

Understand Basic BGP Troubleshoot

Contents

[Introduction](#)

[Prerequisites](#)

[Background Information](#)

[Identify Configuration Implemented for BGP Peers](#)

[How to Understand Basic BGP Outputs](#)

[Related Information](#)

Introduction

This document describes the procedure to perform Border Gateway Protocol (BGP) troubleshooting and understanding of the basic outputs on a Nexus.

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

BGP is an exterior gateway protocol used in large-scale networks to enable the exchange of routing and reachability information between autonomous systems (ASes). It is the core routing protocol that powers the Internet global routing system.

Identify Configuration Implemented for BGP Peers

Validate the feature BGP is enabled on Nexus 9300.

Identify and understand the configuration for the BGP process.

```
switch# show running-config bgp
```

```
!Command: show running-config bgp
!Running configuration last done at: Tue Jul 18 19:45:05 2023
!Time: Tue Jul 18 19:45:44 2023
```

```
version 10.2(4) Bios:version 05.47
feature bgp
```

```
router bgp 64512                                -> The AS number of the local BGP speaker.
  router-id 172.17.255.255                       -> BGP speaker identifier.
  address-family ipv4 unicast                    -> Global address family configuration.
    network 10.100.1.0/24                         -> Specifies a network as local to this autonomous system
  redistribute direct route-map permit_all        -> Routes redistribution from other routing protocols (OS
neighbor 10.1.1.1                                -> IP address of the remote BGP peer.
  remote-as 64512                                -> The AS number of the remote BGP peer.*
  address-family ipv4 unicast                    -> Local address family configuration.
    prefix-list allow_in in                       -> Prefix-list applied at the inbound of the BGP peer. **
    prefix-list allow_out out                     -> Prefix-list applied at the outbound of the BGP peer. *
    soft-reconfiguration inbound always           -> Store the inbound BGP route updates.
neighbor 172.18.255.255
  remote-as 65535
  update-source loopback10                       -> Interface used to source BGP updates.
  ebgp-multihop 3                                -> Maximum hops to reach peer IP address, it modifies the
  address-family ipv4 unicast
    route-map block_route in                     -> Route-map applied at the inbound of the BGP peer. **
    route-map no_local out                       -> Route-map applied at the outbound of the BGP peer. **
    soft-reconfiguration inbound always
```



Note: Same ASN configured at local and remote identifies an iBGP session, a different ASN configured identifies an eBGP session.

The route map has a higher hierarchy value than a prefix list applied to the same peer.


How to Understand Basic BGP Outputs

BGP peer status

```
switch# show ip bgp summary
```

```
BGP summary information for VRF default, address family IPv4 Unicast
BGP router identifier 172.17.255.255, local AS number 64512 -> Local BGP ID and ASN
BGP table version is 67, IPv4 Unicast config peers 2, capable peers 2
20 network entries and 19 paths using 5424 bytes of memory
BGP attribute entries [6/2112], BGP AS path entries [2/20]
BGP community entries [0/0], BGP clusterlist entries [0/0]
13 received paths for inbound soft reconfiguration
12 identical, 0 modified, 1 filtered received paths using 96 bytes
```

```
Neighbor      V AS      MsgRcvd   MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
10.1.1.1      4 64512    346      334      67     0     0 05:25:12 8    -> BGP peer IP address, remote
```

 **Note:** Up/Down section displays the time the BGP peer session has been up or down.
State/PfxRcd section displays the BGP state the session is currently on. It shows the number of prefixes learned from its peer when the state is established.

BGP STATES	
Idle	This is the first state where BGP waits for a “start event”. The start event occurs when someone configures a new BGP neighbor or when we reset an established BGP peering.
Connect	BGP is waiting for the TCP three-way handshake to complete. When it is successful, it continues to the OpenSent state. In case it fails, we continue to the Active state.
Active	BGP tries another TCP three-way handshake to establish a connection with the remote BGP neighbor. If it is successful, moves to the OpenSent state. If the ConnectRetry timer expires then we move back to the Connect state.
OpenSent	In this state, BGP waits for an Open message from the remote BGP neighbor.
OpenConfirm	BGP waits for a keepalive message from the remote BGP neighbor.
Established	The BGP neighbor adjacency is complete and the BGP routers send update packets to exchange routing information.

