DCI from the L3VPN side to the Data Center side. Here, instead of advertising the remote Data Center's original BGP EVPN routes, you can configure the DCI router to advertise to its BGP EVPN neighbor the routes that are re-originated after importing them from the L3VPN BGP neighbor. For this case of VPNv4 or VPNv6 routes being propagated to the BGP EVPN neighbors (Data Center neighbors), re-originating the routes refers to replacing the normal route-target values associated with the BGP EVPN neighbors. The converse holds true for the routing information traffic propagation from the BGP EVPN control plane to BGP L3VPN control plane. You can configure this re-origination by using the re-originate keyword in the import re-originate command. Configuring this command, by default, also enables advertisement of L2VPN EVPN prefixes to the EVPN BGP neighbors. You can suppress native L2VPN EVPN address-family NLRI advertisements towards the EVPN Neighbor using the advertise l2vpn evpn disable command under the EVPN BGP address-family configuration mode.

Route Address-Family and Encoded Address-Family

When an address-family is configured for a BGP neighbor, it means that the specified address-family routes encoded with the NLRI for that address-family is advertised to the neighbor. This does not hold for data center BGP neighbors because they use only EVPN address-family. Here, BGP neighbors advertise VPNv4 or VPNv6 unicast routes using the EVPN NLRI encoding. Thus, here the encoded address-family and route address family can be possibly different. You can advertise the VPNv4 or VPNv6 address-family using the advertise vpnv4 unicast or advertise vpnv6 unicast command. For example, a EVPN address-family BGP neighbor configured with the advertise vpnv4 unicast command sends VPNv4 unicast routes in an EVPN encoded NLRI.

Local VPNv4 or VPNv6 Routes Advertisement

On the DCI router, the locally sourced VPNv4 or VPNv6 routes can be advertised to the BGP EVPN neighbors with the normal route targets (RTs) configured for the VRF or the stitching RTs associated with the BGP EVPN neighbors. By default, these routes are advertised with the normal route targets. You can configure these local VPNv4 or VPNv6 route advertisements to be advertised with stitching RTs to the BGP EVPN neighbors by using the advertise vpnv4 unicast local stitching-rt or advertise vpnv6 unicast local stitching-rt command as required.

Data Center VXLAN with Support for MP-BGP

The Data Center VXLAN uses MP-BGP for control-plane learning of end-host Layer 2 and Layer 3 reachability information. The DCI router is configured with a VXLAN Tunnel EndPoint (VTEP). For VTEP configuration details, see the chapter Implementing Layer 3 VXLAN Gateway. You also need to run the host-reachability protocol bgp command to specify that control-plane learning within Data center site is through BGP routing protocol.

The DCI Gateway router and the EVPN BGP neighbor (Data Center BGP neighbor) exchange BGP EVPN NLRIs of route type 5 that carry L3 routing information and associated VXLAN encapsulation information. Some of the VXLAN information is carried in the EVPN NLRI and the rest is carried in RFC 5512 Tunnel Type Encapsulation EXTCOMM and Router MAC EXTCOMM defined in draft-ietf-bess-evpn-inter-subnet-forwarding-00. BGP downloads VXLAN encapsulation as RIB remote next hop opaque attribute to L3RIB.

Default-Originate Forwarding to BGP EVPN Neighbor

Instead of advertising the specific networks available in the remote Data Center, you can configure the DCI gateway to advertise a default route to the directly connected Data Center neighbor. To send the default route for a VRF instance to the Data Center BGP EVPN neighbor, the VPN default-originate information that is typically forwarded to the L3VPN BGP neighbor, is also configured to be forwarded to the BGP EVPN neighbor in the Data Center. To do so, you need to configure allow vpn default-originate command in the BGP VRF configuration mode and also configure default-originate command under EVPN BGP neighbor in L2VPN EVPN address-family configuration mode. This configures BGP to forward only one default route information for a VRF instance from the DCI Gateway to the BGP neighbor that has L2VPN EVPN address-family. This default route information is encoded in the EVPN "IP Prefix Route" NLRI.

