

Configure BGP DN for Multiple AF on Catalyst 9000 Series Switches

Contents

[Introduction](#)
[Prerequisites](#)
[Requirements](#)
[Components Used](#)
[Background Information](#)
[Terminology](#)
[Configure](#)
[Network Diagram](#)
[Underlay \(OSPF\)](#)
[Overlay \(BGP\)](#)
[Spine-01](#)
[Leaf-01](#)
[Optional Configurations](#)
[Alternate-AS](#)
[Verify](#)
[Spine \(RR\)](#)
[Troubleshoot](#)
[DN Peering](#)
[DN Debugs](#)
[Static Neighbor](#)
[Related Information](#)

Introduction

This document describes EVPN simplification CLI for BGP Dynamic Neighbor in EVPN and MVPN address families on Catalyst 9000 series switches.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Basic BGP configuration
- BGP address families

Components Used

The information in this document is based on these software and hardware versions:

- Catalyst 9300
- Catalyst 9400
- Catalyst 9500
- Catalyst 9600

- Cisco IOS® XE 17.11.1 and later

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

EVPN deployments involve BGP neighborships with many peers, often within the same subnet. Before the introduction of the BGP Dynamic Neighbor (DN) feature at least 255 configuration lines would be required and maintained to establish BGP sessions with all peers in a /24 subnet. **With BGP DN this can be achieved with only 3 lines.** The only functional difference between the static neighbor & dynamic neighbor configurations is that the latter only creates connections with other BGP peers passively. Therefore, this feature can only be configured by one of the peers.

The ability to group together many peers within a subnet range is highly desirable, if not necessary for the deployment to be feasible, and has already been widely adopted for the IPv4, IPv6 and VRF address-families. This feature is desirable for other address-families, especially EVPN, as it helps avoid the writing and maintenance of extensive and complex configurations in Spine-Leaf topologies where many leaves are in the same subnet.

Note: This feature does not introduce any new CLIs. It allows activation of peer-groups configured with 'bgp listen range' under previously unsupported address-families.

