

# Configure BGP Route Reflectors on Nexus 9000

## Contents

---

### [Introduction](#)

### [Prerequisites](#)

#### [Requirements](#)

#### [Components Used](#)

### [Background Information](#)

### [Configure](#)

#### [Network Diagram](#)

#### [Objective](#)

#### [Initial Verification](#)

### [Configurations](#)

### [Validation](#)

### [Summary](#)

### [Related Information](#)

---

## Introduction

This document describes the procedure to configure Border Gateway Protocol (BGP) Route Reflectors 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 on production, do ensure that you understand the potential impact of any command you plan to apply.

## Background Information

The mechanism BGP uses to avoid routing loops is by 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.

If BGP prefix is received from an IBGP peer, then this prefix cannot be advertised to another IBGP neighbor. BGP uses its split-horizon rule to prevent loops within the same ASN.

This rule forces you to have a full-mesh topology to exchange prefixes with every IBGP neighbor. For big networks having a full-meshed scenario is not scalable as it uses excessive resources to create the peering's.

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

- Route Reflection
- Confederations

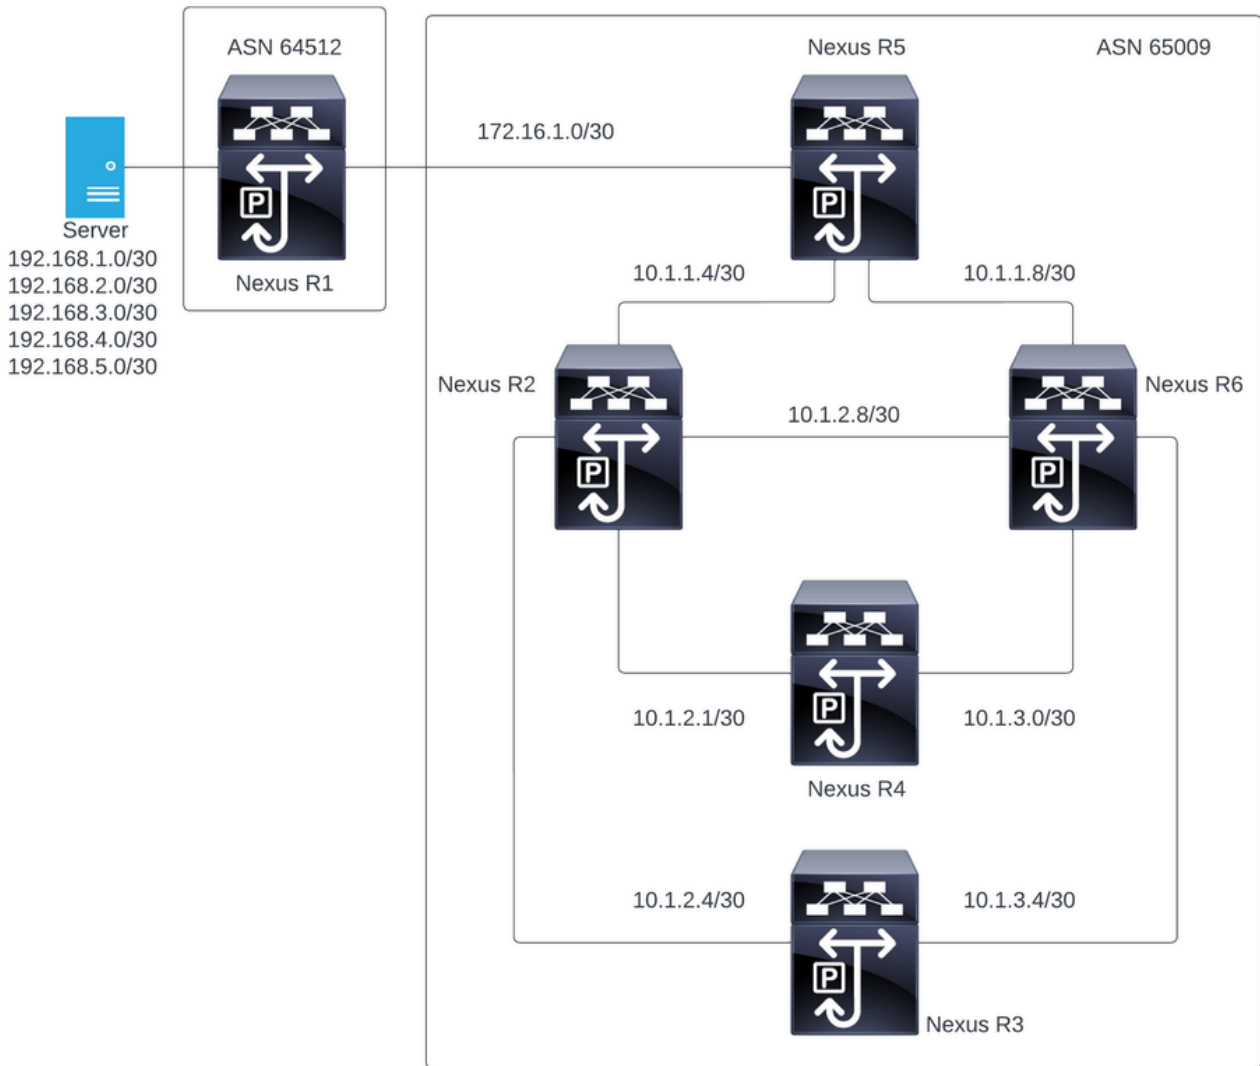
This document focus on the usage of BGP Route Reflector only.