Understand BGP peer information.

```

switch# show ip bgp neighbors 10.1.1.1
BGP neighbor is 10.1.1.1, remote AS 64512, ibgp link, Peer index 3
  BGP version 4, remote router ID 172.16.255.255
  Neighbor previous state = OpenConfirm
  BGP state = Established, up for 00:05:29
  Neighbor vrf: default
  Peer is directly attached, interface Ethernet1/49
  Last read 00:00:28, hold time = 180, keepalive interval is 60 seconds
  Last written 00:00:28, keepalive timer expiry due 00:00:31
  Received 363 messages, 0 notifications, 0 bytes in queue
  Sent 354 messages, 1 notifications, 0(0) bytes in queue
  Enhanced error processing: On
    0 discarded attributes
  Connections established 2, dropped 1
  Last update recd 00:05:28, Last update sent = 00:05:28
    Last reset by us 00:06:21, due to holdtimer expired error
  Last error length sent: 0
  Reset error value sent: 0
  Reset error sent major: 4 minor: 0
  Notification data sent:
  Last reset by peer never, due to No error
  Last error length received: 0
  Reset error value received 0
  Reset error received major: 0 minor: 0
  Notification data received:

Neighbor capabilities:
Dynamic capability: advertised (mp, refresh, gr) received (mp, refresh, gr)
Dynamic capability (old): advertised received
    
```

-> Peer IP address, remote AS, Peer index
-> BGP version, Peer Router ID
-> Previous BGP state
-> Current BGP state and up time
-> VRF used for the peer
-> Interface used to source traffic
-> Amount of time from last read/written
-> Amount of time from last received/sent
-> Counters informing the state of the peer
-> Timer of the last major update
-> Last reset timer and reason

Route refresh capability (new): advertised received
Route refresh capability (old): advertised received
4-Byte AS capability: advertised received
Address family IPv4 Unicast: advertised received
Graceful Restart capability: advertised received

Graceful Restart Parameters:

Address families advertised to peer:
 IPv4 Unicast
Address families received from peer:
 IPv4 Unicast
Forwarding state preserved by peer for:
Restart time advertised to peer: 120 seconds
Stale time for routes advertised by peer: 300 seconds
Restart time advertised by peer: 120 seconds
Extended Next Hop Encoding Capability: advertised received
Receive IPv6 next hop encoding Capability for AF:
 IPv4 Unicast VPNv4 Unicast

Message statistics:

	Sent	Rcvd
Opens:	2	2
Notifications:	1	0
Updates:	22	20
Keepalives:	340	339
Route Refresh:	1	0
Capability:	2	2
Total:	354	363
Total bytes:	7949	7524
Bytes in queue:	0	0

For address family: IPv4 Unicast
BGP table version 88, neighbor version 88
8 accepted prefixes (8 paths), consuming 2176 bytes of memory
0 received prefixes treated as withdrawn
11 sent prefixes (11 paths)
Inbound soft reconfiguration allowed(always)
Third-party Nexthop will not be computed.
Inbound ip prefix-list configured is allow_in, handle obtained
Outbound ip prefix-list configured is allow_out, handle obtained
Last End-of-RIB received 00:00:01 after session start
Last End-of-RIB sent 00:00:01 after session start
First convergence 00:00:01 after session start with 11 routes sent

-> Amount of prefixes ad

Local host: 10.1.1.2, Local port: 28262
Foreign host: 10.1.1.1, Foreign port: 179
fd = 73