Terminology

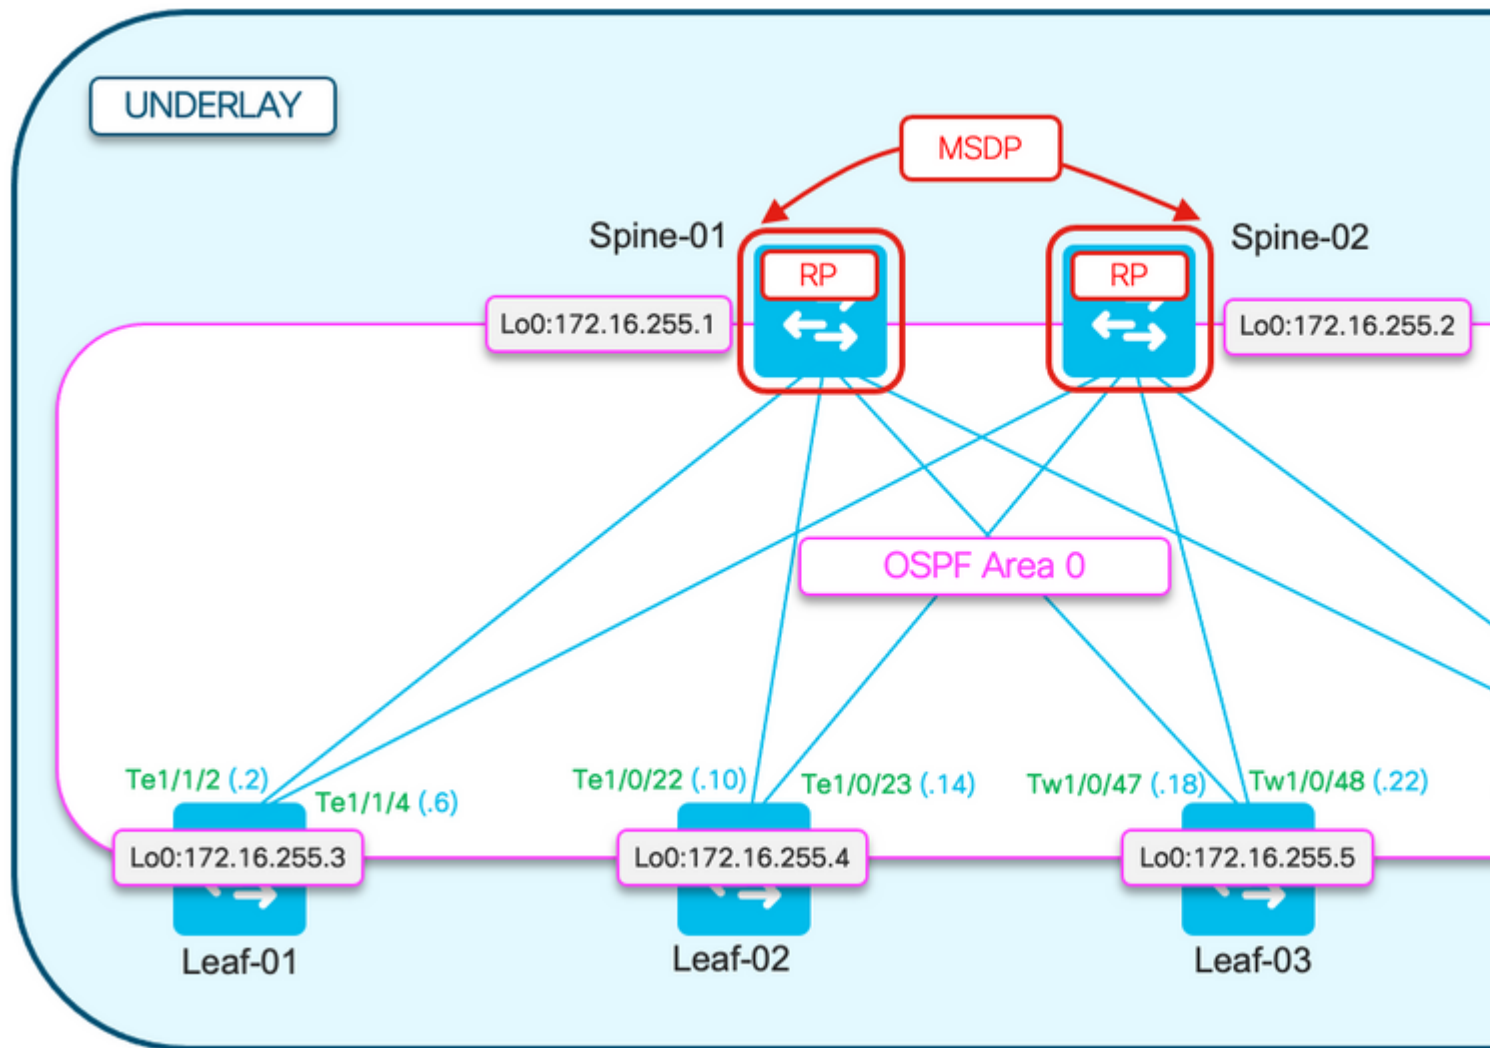
DN	Dynamic Neighbor	Passive mode where BGP listens for connection attempts and forms neighbors
AF	Address Family	Defines which type prefixes and routing info BGP handles
AS	Autonomous System	A set of Internet routable IP prefixes that belong to a network or a collection of networks that are all managed, controlled and supervised by a single entity or organization
EVPN	Ethernet Virtual Private Network	Extension that allows BGP to transport Layer 2 MAC and Layer 3 IP information is EVPN and uses Multi-Protocol Border Gateway Protocol (MP-BGP) as the protocol to distribute reachability information that pertains to the VXLAN overlay network.
MVPN	Multicast Virtual Private Network	Multicast VPN (MVPN) feature provides the ability to support multicast over a Layer 3 Virtual Private Network (VPN)
VXLAN	Virtual Extensible LAN (Local Area	VXLAN is designed to overcome the inherent limitations of VLANs and STP. It is a proposed IETF standard [RFC 7348] to provide the same

	Network)	Ethernet Layer 2 network services as VLANs do, but with greater flexibility. Functionally, it is a MAC-in- UDP encapsulation protocol that runs as a virtual overlay on a Layer 3 underlay network.
--	----------	---