Using Route Reflectors overcomes the full-mesh limitation. A Route Reflector is a designated device within the AS (Autonomous System) that reflects iBGP prefixes to other routers in the same AS.

## **Configure**

To check the BGP Route Reflectors configuration guide, see [Configuring Advanced BGP](#)

## **Network Diagram**



In this figure, the Nexus R2, Nexus R3, Nexus 4, Nexus R5 and Nexus R6 establish an iBGP neighborship. And, the Nexus R1 and Nexus R5 establish an eBGP neighborship.

## Objective

Exchange the networks generated by Nexus R1 to all the Nexus within the same ASN 65000.

## Initial Verification

Nexus R1 performs the advertisement of the prefixes:

```
R1# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 78, Local Router ID is 192.168.1.1
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-i
njected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - b
est2
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>r192.168.1.0/30	0.0.0.0	0	100	32768	?
*>r192.168.2.0/30	0.0.0.0	0	100	32768	?
*>r192.168.3.0/30	0.0.0.0	0	100	32768	?
*>r192.168.4.0/30	0.0.0.0	0	100	32768	?
*>r192.168.5.0/30	0.0.0.0	0	100	32768	?

R1#

Nexus R5 receives the advertisement from Nexus R1 which is an eBGP session.

Nexus R5

```
R5# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 123, Local Router ID is 172.16.1.2
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
*>r10.1.1.4/30	0.0.0.0	0	100	32768	?
*>r10.1.1.8/30	0.0.0.0	0	100	32768	?
*>r172.16.1.0/30	0.0.0.0	0	100	32768	?
*>e192.168.1.0/30	172.16.1.1	0		0	64512 ?
*>e192.168.2.0/30	172.16.1.1	0		0	64512 ?
*>e192.168.3.0/30	172.16.1.1	0		0	64512 ?
*>e192.168.4.0/30	172.16.1.1	0		0	64512 ?
*>e192.168.5.0/30	172.16.1.1	0		0	64512 ?

R5#

Nexus R5 advertise the prefixes to its iBGP peers Nexus R2 and R6.

```
R2# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 33, Local Router ID is 10.1.1.6
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
*>i10.1.1.4/30	10.1.1.5	0	100	0	?
*>i10.1.1.8/30	10.1.1.5	0	100	0	?
*>i172.16.1.0/30	10.1.1.5	0	100	0	?
*>i192.168.1.0/30	172.16.1.1	0	100	0	64512 ?
*>i192.168.2.0/30	172.16.1.1	0	100	0	64512 ?
*>i192.168.3.0/30	172.16.1.1	0	100	0	64512 ?
*>i192.168.4.0/30	172.16.1.1	0	100	0	64512 ?
*>i192.168.5.0/30	172.16.1.1	0	100	0	64512 ?

R2#

```
R6# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 33, Local Router ID is 10.1.1.10
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
*>i10.1.1.4/30	10.1.1.9	0	100	0	?
*>i10.1.1.8/30	10.1.1.9	0	100	0	?
*>i172.16.1.0/30	10.1.1.9	0	100	0	?
*>i192.168.1.0/30	172.16.1.1	0	100	0	64512 ?
*>i192.168.2.0/30	172.16.1.1	0	100	0	64512 ?
*>i192.168.3.0/30	172.16.1.1	0	100	0	64512 ?
*>i192.168.4.0/30	172.16.1.1	0	100	0	64512 ?
*>i192.168.5.0/30	172.16.1.1	0	100	0	64512 ?