With the advertisement of a default route to the connected Data Center, the DCI Gateway should not advertise specific prefixes of the remote Data Center to the BGP EVPN neighbor. To prevent forwarding of VRF prefixes, you need to configure the DCI gateway with a EVPN BGP neighbor policy that drops forwarding of all prefixes.

Configure Data Center Interconnect Router

Perform the following steps to configure the Data Center Interconnect (DCI) router:

- Configure VRF and route targets import/export rules
- Configure Bridge Domain for DCI Gateway
- Configure VTEP.
- Configure EVPN BGP neighbor and route advertisement
- Configure L3VPN BGP neighbor relationship and route advertisements

Configure VRF and route targets import/export rules

Perform the following to configure VRF and define route targets to be used for import and export of forwarding information.

```
Router# configure
Router(config)# vrf data-center-10
Router(config-vrf-af)# address-family ipv4 unicast
Router(config-vrf-af)# import route-target 1:1
Router(config-vrf-af)# export route-target 1:2
Router(config-vrf-af)# import route-target 10:1 stitching
Router(config-vrf-af)# export route-target 10:2 stitching
Router(config-vrf-af)# commit
```

Configure Bridge Domain for DCI Gateway

Perform the following to configure the bridge domain on the DCI Gateway.



Note For DCI VxLAN L3 Gateway, only routed interface BVI and member vni can be configured in the bridge-domain. All other L2 services such as EVI, PW, or AC are not supported in the bridge-domain.

```
Router# configure
Router(config)# interface bvi 1
Router(config-if)# vrf cust1
Router(config-if)# ipv4 address 40.1.1.1 255.255.255.255
Router(config)# exit
Router(config)# l2vpn
Router(config-l2vpn)# bridge group bg1
Router(config-l2vpn-bg)# bridge-domain bd1
Router(config-l2vpn-bg-bd)# routed interface BVI1
Router(config-l2vpn-bg-bd)# member vni 5001
Router(config-l2vpn-bg-bd)# commit
```

Configure VTEP (VxLAN Terminal EndPoint) on the DCI Gateway.

Perform the following to configure VTEP (VxLAN Terminal EndPoint) on the DCI Gateway.

```
Router# configure
Router(config)# interface loopback 0
Router(config-if)# ipv4 address 40.1.1.1 255.255.255.255
Router(config)# exit
Router(config)# interface nve 1
```

```

Router(config-if)# source interface loopback 0
Router(config-if)# member vni 5001
Router(config-nve-vni)# vrf cust1
Router(config-nve-vni)# host reachability protocol bgp
Router(config-nve-vni)# commit

```

Configure EVPN BGP neighbor and route advertisements

Perform the following on the DCI router to configure BGP neighbor relationship and route advertisements with the EVPN BGP neighbor.

```

Router# configure
Router(config)# router bgp 100
Router(config-bgp)# address-family l2vpn evpn
Router(config-bgp-af)# exit
Router(config-bgp)# neighbor 1.1.1.1
Router(config-bgp-nbr)# remote-as 100
Router(config-bgp-nbr)# address-family l2vpn evpn
Router(config-bgp-nbr)# default-originate /*optional configuration*/
Router(config-bgp-nbr-af)# import stitching-rt reoriginate
Router(config-bgp-nbr-af)# advertise vpnv4 unicast re-originated
Router(config-bgp-nbr-af)# advertise vpnv6 unicast re-originated
Router(config-bgp-nbr-af)# advertise l2vpn evpn disable/*optional configuration*/
Router(config-bgp-nbr-af)# commit

```

Configure L3VPN BGP neighbor relationship and route advertisements

Perform the following to configure BGP neighbor relationship and route advertisements with the L3VPN BGP neighbor.

```

Router# configure
Router(config)# router bgp 100
Router(config-bgp-nbr)# address-family vpnv4
Router(config-bgp-af)# exit
Router(config-bgp)# neighbor 1.1.1.1
Router(config-bgp-nbr)# remote-as 100
Router(config-bgp-nbr)# address-family vpnv4
Router(config-bgp-nbr-af)# import reoriginate stitching-rt
Router(config-bgp-nbr-af)# advertise vpnv4 unicast re-originated