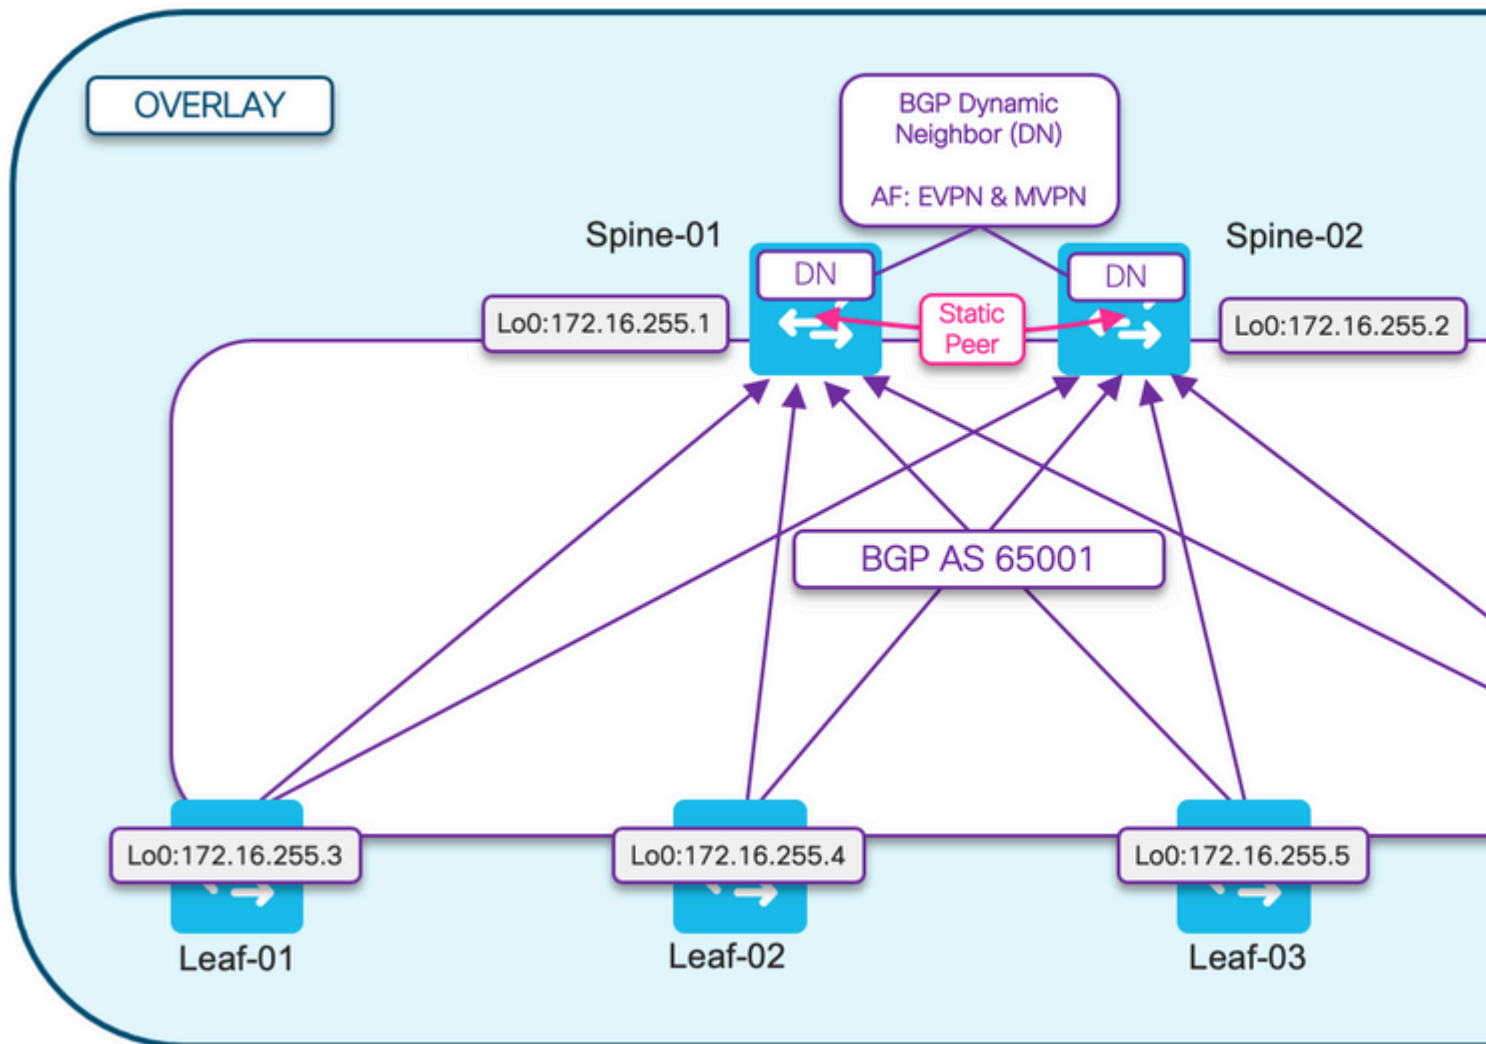
Configure

Network Diagram

Underlay (OSPF)



Overlay (BGP)



Spine-01

<#root>

Spine-01#

sh run | section r b

router bgp 65001

bgp router-id 172.16.255.1

bgp log-neighbor-changes

bgp listen range 172.16.255.0/29 peer-group DN-GROUP <-- Listen for connections from peers in this subnet

bgp listen block 172.16.255.2 <-- Excluded peer RR from DN subnet range to allow

bgp listen limit 4 <-- Limit the number of peerings allowed

no bgp default ipv4-unicast

neighbor DN-GROUP peer-group

neighbor DN-GROUP remote-as 65001

<-- Peers only allowed in this AS

neighbor DN-GROUP update-source Loopback0

neighbor 172.16.255.2 remote-as 65001

<-- Static Neighbor config

```

neighbor 172.16.255.2 update-source Loopback0          <-- Static Neighbor config

!
address-family ipv4
exit-address-family
!

address-family ipv4 mvpn
  neighbor DN-GROUP activate                          <--- Dynamic peer group must be activated in MVPN AF

  neighbor DN-GROUP send-community both
  neighbor DN-GROUP route-reflector-client

  neighbor 172.16.255.2 activate                      <--- Static peering activation in MVPN AF

  neighbor 172.16.255.2 send-community extended
exit-address-family
!

address-family l2vpn evpn
  neighbor DN-GROUP activate                          <--- Dynamic peer group must be activated in EVPN AF

  neighbor DN-GROUP send-community both
  neighbor DN-GROUP route-reflector-client