R6#

By the loop avoidance rule, BGP do not advertise any prefix learned by an iBGP peer to other iBGP peers.

```
R2# show ip bgp neighbors 10.1.2.2 advertised-routes
```

```
Peer 10.1.2.2 routes for address family IPv4 Unicast:
BGP table version is 88, Local Router ID is 10.150.0.1
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
*>r10.1.1.4/30	0.0.0.0	0	100	32768	?
*>r10.1.2.0/30	0.0.0.0	0	100	32768	?
*>r10.1.2.4/30	0.0.0.0	0	100	32768	?
*>r10.1.2.8/30	0.0.0.0	0	100	32768	?
*>r10.150.0.0/30	0.0.0.0	0	100	32768	?

R2#

Nexus R3 and Nexus R4 do not receive the prefixes generated by Nexus R1.

```
R3# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 28, Local Router ID is 10.100.100.1
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
*>l10.100.100.0/24	0.0.0.0		100	32768	i

```
R4# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 6, Local Router ID is 10.200.200.1
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
*>110.200.200.0/24	0.0.0.0		100	32768	i

```
R4#
```


## Configurations

For the above topology, the selected devices to be turned into a Route Reflector are the Nexus R2 and Nexus R6 as this devices have connection to all iBGP peers within the same ASN.

---

 **Note:** The Route Reflector configuration causes a reset to every bgp peer modified. This configuration must be done under a Maintenance Window.

---

 **Note:** By the topology design the Nexus R2 and R6 are made Route Reflectors for ASN 65000.

---

Nexus R1

```
R1# show run bgp
```

```
!Command: show running-config bgp
!Running configuration last done at: Wed Jan 31 02:43:31 2024
!Time: Wed Jan 31 23:24:28 2024
```

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

```
router bgp 64512
  address-family ipv4 unicast
    redistribute direct route-map REDISTRIBUTE_LOCAL
  neighbor 172.16.1.2
    remote-as 65000
    address-family ipv4 unicast
```

```
R1#
```

Nexus R2

```
R2# show run bgp
```

```
!Command: show running-config bgp
!Running configuration last done at: Wed Jan 31 03:10:49 2024
```

!Time: Wed Jan 31 23:24:15 2024

version 10.2(5) Bios:version 05.47  
feature bgp

```
router bgp 65000
  cluster-id 10.150.0.1
  address-family ipv4 unicast
    redistribute direct route-map ALLOW
  neighbor 10.1.1.5
    remote-as 65000
    address-family ipv4 unicast
      route-reflector-client
  neighbor 10.1.2.2
    remote-as 65000
    address-family ipv4 unicast
      route-reflector-client
  neighbor 10.1.2.6
    remote-as 65000
    address-family ipv4 unicast
      route-reflector-client
  neighbor 10.1.2.10
    remote-as 65000
    address-family ipv4 unicast
      route-reflector-client
```

R2#

Nexus R3

R3# show run bgp

```
!Command: show running-config bgp
!Running configuration last done at: Wed Jan 31 02:49:05 2024
!Time: Wed Jan 31 23:10:07 2024
```

version 10.2(5) Bios:version 05.47  
feature bgp

```
router bgp 65000
  address-family ipv4 unicast
    network 10.100.100.0/24
  neighbor 10.1.2.5
    remote-as 65000
  address-family ipv4 unicast
  neighbor 10.1.3.5
    remote-as 65000
  address-family ipv4 unicast
```

R3#

Nexus R4

R4# show run bgp

```
!Command: show running-config bgp
!Running configuration last done at: Wed Jan 31 02:56:11 2024
!Time: Wed Jan 31 23:00:44 2024
```

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