```

```

Router# configure
Router(config)# router bgp 100
Router(config-bgp-nbr)# address-family vpnv6
Router(config-bgp-af)# exit
Router(config-bgp)# neighbor 1.1.1.1
Router(config-bgp-nbr)# remote-as 100
Router(config-bgp-nbr)# address-family vpnv6
Router(config-bgp-nbr-af)# import reoriginate stitching-rt
Router(config-bgp-nbr-af)# advertise vpnv6 unicast re-originated

```

Verification

You can use the following show commands to verify the DCI Gateway configurations:

```

Router# show bgp l2vpn evpn

BGP router identifier 30.30.30.30, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active

```

```

Table ID: 0x0   RD version: 0
BGP main routing table version 16
BGP NSR Initial initsync version 1 (Reached)
BGP NSR/ISSU Sync-Group versions 16/0
BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best
               i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 100:1
*>i[2][10000][48][0226.51bd.c81c][32][200::1001]/232
   11.0.0.1          100          0 i
*>i[2][10000][48][0226.51bd.c81c][32][200:1::1001]/232
   11.0.0.1          100          0 i
*>i[2][10000][48][0226.51bd.c81c][32][200.1.1.1]/136
   11.0.0.1          100          0 i
*>i[2][10000][48][0226.51bd.c81c][32][200.1.1.2]/136
   11.0.0.1          100          0 i
*>i[5][4231][32][100.1.1.1]/80
   11.0.0.1          100          0 i
*>i[5][4231][32][100.1.1.2]/80
   11.0.0.1          100          0 i
*>i[5][4231][112][fec0::1001]/176
   11.0.0.1          100          0 i
*>i[5][4232][112][fec0::1:1001]/176
   11.0.0.1          100          0 i

Processed 8 prefixes, 8 paths

```

```
Router# show bgp l2vpn evpn rd 100:1 [5][4231][112][fec0::1001]/176 detail
```

```

BGP routing table entry for [5][4231][112][fec0::1001]/176, Route Distinguisher: 100:1
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          5         5
  Flags: 0x04040001+0x00000000;
Last Modified: Aug 21 00:16:58.000 for 00:17:46
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  Flags: 0x4000600025060005, import: 0x3f
  Not advertised to any peer
  Local
    11.0.0.1 (metric 2) from 20.0.0.1 (11.0.0.1)
    Received Label 16001
    Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
reoriginate, not-in-vrf
    Received Path ID 0, Local Path ID 1, version 5
    Extended community: Flags 0x2: Encapsulation Type:8 Router MAC:aabb.cccd.eeff RT:65540:1
RT:40.40.40.40:1 RT:100:1
  Originator: 11.0.0.1, Cluster list: 20.20.20.20
  EVPN ESI: ffff.ffff.ffff.ffff.fff01, Gateway Address : fec0::254

```

```
Router# show bgp l2vpn evpn neighbors 20.0.0.1 detail
```

```

BGP neighbor is 20.0.0.1
Remote AS 100, local AS 100, internal link
Remote router ID 20.20.20.20
BGP state = Established, up for 00:08:58
NSR State: NSR Ready
Last read 00:00:34, Last read before reset 00:00:00