  neighbor 172.16.255.2 activate                      <-- Static peering activation in EVPN AF

  neighbor 172.16.255.2 send-community extended
exit-address-family

```

Note: It is possible to have static and dynamic neighbors as part of same peer group, but static neighbors need to be activated individually.

Tip: It is recommended to limit the number of peers to number of neighbors you expect to have to prevent unexpected peers. Set both the subnet range and listen limit as tightly as possible.

Leaf-01

```

<#root>

Leaf-01#

show run | sec r bgp

router bgp 65001
  bgp router-id 172.16.255.3
  bgp log-neighbor-changes
  no bgp default ipv4-unicast
  neighbor 172.16.255.1 remote-as 65001
  neighbor 172.16.255.1 update-source Loopback0
  neighbor 172.16.255.2 remote-as 65001
  neighbor 172.16.255.2 update-source Loopback0
!
address-family ipv4

```

```

exit-address-family
!

address-family ipv4 mvpn

neighbor 172.16.255.1 activate <-- Route Reflector 1

neighbor 172.16.255.1 send-community both
neighbor 172.16.255.2 activate <-- Route Reflector 2

neighbor 172.16.255.2 send-community both
exit-address-family
!

address-family l2vpn evpn

neighbor 172.16.255.1 activate <-- Route Reflector 2

neighbor 172.16.255.1 send-community both
neighbor 172.16.255.2 activate <-- Route Reflector 2

neighbor 172.16.255.2 send-community both
exit-address-family

```

Note: The configuration for the other Spine Route Reflector & VTEPs are the same, so are not repeated in this section

