

Configure BGP Confederations on Nexus 9000

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Background Information](#)

[Configure Example](#)

[Network Diagram](#)

[Configuration](#)

[Verify](#)

Introduction

This document describes the procedure to configure Border Gateway Protocol (BGP) Confederations on Nexus 9000 series.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Nexus Switches
- BGP

Components Used

This document is not restricted to specific software and hardware versions.

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

The mechanism BGP uses to avoid routing loops is adding its own Autonomous System Number (ASN) to the AS PATH attribute. For an IBGP topology, the updates do not modify the AS-PATH attribute as they belong to the same ASN, which can cause a routing loop.

BGP uses its split-horizon rule that states a BGP device receives a prefix via an IBGP peer, and cannot advertise the prefix to another IBGP neighbor. This rule forces you to have a full-mesh topology to exchange prefixes with every IBGP neighbor. For big networks to have a full-meshed scenario is not scalable as it uses excessive resources to create the peerings.

The alternatives to a full-meshed topology for iBGP are:

- Route Reflection
- Confederations

This document focus on the usage of BGP Confederations only.


BGP Confederation eliminates the need for a full-mesh topology by dividing an ASN into multiple sub-autonomous systems and grouping them into a confederation.


Configure Example


To check the BGP Confederation configuration guide, see [Advance BGP Configuration Guide Nexus 9000](#).


Network Diagram



 **Note:** In the topology, Nexus 1, Nexus 2, Nexus 3, and Nexus 4 are part of the same Confederation ASN 64513, dividing the confederation into 2 smaller ASN (64512 and 64514)

 **Note:** Nexus 1 and Nexus 3 establish an iBGP neighborhood.

 **Note:** Nexus 2 and Nexus 4 establish an iBGP neighborhood.

 **Note:** Nexus 1 and Nexus 2 establish an eBGP neighborhood.

Configuration



Note: The Confederation configuration causes a reset to every established bgp peer. This configuration must be done under a Maintenance Window.

Nexus 1

```
Nexus1# show running-config bgp
```

```
!Command: show running-config bgp
!Running configuration last done at: Thu Jul 20 21:13:08 2023
!Time: Thu Jul 20 23:28:41 2023
```

```
version 10.2(5) Bios:version 05.47
feature bgp
```

```
router bgp 64512                                --> Local ASN
  confederation identifier 64513                --> BGP Confederation identifier
  confederation peers 64514                    --> ASN that belong to the Confederation
  log-neighbor-changes
  neighbor 10.0.5.2
    remote-as 64515
    address-family ipv4 unicast
  neighbor 172.16.255.2
    remote-as 64514
    update-source loopback0
    address-family ipv4 unicast
  neighbor 172.16.255.3
    remote-as 64512
    log-neighbor-changes
    update-source loopback0
    address-family ipv4 unicast
```

Nexus 2

```
Nexus2# show running-config bgp
```

```
!Command: show running-config bgp
!Running configuration last done at: Thu Jul 20 21:12:54 2023
!Time: Thu Jul 20 23:31:04 2023
```

```
version 10.2(5) Bios:version 05.47
feature bgp
```

```
router bgp 64514                                --> Local ASN
  confederation identifier 64513                --> BGP Confederation identifier
  confederation peers 64512                    --> ASN that belong to the Confederation
  log-neighbor-changes
  neighbor 172.16.255.1
    remote-as 64512
    update-source loopback0
    address-family ipv4 unicast
  neighbor 172.16.255.4
    remote-as 64514
```

```
update-source loopback0
address-family ipv4 unicast
```

Nexus 3

```
Nexus3# show running-config bgp
```

```
!Command: show running-config bgp
!Running configuration last done at: Thu Jul 20 21:14:54 2023
!Time: Thu Jul 20 23:31:45 2023
```

```
version 10.2(5) Bios:version 05.47
feature bgp
```

```
router bgp 64512                                --> Local ASN
  confederation identifier 64513                --> BGP Confederation identifier
  log-neighbor-changes
  neighbor 172.16.255.1
    remote-as 64512
    log-neighbor-changes
    update-source loopback0
    address-family ipv4 unicast
```