```



```

Hold time is 180, keepalive interval is 60 seconds
Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
Last write 00:00:36, attempted 19, written 19
Second last write 00:01:36, attempted 143, written 143
Last write before reset 00:00:00, attempted 0, written 0
Second last write before reset 00:00:00, attempted 0, written 0
Last write pulse rcvd Aug 21 00:25:03.667 last full not set pulse count 33
Last write pulse rcvd before reset 00:00:00
Socket not armed for io, armed for read, armed for write
Last write thread event before reset 00:00:00, second last 00:00:00
Last KA expiry before reset 00:00:00, second last 00:00:00
Last KA error before reset 00:00:00, KA not sent 00:00:00
Last KA start before reset 00:00:00, second last 00:00:00
Precedence: internet
Non-stop routing is enabled
Entered Neighbor NSR TCP mode:
  TCP Initial Sync :           Aug 21 00:18:07.291
  TCP Initial Sync Phase Two : Aug 21 00:18:07.319
  TCP Initial Sync Done :      Aug 21 00:18:08.334
Multi-protocol capability received
Neighbor capabilities:          Adv          Rcvd
Route refresh:                  Yes         Yes
4-byte AS:                      Yes         Yes
Address family VPNv4 Unicast:   Yes         No
Address family VPNv6 Unicast:   Yes         No
Address family L2VPN EVPN:      Yes         Yes
Message stats:
  InQ depth: 0, OutQ depth: 0
  InQ depth: 0, OutQ depth: 0
  Last_Sent          Sent  Last_Rcvd          Rcvd
Open:                Aug 21 00:16:38.087    1  Aug 21 00:16:40.123    1
Notification:        ---          0  ---                  0
Update:              Aug 21 00:24:01.421    9  Aug 21 00:24:03.652   13
Keepalive:           Aug 21 00:25:01.434    8  Aug 21 00:25:03.667    9
Route_Refresh:       Aug 21 00:24:01.377    3  ---                  0
Total:                21
Minimum time between advertisement runs is 0 secs
Inbound message logging enabled, 3 messages buffered
Outbound message logging enabled, 3 messages buffered

```

```

For Address Family: VPNv4 Unicast
BGP neighbor version 35
Update group: 0.3 Filter-group: 0.1 No Refresh request being processed
Advertise Reorigination Enabled
Advertise AFI EoR can be sent
Route refresh request: received 0, sent 0
0 accepted prefixes, 0 are bestpaths
Cumulative no. of prefixes denied: 0.
Prefix advertised 4, suppressed 0, withdrawn 0
Maximum prefixes allowed 2097152
Threshold for warning message 75%, restart interval 0 min
AIGP is enabled
An EoR was not received during read-only mode
Last ack version 35, Last synced ack version 35
Outstanding version objects: current 0, max 1
Additional-paths operation: None
Send Multicast Attributes

```

```

For Address Family: VPNv6 Unicast
BGP neighbor version 29
Update group: 0.3 Filter-group: 0.1 No Refresh request being processed
Advertise Reorigination Enabled
Advertise AFI EoR can be sent
Route refresh request: received 0, sent 0
0 accepted prefixes, 0 are bestpaths

```

```

Cumulative no. of prefixes denied: 0.
Prefix advertised 0, suppressed 0, withdrawn 0
Maximum prefixes allowed 1048576
Threshold for warning message 75%, restart interval 0 min
AIGP is enabled
An EoR was not received during read-only mode
Last ack version 29, Last synced ack version 29
Outstanding version objects: current 0, max 0
Additional-paths operation: None
Send Multicast Attributes
Advertise VPNv4 routes enabled with Reoriginate,Local with stitching-RT option

```

```

For Address Family: L2VPN EVPN
BGP neighbor version 18
Update group: 0.2 Filter-group: 0.1 No Refresh request being processed
Route refresh request: received 0, sent 3
8 accepted prefixes, 8 are bestpaths
Cumulative no. of prefixes denied: 0.
Prefix advertised 4, suppressed 0, withdrawn 6
Maximum prefixes allowed 2097152
Threshold for warning message 75%, restart interval 0 min
AIGP is enabled
An EoR was received during read-only mode
Last ack version 18, Last synced ack version 18
Outstanding version objects: current 0, max 2
Additional-paths operation: None
Send Multicast Attributes
Advertise VPNv4 routes enabled with Reoriginate, option
Advertise VPNv6 routes is enabled with Reoriginate, option
Import Stitching is enabled for this neighbor address-family
Import Reoriginate is enabled for this neighbor address-family

Connections established 1; dropped 0
Local host: 30.0.0.1, Local port: 59405, IF Handle: 0x00000000
Foreign host: 20.0.0.1, Foreign port: 179
Last reset 00:00:00