```
router bgp 65000
  address-family ipv4 unicast
    network 10.200.200.0/24
  neighbor 10.1.2.1
    remote-as 65000
  address-family ipv4 unicast
  neighbor 10.1.3.1
    remote-as 65000
  address-family ipv4 unicast
```

R4#

Nexus R5

R5# show run bgp

```
!Command: show running-config bgp
!Running configuration last done at: Wed Jan 31 02:48:38 2024
!Time: Wed Jan 31 23:24:15 2024
```

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

```
router bgp 65000
  address-family ipv4 unicast
    redistribute direct route-map ALLOW
  neighbor 10.1.1.6
    remote-as 65000
  address-family ipv4 unicast
  neighbor 10.1.1.10
    remote-as 65000
  address-family ipv4 unicast
  neighbor 172.16.1.1
    remote-as 64512
  address-family ipv4 unicast
```

R5#

Nexus R6

R6# show run bgp

```
!Command: show running-config bgp
```



!Running configuration last done at: Wed Jan 31 03:12:40 2024  
!Time: Wed Jan 31 23:24:18 2024

version 10.2(5) Bios:version 05.47  
feature bgp

```
router bgp 65000
  cluster-id 10.160.0.1
  address-family ipv4 unicast
    redistribute direct route-map ALLOW
  neighbor 10.1.1.9
    remote-as 65000
    address-family ipv4 unicast
      route-reflector-client
  neighbor 10.1.2.9
    remote-as 65000
    address-family ipv4 unicast
      route-reflector-client
  neighbor 10.1.3.2
    remote-as 65000
    address-family ipv4 unicast
      route-reflector-client
  neighbor 10.1.3.6
    remote-as 65000
    address-family ipv4 unicast
      route-reflector-client
```

R6#

## Validation

After turning Nexus R2 and Nexus R6 as Route Reflector, the prefixes are advertised to the iBGP peers. .

R2# show ip bgp neighbors 10.1.2.2 advertised-routes

Peer 10.1.2.2 routes for address family IPv4 Unicast:  
BGP table version is 22, Local Router ID is 10.150.0.1  
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
*>r10.1.1.4/30	0.0.0.0	0	100	32768	?
*>i10.1.1.8/30	10.1.1.5	0	100	0	?
*>r10.1.2.0/30	0.0.0.0	0	100	32768	?
*>r10.1.2.4/30	0.0.0.0	0	100	32768	?
*>r10.1.2.8/30	0.0.0.0	0	100	32768	?
*>i10.100.100.0/24	10.1.2.6		100	0	i
*>r10.150.0.0/30	0.0.0.0	0	100	32768	?
*>i172.16.1.0/30	10.1.1.5	0	100	0	?
*>i192.168.1.0/30	172.16.1.1	0	100	0	64512 ?
*>i192.168.2.0/30	172.16.1.1	0	100	0	64512 ?
*>i192.168.3.0/30	172.16.1.1	0	100	0	64512 ?
*>i192.168.4.0/30	172.16.1.1	0	100	0	64512 ?