Nexus 4

```
Nexus4# show running-config bgp
```

```
!Command: show running-config bgp
!Running configuration last done at: Thu Jul 20 15:11:15 2023
!Time: Thu Jul 20 17:32:58 2023
```

```
version 10.2(5) Bios:version 05.47
feature bgp
```

```
router bgp 64514                                --> Local ASN
  confederation identifier 64513                --> BGP Confederation identifier
  log-neighbor-changes
  address-family ipv4 unicast
    network 172.30.1.0/24
    network 172.30.2.0/24
    network 172.30.3.0/24
    network 172.30.4.0/24
    network 172.30.5.0/24
    network 172.30.6.0/24
    network 172.30.7.0/24
    network 172.30.8.0/24
  neighbor 172.16.255.2
    remote-as 64514
    update-source loopback0
    address-family ipv4 unicast
```

Nexus 5

```
Nexus5# show running-config bgp
```

```
!Command: show running-config bgp
!Running configuration last done at: Thu Jul 20 21:13:32 2023
!Time: Thu Jul 20 23:33:34 2023
```

```
version 10.2(5) Bios:version 05.47
feature bgp
```

```
router bgp 64515                                --> Local ASN
  log-neighbor-changes
  address-family ipv4 unicast
    network 192.168.1.0/24
    network 192.168.2.0/24
    network 192.168.3.0/24
    network 192.168.4.0/24
    network 192.168.5.0/24
    network 192.168.6.0/24
    network 192.168.7.0/24
    network 192.168.8.0/24
  neighbor 10.0.5.1
    remote-as 64513
    address-family ipv4 unicast
```

Verify

Nexus 1

Step 1: BGP neighbors are established and exchanging prefixes.

```
Nexus1# show ip bgp summary
```

```
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 172.16.255.1, local AS number 64512
BGP table version is 28, IPv4 Unicast config peers 3, capable peers 3
16 network entries and 16 paths using 4416 bytes of memory
BGP attribute entries [2/704], BGP AS path entries [2/12]
BGP community entries [0/0], BGP clusterlist entries [0/0]
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.5.2	4	64515	328	321	28	0	0	05:00:34	8
172.16.255.2	4	64514	270	264	28	0	0	04:16:51	8
172.16.255.3	4	64512	318	318	28	0	0	05:12:28	0

Step 2: The prefixes learned by the confederation peer are identified by looking the status *c-confed*.

```
Nexus1# show ip bgp
```

```
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 28, Local Router ID is 172.16.255.1
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>c172.30.1.0/24	172.16.255.4		100	0	(64514) i
*>c172.30.2.0/24	172.16.255.4		100	0	(64514) i
*>c172.30.3.0/24	172.16.255.4		100	0	(64514) i
*>c172.30.4.0/24	172.16.255.4		100	0	(64514) i
*>c172.30.5.0/24	172.16.255.4		100	0	(64514) i
*>c172.30.6.0/24	172.16.255.4		100	0	(64514) i
*>c172.30.7.0/24	172.16.255.4		100	0	(64514) i
*>c172.30.8.0/24	172.16.255.4		100	0	(64514) i
*>e192.168.1.0/24	10.0.5.2			0	64515 i
*>e192.168.2.0/24	10.0.5.2			0	64515 i
*>e192.168.3.0/24	10.0.5.2			0	64515 i
*>e192.168.4.0/24	10.0.5.2			0	64515 i
*>e192.168.5.0/24	10.0.5.2			0	64515 i
*>e192.168.6.0/24	10.0.5.2			0	64515 i
*>e192.168.7.0/24	10.0.5.2			0	64515 i
*>e192.168.8.0/24	10.0.5.2			0	64515 i

Nexus 2

Step 1: BGP neighbors are established and exchanging prefixes.

```
Nexus2# show ip bgp summary
```

```
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 172.16.255.2, local AS number 64514
BGP table version is 27, IPv4 Unicast config peers 2, capable peers 2
16 network entries and 16 paths using 3904 bytes of memory
BGP attribute entries [2/344], BGP AS path entries [1/12]
BGP community entries [0/0], BGP clusterlist entries [0/0]
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.255.1	4	64512	274	265	27	0	0	04:20:26	8
172.16.255.4	4	64514	269	262	27	0	0	04:16:00	8

Step 2: The prefixes learned by the confederation peer are identified by looking the status *c-confed*.