```

Router# **show bgp sessions**

Neighbor	VRF	Spk	AS	InQ	OutQ	NBRState	NSRState
20.0.0.1	default	0	100	0	0	Established	NSR Ready[PP]
32.0.0.2	default	0	200	0	0	Established	NSR Ready

Router# **show bgp vpnv4 unicast**

```

BGP router identifier 30.30.30.30, local AS number 100
BGP generic scan interval 60 secs
Non-stop routing is enabled
BGP table state: Active
Table ID: 0x0 RD version: 0
BGP main routing table version 39
BGP NSR Initial initsync version 4 (Reached)
BGP NSR/ISSU Sync-Group versions 39/0
BGP scan interval 60 secs

```

```

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete

```

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher: 1:1					
*> 1.1.1.0/24	32.0.0.2			0 200 300	i
*> 1.1.2.0/24	32.0.0.2			0 200 300	i
Route Distinguisher: 30.30.30.30:0 (default for vrf foo)					
*> 1.1.1.0/24	32.0.0.2			0 200 300	i

```
*> 1.1.2.0/24          32.0.0.2          0 200 300 i
*>i100.1.1.1/32       11.0.0.1          100   0 i
*>i100.1.1.2/32       11.0.0.1          100   0 i
*>i200.1.1.1/32       11.0.0.1          100   0 i
*>i200.1.1.2/32       11.0.0.1          100   0 i
```

Router# **show bgp vpnv4 unicast rd 30.30.30.30:0 1.1.1.0/24 detail**

```
BGP routing table entry for 1.1.1.0/24, Route Distinguisher: 30.30.30.30:0
Versions:
  Process          bRIB/RIB  SendTblVer
  Speaker          26        26
  Flags: 0x04103001+0x00000000;
Last Modified: Aug 21 00:24:01.000 for 00:04:58
Paths: (1 available, best #1)
  Advertised to peers (in unique update groups):
    20.0.0.1
  Path #1: Received by speaker 0
  Flags: 0x4000c00005060001, import: 0x80
  Advertised to peers (in unique update groups):
    20.0.0.1
  200 300
  32.0.0.2 from 32.0.0.2 (40.40.40.40)
  Received Label 24001
  Origin IGP, localpref 100, valid, external, best, group-best, import-candidate,
imported, reoriginated with stitching-rt
  Received Path ID 0, Local Path ID 1, version 26
  Extended community: RT:100:2
  Source AFI: VPNv4 Unicast, Source VRF: default, Source Route Distinguisher: 1:1
```

Router# **show bgp vrf foo**

```
BGP VRF foo, state: Active
BGP Route Distinguisher: 30.30.30.30:0
VRF ID: 0x60000002
BGP router identifier 30.30.30.30, local AS number 100
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000011 RD version: 35
BGP main routing table version 35
BGP NSR Initial initsync version 4 (Reached)
BGP NSR/ISSU Sync-Group versions 31/0

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 30.30.30.30:0 (default for vrf foo)
*> 1.1.1.0/24       32.0.0.2          0 200 300 i
*> 1.1.2.0/24       32.0.0.2          0 200 300 i
*>i100.1.1.1/32     11.0.0.1          100   0 i
*>i100.1.1.2/32     11.0.0.1          100   0 i
*>i200.1.1.1/32     11.0.0.1          100   0 i
*>i200.1.1.2/32     11.0.0.1          100   0 i
```