```
*>i192.168.5.0/30    172.16.1.1        0        100        0 64512 ?
```

R2#

Nexus R3 and Nexus R4 receives the prefixes generated by Nexus R1.

R3# show ip bgp

BGP routing table information for VRF default, address family IPv4 Unicast

BGP table version is 108, Local Router ID is 10.100.100.1

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
* i10.1.1.4/30	10.1.2.9	0	100	0	?
*>i	10.1.2.5	0	100	0	?
*>i10.1.1.8/30	10.1.3.5	0	100	0	?
* i	10.1.2.10	0	100	0	?
* i10.1.2.0/30	10.1.2.9	0	100	0	?
*>i	10.1.2.5	0	100	0	?
* i10.1.2.4/30	10.1.2.9	0	100	0	?
*>i	10.1.2.5	0	100	0	?
* i10.1.2.8/30	10.1.3.5	0	100	0	?
*>i	10.1.2.5	0	100	0	?
* i10.1.3.0/30	10.1.2.10	0	100	0	?
*>i	10.1.3.5	0	100	0	?
* i10.1.3.4/30	10.1.2.10	0	100	0	?
*>i	10.1.3.5	0	100	0	?
*>l10.100.100.0/24	0.0.0.0		100	32768	i
* i10.150.0.0/30	10.1.2.9	0	100	0	?
*>i	10.1.2.5	0	100	0	?
* i10.160.0.0/30	10.1.2.10	0	100	0	?
*>i	10.1.3.5	0	100	0	?
* i10.200.200.0/24	10.1.3.2		100	0	i
*>i	10.1.2.2		100	0	i
* i172.16.1.0/30	10.1.1.9	0	100	0	?
*>i	10.1.1.5	0	100	0	?
* i192.168.1.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?
* i192.168.2.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?
* i192.168.3.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?
* i192.168.4.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?
* i192.168.5.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?

R3#

R4# show ip bgp

BGP routing table information for VRF default, address family IPv4 Unicast

BGP table version is 78, Local Router ID is 10.200.200.1

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
* i10.1.1.4/30	10.1.2.9	0	100	0	?
*>i	10.1.2.1	0	100	0	?
*>i10.1.1.8/30	10.1.3.1	0	100	0	?
* i	10.1.2.10	0	100	0	?
* i10.1.2.0/30	10.1.2.9	0	100	0	?
*>i	10.1.2.1	0	100	0	?
* i10.1.2.4/30	10.1.2.9	0	100	0	?
*>i	10.1.2.1	0	100	0	?
* i10.1.2.8/30	10.1.3.1	0	100	0	?
*>i	10.1.2.1	0	100	0	?
* i10.1.3.0/30	10.1.2.10	0	100	0	?
*>i	10.1.3.1	0	100	0	?
* i10.1.3.4/30	10.1.2.10	0	100	0	?
*>i	10.1.3.1	0	100	0	?
* i10.100.100.0/24	10.1.3.6		100	0	i
*>i	10.1.2.6		100	0	i
* i10.150.0.0/30	10.1.2.9	0	100	0	?
*>i	10.1.2.1	0	100	0	?
* i10.160.0.0/30	10.1.2.10	0	100	0	?
*>i	10.1.3.1	0	100	0	?
*>i10.200.200.0/24	0.0.0.0		100	32768	i
* i172.16.1.0/30	10.1.1.9	0	100	0	?
*>i	10.1.1.5	0	100	0	?
* i192.168.1.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?
* i192.168.2.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?
* i192.168.3.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?
* i192.168.4.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?
* i192.168.5.0/30	172.16.1.1	0	100	0	64512 ?
*>i	172.16.1.1	0	100	0	64512 ?

R4#

The prefixes learned by the Route Reflector display the BGP peer IP and the Cluster ID.

R4# show ip bgp 192.168.2.0

BGP routing table information for VRF default, address family IPv4 Unicast

BGP routing table entry for 192.168.2.0/30, version 63

Paths: (2 available, best #2)

Flags: (0x8000001a) (high32 00000000) on xmit-list, is in urib, is best urib route, is in HW

Path type: internal, path is valid, not best reason: Neighbor Address, no labeled nexthop

AS-Path: 64512 , path sourced external to AS

172.16.1.1 (metric 0) from 10.1.3.1 (10.160.0.1) >>>>>> Peer IP (Cluster ID)

Origin incomplete, MED 0, localpref 100, weight 0

Originator: 172.16.1.2 Cluster list: 10.160.0.1

Advertised path-id 1

Path type: internal, path is valid, is best path, no labeled nexthop, in rib

AS-Path: 64512 , path sourced external to AS

172.16.1.1 (metric 0) from 10.1.2.1 (10.150.0.1)


Origin incomplete, MED 0, localpref 100, weight 0

Originator: 172.16.1.2 Cluster list: 10.150.0.1

Path-id 1 not advertised to any peer

R4#

---


 **Note:** The routes received by Nexus R3 and Nexus R4 by the Route Reflector are not advertised to other iBGP peer as part of the Loop avoidance rule.

---

## Summary

The Route Reflector configuration does not require a complex configuration to reflect the routes to their iBGP clients, just add under the address family the keyword "route-reflector-client", this notifies the system to act as a Route Reflector for that neighbor.

---

 **Note:** To review BGP confederations in Nexus, see [Configure BGP Confederations on Nexus 9000](#)

---

## Related Information

- [Cisco Technical Support & Downloads](#)