```
Nexus2# show ip bgp
```

```
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 27, Local Router ID is 172.16.255.2
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i172.30.1.0/24	172.16.255.4		100	0	i
*>i172.30.2.0/24	172.16.255.4		100	0	i
*>i172.30.3.0/24	172.16.255.4		100	0	i
*>i172.30.4.0/24	172.16.255.4		100	0	i
*>i172.30.5.0/24	172.16.255.4		100	0	i
*>i172.30.6.0/24	172.16.255.4		100	0	i
*>i172.30.7.0/24	172.16.255.4		100	0	i
*>i172.30.8.0/24	172.16.255.4		100	0	i
*>c192.168.1.0/24	10.0.5.2		100	0	(64512) 64515 i

```
*>c192.168.2.0/24      10.0.5.2          100          0 (64512) 64515 i
*>c192.168.3.0/24      10.0.5.2          100          0 (64512) 64515 i
*>c192.168.4.0/24      10.0.5.2          100          0 (64512) 64515 i
*>c192.168.5.0/24      10.0.5.2          100          0 (64512) 64515 i
*>c192.168.6.0/24      10.0.5.2          100          0 (64512) 64515 i
*>c192.168.7.0/24      10.0.5.2          100          0 (64512) 64515 i
*>c192.168.8.0/24      10.0.5.2          100          0 (64512) 64515 i
```

Nexus 3

Step 1: BGP neighbor is established and exchanging prefixes.

```
Nexus3# show ip bgp summary
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 172.16.255.3, local AS number 64512
BGP table version is 34, IPv4 Unicast config peers 1, capable peers 1
16 network entries and 16 paths using 4416 bytes of memory
BGP attribute entries [2/704], BGP AS path entries [2/12]
BGP community entries [0/0], BGP clusterlist entries [0/0]
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/
172.16.255.1	4	64512	332	323	34	0	0	05:17:52	16

Step 2: From Nexus 3 perspective, the prefixes are learned as normal iBGP prefixes. From its perspective, it does not have any confederation neighbor, and its only peer is an iBGP peer.

```
Nexus3# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 34, Local Router ID is 172.16.255.3
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-i
njected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - b
est2
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i172.30.1.0/24	172.16.255.4		100	0	(64514) i
*>i172.30.2.0/24	172.16.255.4		100	0	(64514) i
*>i172.30.3.0/24	172.16.255.4		100	0	(64514) i
*>i172.30.4.0/24	172.16.255.4		100	0	(64514) i
*>i172.30.5.0/24	172.16.255.4		100	0	(64514) i
*>i172.30.6.0/24	172.16.255.4		100	0	(64514) i
*>i172.30.7.0/24	172.16.255.4		100	0	(64514) i
*>i172.30.8.0/24	172.16.255.4		100	0	(64514) i
*>i192.168.1.0/24	10.0.5.2		100	0	64515 i
*>i192.168.2.0/24	10.0.5.2		100	0	64515 i
*>i192.168.3.0/24	10.0.5.2		100	0	64515 i
*>i192.168.4.0/24	10.0.5.2		100	0	64515 i
*>i192.168.5.0/24	10.0.5.2		100	0	64515 i
*>i192.168.6.0/24	10.0.5.2		100	0	64515 i
*>i192.168.7.0/24	10.0.5.2		100	0	64515 i
*>i192.168.8.0/24	10.0.5.2		100	0	64515 i

Nexus 4

Step 1: BGP neighbor is established and exchanging prefixes.

```
Nexus4# show ip bgp summary
```

```
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 172.16.255.4, local AS number 64514
BGP table version is 18, IPv4 Unicast config peers 1, capable peers 1
16 network entries and 16 paths using 4416 bytes of memory
BGP attribute entries [2/704], BGP AS path entries [1/12]
BGP community entries [0/0], BGP clusterlist entries [0/0]
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.255.2	4	64514	268	267	18	0	0	04:21:40	8

Step 2: From Nexus 4 perspective, the prefixes are learned as normal iBGP prefixes. From its perspective, it does not have any confederation neighbor, and its only peer is an iBGP peer.