Processed 6 prefixes, 6 paths

Router# **show bgp vrf foo ipv4 unicast 100.1.1.1/32 detail**

```
BGP routing table entry for 100.1.1.1/32, Route Distinguisher:
30.30.30.30:0
```

```

Versions:
Process          bRIB/RIB  SendTblVer
Speaker         43       43
  Local Label: 24001 (with rewrite);
  Flags: 0x05081001+0x00000200;
Last Modified: Dec  8 18:04:21.000 for 05:20:30
Paths: (1 available, best #1)
  Advertised to PE peers (in unique update groups):
    32.0.0.2
  Path #1: Received by speaker 0
  Flags: 0x400061000d060005, import: 0x80
  Advertised to PE peers (in unique update groups):
    32.0.0.2
Local
  11.0.0.1 (metric 2) from 20.0.0.1 (11.0.0.1)
    Received Label 1234
    Origin IGP, localpref 100, valid, internal, best, group-best, import-candidate,
imported, reoriginated
    Received Path ID 0, Local Path ID 1, version 43
    Extended community: Encapsulation Type:8 Router MAC:aabb.ccdd.eeff RT:1:2
    Originator: 11.0.0.1, Cluster list: 20.20.20.20
    RIB RNH: table_id 0xe0000011, Encap 8, VNI 1234, MAC Address: aabb.ccdd.eeff, IP
Address: 11.0.0.1, IP table_id 0xe0000000
  Source AFI: L2VPN EVPN, Source VRF: default, Source Route
Distinguisher: 100:1

Router# show bgp vpnv4 unicast update-group

```

```

Update group for VPNv4 Unicast, index 0.1:
Attributes:
  Outbound policy: pass
  First neighbor AS: 200
  Send communities
  Send GSHUT community if originated
  Send extended communities
  4-byte AS capable
  Send Re-originated VPN routes
  Send multicast attributes
  Minimum advertisement interval: 30 secs
Update group desynchronized: 0
Sub-groups merged: 0
Number of refresh subgroups: 0
Messages formatted: 8, replicated: 8
All neighbors are assigned to sub-group(s)
  Neighbors in sub-group: 0.2, Filter-Groups num:1
  Neighbors in filter-group: 0.2(RT num: 0)
    32.0.0.2

```

```

Update group for VPNv4 Unicast, index 0.3:
Attributes:
  Neighbor sessions are IPv4
  Internal
  Common admin
  First neighbor AS: 100
  Send communities
  Send GSHUT community if originated
  Send extended communities
  4-byte AS capable
  Send AIGP
  Send Re-originated VPN routes
  Send multicast attributes
  Minimum advertisement interval: 0 secs
Update group desynchronized: 0

```

```

Sub-groups merged: 0
Number of refresh subgroups: 0
Messages formatted: 2, replicated: 2
All neighbors are assigned to sub-group(s)
  Neighbors in sub-group: 0.1, Filter-Groups num:1
    Neighbors in filter-group: 0.1(RT num: 0)
      20.0.0.1

```

```
Router# show bgp l2vpn evpn update-group
```

```

Update group for L2VPN EVPN, index 0.2:
Attributes:
  Neighbor sessions are IPv4
  Internal
  Common admin
  First neighbor AS: 100
  Send communities
  Send GSHUT community if originated
  Send extended communities
  4-byte AS capable
  Send AIGP
  Send multicast attributes
  Minimum advertisement interval: 0 secs
Update group desynchronized: 0
Sub-groups merged: 0
Number of refresh subgroups: 0
Messages formatted: 4, replicated: 4
All neighbors are assigned to sub-group(s)
  Neighbors in sub-group: 0.1, Filter-Groups num:1
    Neighbors in filter-group: 0.1(RT num: 0)
      20.0.0.1

```

Example for configuring Data Center Interconnection Layer 3 Gateway

The following configurations provide an example Data Center Interconnection (DCI) Layer 3 Gateway configuration.