Note: To peer RRs that use DN you must use static configuration. They are in passive mode, and do not generate OPEN messages toward each other

Optional Configurations

Alternate-AS

Allows BGP sessions to be established with peers in the specified subnet that have an AS number listed (up to five alternate eBGP AS are supported for total of 6 ASes).

```

<#root>

Spine-01#

show run | section r b

router bgp 65001
  bgp listen range 172.16.255.0/29 peer-group DN-GROUP
  bgp listen limit 20
  neighbor DN-GROUP peer-group

neighbor DN-GROUP remote-as 1 alternate-as 2 3 4 5 6<-- Allow peering with more than one AS

```

neighbor DN-GROUP update-source Loopback0

Note: This might not be enough alternates for high-scale EVPN use-cases (which might require as many as 1000 remote ASes per Dynamic Neighbor peer-group). Future releases shall increase this limit to accommodate such deployments.

Verify

Spine (RR)

Verify the Spine, who is the BGP listener, has formed dynamic peers

- Spine-02 looks the same, so only showing Spine-01 in this example output

```
<#root>
```

```
Spine-01#
```

```
show ip bgp all summary
```

```
For address family: L2VPN E-VPN
```

```
<-- Address Family type EVPN
```

```
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 64, main routing table version 64
19 network entries using 7296 bytes of memory
24 path entries using 5568 bytes of memory
13/10 BGP path/bestpath attribute entries using 3848 bytes of memory
13 BGP extended community entries using 1670 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 18382 total bytes of memory
BGP activity 22/3 prefixes, 56/32 paths, scan interval 60 secs
19 networks peaked at 19:07:48 May 10 2023 UTC (00:48:59.007 ago)
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
*									
172.16.255.3	4	65001	286	328	64	0	0	04:05:53	9
*									
172.16.255.4	4	65001	309	340	64	0	0	04:05:53	13
*									
172.16.255.5	4	65001	276	315	64	0	0	04:05:43	1
*									
172.16.255.6	4	65001	276	313	64	0	0	04:05:53	1

```
* Dynamically created based on a listen range command
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
----------	---	----	---------	---------	--------	-----	------	---------	--------------

```
Dynamically created neighbors: 4, Subnet ranges: 1
```

```
BGP peergroup DN-GROUP listen range group members:
```

```

172.16.255.0/29 <-- Peering subnet configured
For address family: MVPNv4 Unicast <-- Address Family type MVPN
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 1, main routing table version 1

Neighbor      V          AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
*
172.16.255.3  4          65001   286    328      1     0     0 04:05:53      0
*
172.16.255.4  4          65001   309    340      1     0     0 04:05:54      0
*
172.16.255.5  4          65001   276    315      1     0     0 04:05:43      0
*
172.16.255.6  4          65001   276    313      1     0     0 04:05:53      0
* Dynamically created based on a listen range command
Dynamically created neighbors: 4, Subnet ranges: 1
BGP peergroup DN-GROUP listen range group members:
 172.16.255.0/29 <-- Peering subnet configured

Total dynamically created neighbors: 4/(4 max), Subnet ranges: 1 <-- 4 total neighbors from the max

```

Note: You can use per-AF BGP commands to see only a single AF's info:

- show bgp l2vpn evpn summary
 - show ip bgp ipv4 mvpn all summary
-

Observe Spine-01 (RR) TCP ports:

- Because Spine is passive, its local port is always the fixed BGP port 179 whereas the VTEPs show ephemeral ports
- 172.16.255.1 > 172.16.255.2 is defined as a static peer, but this RR happened to receive the OPEN so used port 179 as well

```
<#root>
```

```
Spine-01#
```

```
show tcp brief
```

```
TCB
```

```
Local Address
```

```
Foreign Address (state)
```

```
72632453A5C0 172.16.255.1.
```



```

179
    172.16.255.5.15285          ESTAB
<-- Dynamic peer (Leaf-03)
7263245C73C0 172.16.255.1.

179
    172.16.255.6.52693          ESTAB
<-- Dynamic peer (Border Leaf)
7263245B3D30 172.16.255.1.

179
    172.16.255.4.44091          ESTAB
<-- Dynamic peer (Leaf-02)
7263272DE460 172.16.254.1.22430      172.16.254.2.

639
    ESTAB
<-- MSDP peer      (Spine-02)
7263246DD3A8 172.16.255.1.179          172.16.255.2.21866      ESTAB
<-- Static Peer   (Spine-02 RR)
7263245BDAC0 172.16.255.1.

179
    172.16.255.3.31383          ESTAB
<-- Dynamic peer (Leaf-01)

```

Troubleshoot

DN Peering

If the peer group is added under the L2VPN EVPN address-family, it does not take effect if the peer-group is activated under another AF.