```
Nexus4# show ip bgp
```

```
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 18, Local Router ID is 172.16.255.4
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>172.30.1.0/24	0.0.0.0		100	32768	i
*>172.30.2.0/24	0.0.0.0		100	32768	i
*>172.30.3.0/24	0.0.0.0		100	32768	i
*>172.30.4.0/24	0.0.0.0		100	32768	i
*>172.30.5.0/24	0.0.0.0		100	32768	i
*>172.30.6.0/24	0.0.0.0		100	32768	i
*>172.30.7.0/24	0.0.0.0		100	32768	i
*>172.30.8.0/24	0.0.0.0		100	32768	i
*>i192.168.1.0/24	10.0.5.2		100	0 (64512)	64515 i
*>i192.168.2.0/24	10.0.5.2		100	0 (64512)	64515 i
*>i192.168.3.0/24	10.0.5.2		100	0 (64512)	64515 i
*>i192.168.4.0/24	10.0.5.2		100	0 (64512)	64515 i
*>i192.168.5.0/24	10.0.5.2		100	0 (64512)	64515 i
*>i192.168.6.0/24	10.0.5.2		100	0 (64512)	64515 i
*>i192.168.7.0/24	10.0.5.2		100	0 (64512)	64515 i
*>i192.168.8.0/24	10.0.5.2		100	0 (64512)	64515 i

Nexus 5

Step 1: BGP neighbor is established and exchanging prefixes. In this case the only neighbor is Nexus 3, and this is an eBGP neighbor.

```
Nexus5# show ip bgp summary
```

```
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 172.16.255.5, local AS number 64515
BGP table version is 18, IPv4 Unicast config peers 1, capable peers 1
```


16 network entries and 16 paths using 3904 bytes of memory
 BGP attribute entries [2/344], BGP AS path entries [1/6]
 BGP community entries [0/0], BGP clusterlist entries [0/0]

```
Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.5.1      4 64513    317    317     18   0   0 05:11:26 8
```

Step 2: From Nexus 5 perspective, the prefixes received from Nexus 3 are received using the AS PATH of the Confederation ASN, but are received as normal eBGP prefixes.

```
Nexus5# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 18, Local Router ID is 172.16.255.5
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-i
njected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - b
est2
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>e172.30.1.0/24	10.0.5.1			0	64513 i
*>e172.30.2.0/24	10.0.5.1			0	64513 i
*>e172.30.3.0/24	10.0.5.1			0	64513 i
*>e172.30.4.0/24	10.0.5.1			0	64513 i
*>e172.30.5.0/24	10.0.5.1			0	64513 i
*>e172.30.6.0/24	10.0.5.1			0	64513 i
*>e172.30.7.0/24	10.0.5.1			0	64513 i
*>e172.30.8.0/24	10.0.5.1			0	64513 i
*>1192.168.1.0/24	0.0.0.0		100	32768	i
*>1192.168.2.0/24	0.0.0.0		100	32768	i
*>1192.168.3.0/24	0.0.0.0		100	32768	i
*>1192.168.4.0/24	0.0.0.0		100	32768	i
*>1192.168.5.0/24	0.0.0.0		100	32768	i
*>1192.168.6.0/24	0.0.0.0		100	32768	i
*>1192.168.7.0/24	0.0.0.0		100	32768	i
*>1192.168.8.0/24	0.0.0.0		100	32768	i

 **Note:** Nexus 5 is not aware of the Confederation used between Nexus 1 - Nexus 4. From its perspective these are normal eBGP prefixes.