VTEP-related configuration:

```

interface Loopback1
  ipv4 address 40.1.1.1 255.255.255.255
!

interface nve1
  source-interface Loopback1
  member vni 1
  vrf cust1
  host-reachabilty protocol bgp
!
interface BVI1
  vrf cust1
  ipv4 address 10.99.1.30 255.255.255.0
  ipv6 address 10:99:1::30/64
!

l2vpn
  bridge group bg1
  bridge-domain bd1
  routed interface BVI1
  member vni 1

```

```
!
!
```

VRF-related configuration

```
vrf data-center-10
  import route-target 1:1
  export route-target 1:2
  import route-target 10:10 stitching
  export route-target 10:20 stitching
```

Data Center EVPN BGP neighbor-related configuration

```
router bgp 1
  neighbor 1.1.1.1
    address-family l2vpn evpn
      import stitching-rt reoriginate
      advertise vpv4 unicast reoriginated
      advertise vpv6 unicast reoriginated
      advertise vpv4 unicast local stitching-rt
      advertise vpv6 unicast local stitching-rt
      advertise l2vpn evpn disable
```

L3VPN BGP neighbor-related configuration

```
router bgp 2
  neighbor 10.10.10.10
    address-family vpv4
      import reoriginate stitching-rt
      advertise vpv4 unicast reoriginated
```

The following example configuration shows how to configure the DCI router to forward default route to its Data Center neighbor.

```
router bgp 1
  address-family vpv4 unicast
  address-family vpv6 unicast
  address-family l2vpn evpn
  exit
  neighbor 1.1.1.1
    address-family l2vpn evpn
    default-originate
  exit
  vrf foo
    rd 2:1
    address-family ipv4 unicast
    allow vpn default-originate
  exit
  address-family ipv6 unicast
  allow vpn default-originate
  exit
exit
!
```

OpFlex


OpFlex is an open and extensible policy protocol used for transferring the policy information between a network policy controller such as the Cisco Application Policy Infrastructure Controller (APIC) and network elements such as routers that are configured as Data Center Interconnect (DCI) gateway. The policies are distributed using the Cisco® Application Centric Infrastructure (ACI) infrastructure within the fabric to the spine nodes. The spine nodes send policies to the DCI gateway through the OpFlex framework. An OpFlex framework resides between the spines and the DCIs. It enables the distribution of the DCI policy model from the fabric to the DCI gateways. DCI gateway acts as an OpFlex agent and the spine acts a policy repository.

Fabric tenant interconnect (FTI) is the OpFlex agent application that runs on the DCI to generate and apply the tenant device configuration on the DCI. Policies configure the DCI service for a given tenant on the DCI gateway.

OpFlex Topology

Consider the topology where OpFlex framework is used between the DCI gateway and the Cisco ACI spine switches to automate fabric-facing tenant provisioning on the DCI gateway. When you configure a new external Layer 3 outside (L3Out) policy for a tenant on the Cisco Application Policy Infrastructure Controller (APIC), the controller programs all related information associated with that tenant, such as VRF instance name and BGP extended community route-target attributes for the Cisco ACI spine switches. The OpFlex framework running on the spine switches reads the L3Out managed object and converts it to the OpFlex model. This information is then pushed to the DCI gateway, which acts as a policy element for the OpFlex framework. On the DCI, the fabric facing configuration for the tenant VFR is auto-generated.

Figure 1: OpFlex Topology



Restrictions

The OpFlex feature is supported with the following restrictions:

- OpFlex feature is not supported on ASR9K with power PC based route-processor.
- FTI cannot generate configuration for multiple RTs of one address family in a tenant VRF provisioned in one fabric.
- On exhaustion of FTI configuration pools, the OpFlex notifications to add tenants are ignored. If existing tenants are deleted, the new tenants must be added again to enable OpFlex notifications to be re-sent to the DCI.
- FTI supports only Type 0 RT format: 2 byte ASN + 4 byte value. Type 1 and Type 2 RT formats are not supported.
- XML configuration and oper schema are not supported for FTI configuration and show commands.

Configure OpFlex

Perform the following tasks to configure the OpFlex session to automate fabric-facing tenant provisioning on the DCI gateway. This includes the one-time configuration that must be done on the DCI to enable DCI hand-off from an ACI fabric.