- In this case, peer-group is present under MVPN then activated under EVPN.
- The EVPN neighbors do not form until a hard clear is performed.

```
<#root>
```

```
Spine-01(config)#
```

```
router bgp
```

```
65001
```

```
Spine-01(config-router)#
```

```
address-family l2vpn evpn
```

```
Spine-01(config-router-af)#
```

```
neighbor DN-GROUP activate
```

```
Spine-01(config-router-af)#
```

```
end
```

```
% BGP: Dynamic peergroup DN-GROUP does not reset session by default and requires manual reset if this pe
```

```
Spine-01#sh bgp l2vpn evpn summary
```

```
<--- Neighbor info missing
```

```
BGP peergroup DN-GROUP listen range group members:
```

```
172.16.255.0/29
```

```
Total dynamically created neighbors: 4/(4 max), Subnet ranges: 1
```

```
Spine-01#
```

```
clear ip bgp * <-- Hard clear
```

```
%BGP-3-NOTIFICATION_MANY: sent to 5 sessions 6/4 (Administrative Reset) for all peers <-- Resetting all
```

```
%BGP_SESSION-5-ADJCHANGE: neighbor *172.16.255.3 MVPNv4 Unicast topology base removed from session Neig
```

```
%BGP-5-ADJCHANGE: neighbor *172.16.255.3 Down User reset
```

```
%BGP_SESSION-5-ADJCHANGE: neighbor *172.16.255.4 MVPNv4 Unicast topology base removed from session Neig
```

```
%BGP-5-ADJCHANGE: neighbor *172.16.255.4 Down User reset
```

```
%BGP_SESSION-5-ADJCHANGE: neighbor *172.16.255.5 MVPNv4 Unicast topology base removed from session Neig
```

```
%BGP-5-ADJCHANGE: neighbor *172.16.255.5 Down User reset
```

```
%BGP_SESSION-5-ADJCHANGE: neighbor *172.16.255.6 MVPNv4 Unicast topology base removed from session Neig
```

```
%BGP-5-ADJCHANGE: neighbor *172.16.255.6 Down User reset
```

```
%BGP-5-ADJCHANGE: neighbor 172.16.255.2 Down User reset
```

```
%BGP_SESSION-5-ADJCHANGE: neighbor 172.16.255.2 MVPNv4 Unicast topology base removed from session User
```

```
%BGP-5-ADJCHANGE: neighbor *172.16.255.3 Up
```

```
<-- Peers coming back up for both MVPN and EVPN
```

```
%BGP-5-ADJCHANGE: neighbor *172.16.255.5 Up
```

```
%BGP-5-ADJCHANGE: neighbor *172.16.255.4 Up
```

```
%BGP-5-ADJCHANGE: neighbor *172.16.255.6 Up
```

```
Spine-01#
```

```
sh bgp l2vpn evpn all summary
```

```
BGP router identifier 172.16.255.1, local AS number 65001
```

```
BGP table version is 61, main routing table version 61
```

```
41 network entries using 15744 bytes of memory
```

```
48 path entries using 11136 bytes of memory
```

```
25/22 BGP path/bestpath attribute entries using 7400 bytes of memory
```

```
20 BGP extended community entries using 2120 bytes of memory
```

```
0 BGP route-map cache entries using 0 bytes of memory
```

```
0 BGP filter-list cache entries using 0 bytes of memory
```

```
BGP using 36400 total bytes of memory
```

```
BGP activity 145/104 prefixes, 22148/22100 paths, scan interval 60 secs
```

```
41 networks peaked at 21:07:01 May 11 2023 UTC (4d21h ago)
```

```

Neighbor      V          AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
*172.16.255.3
  4          65001    26     12     61     0    0 00:05:03    16
<-- Dynamic Peers now show up in EVPN (* = dynamically created)
*172.16.255.4
  4          65001    27     10     61     0    0 00:05:08    19
*172.16.255.5
  4          65001    14     12     61     0    0 00:05:05     6
*172.16.255.6
  4          65001    17     10     61     0    0 00:05:10     7