Configure BGP

Perform this task to enable address-family under BGP routing process for fabric and WAN peering.

```
Router# configure
Router(config)# router bgp 1234
Router(config-bgp)# bgp router-id 198.51.100.1
Router(config-bgp)# address-family vpnv4 unicast
Router(config-bgp-af)# commit
```

Configure BGP Session on the Fabric Side

Perform this task to configure BGP session on the fabric side.

```
Router# configure
Router(config)# router bgp 200
Router(config-bgp)# neighbor 209.165.201.1
Router(config-bgp-nbr)# remote-as 100
Router(config-bgp-nbr)# update-source loopback2
Router(config-bgp-nbr)# address-family l2vpn evpn
Router(config-bgp-nbr-af)# import stitching-rt reoriginate
Router(config-bgp-nbr-af)# advertise vpnv4 unicast re-originated
Router(config-bgp-nbr-af)# commit
```

Configure BGP Session on the WAN Side

Perform this task to configure BGP session on the WAN side.

```
Router# configure
Router(config)# router bgp 200
Router(config-bgp)# neighbor 209.165.200.226
Router(config-bgp-nbr)# remote-as 100
Router(config-bgp-nbr)# update-source loopback2
Router(config-bgp-nbr)# address-family vpnv4 unicast
Router(config-bgp-nbr-af)# import re-originate stitching-rt
Router(config-bgp-nbr-af)# advertise vpnv4 unicast re-originated
Router(config-bgp-nbr-af)# commit
```

Configure DCI Underlay for Fabric and WAN Interfaces

Perform this task to configure DCI underlay for fabric facing interface and WAN facing interface. Perform this task on both the interfaces.

```
Router# configure
Router(config)# interface GigabitEthernet 0/0/0/0
Router(config-if)# ipv4 address 209.165.200.226 255.255.255.224
Router(config-if)# commit
```

Configure IGP for ACI and WAN Reachability

Perform this task to configure IGP for ACI and WAN reachability.

```
Router# configure
Router(config)# router ospf 100
Router(config-ospf)# area 0
Router(config-ospf-ar)# interface GigabitEthernet 0/0/0/1
Router(config-ospf-ar-if)# exit
Router(config-ospf-ar)# exit
Router(config-ospf)# area 100
Router(config-ospf-ar)# nssa
Router(config-ospf-ar)# interface loopback0
Router(config-ospf-ar-if)# exit
```



```
Router(config-ospf-ar) # interface GigabitEthernet 0/0/0/0
Router(config-ospf-ar) # commit
```

Configure MPLS towards WAN

Perform this task to configure MPLS on the DCI.

```
Router# configure
Router# mpls ldp
Router(config-ldp) # interface GigabitEthernet 0/0/0/1
Router(config-ldp-if) # exit
Router(config-ldp) # exit
Router(config) # interface Loopback0
Router(config-if) # ipv4 address 209.165.200.227 255.255.255.224
Router(config-if) # exit
Router(config) # interface nve 1
Router(config-if) # source-interface loopback 0
Router(config-if) # commit
```

Configure FTI Auto-Configuration Parameters

Perform this task to configure FTI auto-configuration parameters.

```
Router# configure
Router(config) # dci-fabric-interconnect
Router(config-fti) # auto-configuration-pool
Router(config-fti-acp) # bgp-as 1234
Router(config-fti-acp) # bridge group bg1
Router(config-fti-acp) # vrf vrfl ipv4-address 198.51.100.1
Router(config-fti-acp) # bd-pool 1 1000
Router(config-fti-acp) # vni-pool 1 1000
Router(config-fti-acp) # local-vtep nve 1
Router(config-fti-acp) # commit
```

Configure OpFlex Session

This task enables the fabric tenant interconnect to setup an OpFlex session with the spine.

```
Router# configure
Router(config) # dci-fabric-interconnect
Router(config-fti) # fabric 1001
Router(config-fti-fabric) # opflex-peer 192.0.2.1
Router(config-fti-fabric) # exit
Router(config-fti) # identity 203.0.113.1
Router(config-fti) # commit
```