```

```

* Dynamically created based on a listen range command
Neighbor      V          AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd

```

Dynamically created neighbors: 4

, Subnet ranges: 1

BGP peergroup DN-GROUP listen range group members:
172.16.255.0/29

Total dynamically created neighbors: 4/(4 max)

, Subnet ranges: 1

Note: A hard clear on the neighbor is required in order for the configuration changes to take effect. This is consistent with the current behavior for address-families that support BGP Dynamic Neighbor.

DN Debugs

If there is an issue with Dynamic Peers you can use debugs to see more about the problem

```
<#root>
```

```
Spine-01#
```

```
debug bgp all range detail
```

```
Spine-01#
```

```
sh debugging
```

```
IP routing:
```

```
BGP dynamic range debugging is on with detail (Dynamic neighbor details only) <-- Only debug DN type nei
```

<#root>

Spine-01(config)#

router bgp 65001

Spine-01(config-router)#

address-family l2vpn evpn

Spine-01(config-router-af)#

neighbor DN-GROUP activate

% BGP: Dynamic peergroup DN-GROUP does not reset session by default and requires manual reset if this pe

*May 16 18:29:53.744:

BGP:DN: Peergroup DN-GROUP is activated for Dynamic Neighbors for L2VPN E-VPN <-- Debug event

Spine-01#

clear ip bgp *

*May 16 18:31:34.806: %BGP-3-NOTIFICATION_MANY: sent to 5 sessions 6/4 (Administrative Reset) for all pe

*May 16 18:31:34.806: BGP:DN: free: *172.16.255.3

*May 16 18:31:34.806: BGP:DN: free: *172.16.255.4

*May 16 18:31:34.806: BGP:DN: free: *172.16.255.5

*May 16 18:31:34.807: BGP:DN: free: *172.16.255.6

*May 16 18:31:43.855:

BGP:DN: validating TCP SYN from 172.16.255.5: in range 172.16.255.0/29, val =0, peer-group = DN-GROUP <-

*May 16 18:31:43.856:

BGP: 172.16.255.5 passive open to 172.16.255.1

*May 16 18:31:43.856:

BGP:DN: passive open ACCEPT - TCP session : tcb=0x7263248098B8,context 0x726317305E20, ctx_group name 17

*May 16 18:31:43.856:

BGP:DN: group->prefix_length=24, group->prefix=172.16.255.0

*May 16 18:31:43.856:

BGP:DN: 172.16.255.5 get_rc_shim_transport_group: range group =172.16.255.0/29, group address = 0x726317

*May 16 18:31:43.856:

BGP:DN: init dynamic neighbor for peergroup DN-GROUP in L2VPN Evpn

*May 16 18:31:43.856:

BGP:DN: Created a new neighbor *172.16.255.5 in range 172.16.255.0/29, peer-group DN-GROUP, count= 0 <-

*May 16 18:31:43.856:

BGP:DN: accepted: nbr *172.16.255.5/7263245A3508 peer-grp: DN-GROUP/726317306460 <-

Note: The other 3 peers have the same debug result, so not shown here.

Static Neighbor

If a neighbor was previously created dynamically it cannot be converted to static.

```
<#root>
Spine-01(config-router)#
neighbor 172.16.255.3 remote-as 65001

% Cannot configure for dynamically created neighbor
```

You can exclude a neighbor in the DN range with this command under the router BGP mode. This allows a mix of static and dynamic peers in the DN subnet range

```
<#root>
router bgp 65001
  bgp listen block 172.16.255.3
```

Related Information

- [BGP EVPN VXLAN Configuration Guide, Cisco IOS XE Dublin 17.11.x \(Catalyst 9500 Switches\)](#)
- [Technical Support & Documentation - Cisco Systems](#)