Understand BGP table

This output displays the status, next hop, metric, local preference, weight, and AS-path of all prefixes learned in the BGP table.

```
switch# show ip bgp
BGP routing table information for VRF default, address family IPv4 Unicast
BGP table version is 88, Local Router ID is 172.17.255.255
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
*>r10.1.1.0/30	0.0.0.0	0	100	32768	?
*>i10.100.1.0/24	10.1.1.1		100	0	i
*>i10.100.2.0/24	10.1.1.1		100	0	i
*>i10.100.3.0/24	10.1.1.1		150	0	i
*>i10.100.4.0/24	10.1.1.1	0	100	0	?
*>i10.100.5.0/24	10.1.1.1	0	100	0	?
*>i10.100.6.0/24	10.1.1.1	0	100	0	?
*>i10.100.7.0/24	10.1.1.1	0	100	0	?
*>i10.100.8.0/24	10.1.1.1	0	100	0	?
*>r172.17.255.255/32	0.0.0.0	0	100	32768	?
*>e172.30.1.0/24	172.18.255.255	0		0	65535 ?
*>e172.30.2.0/24	172.18.255.255	0		0	65535 ?
*>e172.30.3.0/24	172.18.255.255	0		0	65535 ?
*>e172.30.4.0/24	172.18.255.255	0		0	65535 ?
*>e172.30.5.0/24	172.18.255.255	0		0	65535 65534 65533 ?
*>e172.30.6.0/24	172.18.255.255	0		0	65535 65534 65533 ?
*>e172.30.7.0/24	172.18.255.255	0		0	65535 65534 65533 ?
*>e172.30.8.0/24	172.18.255.255	0		0	65535 65534 65533 ?
*>r192.168.1.0/30	0.0.0.0	0	100	32768	?

Prefixes are advertised to a specific BGP neighbor.

```
switch# show ip bgp neighbors 172.18.255.255 advertised-routes
```

```
Peer 172.18.255.255 routes for address family IPv4 Unicast:
BGP table version is 88, Local Router ID is 172.17.255.255
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
*>i10.100.1.0/24	10.1.1.1		100	0	i
*>i10.100.2.0/24	10.1.1.1		100	0	i
*>i10.100.3.0/24	10.1.1.1		150	0	i
*>i10.100.4.0/24	10.1.1.1	0	100	0	?
*>i10.100.5.0/24	10.1.1.1	0	100	0	?
*>i10.100.6.0/24	10.1.1.1	0	100	0	?
*>i10.100.7.0/24	10.1.1.1	0	100	0	?
*>i10.100.8.0/24	10.1.1.1	0	100	0	?

Prefixes received from a BGP peer before any filters (prefix list and/or route map)*

```
switch# show ip bgp neighbors 172.18.255.255 received-routes
```

```
Peer 172.18.255.255 routes for address family IPv4 Unicast:
BGP table version is 88, Local Router ID is 172.17.255.255
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
* e172.18.255.255/32	172.18.255.255	0		0 65535 ?
*>e172.30.1.0/24	172.18.255.255	0		0 65535 ?
*>e172.30.2.0/24	172.18.255.255	0		0 65535 ?
*>e172.30.3.0/24	172.18.255.255	0		0 65535 ?
*>e172.30.4.0/24	172.18.255.255	0		0 65535 ?
*>e172.30.5.0/24	172.18.255.255	0		0 65535 65534 65533 ?
*>e172.30.6.0/24	172.18.255.255	0		0 65535 65534 65533 ?
*>e172.30.7.0/24	172.18.255.255	0		0 65535 65534 65533 ?
*>e172.30.8.0/24	172.18.255.255	0		0 65535 65534 65533 ?

 **Note** Soft-reconfiguration inbound has to be configured on the neighbor

Prefixes received from a BGP peer after filters (prefix-list and/or route-map)

```
switch# show ip bgp neighbors 172.18.255.255 routes
```

```
Peer 172.18.255.255 routes for address family IPv4 Unicast:
BGP table version is 88, Local Router ID is 172.17.255.255
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
*>e172.30.1.0/24	172.18.255.255	0		0 65535 ?
*>e172.30.2.0/24	172.18.255.255	0		0 65535 ?
*>e172.30.3.0/24	172.18.255.255	0		0 65535 ?
*>e172.30.4.0/24	172.18.255.255	0		0 65535 ?
*>e172.30.5.0/24	172.18.255.255	0		0 65535 65534 65533 ?
*>e172.30.6.0/24	172.18.255.255	0		0 65535 65534 65533 ?
*>e172.30.7.0/24	172.18.255.255	0		0 65535 65534 65533 ?
*>e172.30.8.0/24	172.18.255.255	0		0 65535 65534 65533 ?

Detail path information for a specific prefix.

```
switch# show ip bgp 172.30.6.0
BGP routing table information for VRF default, address family IPv4 Unicast
BGP routing table entry for 172.30.6.0/24, version 28
Paths: (3 available, best #3)
Flags: (0x8000001a) (high32 00000000) on xmit-list, is in urib, is best urib route, is in HW
```

-> Prefi
-> Numbe

```
Path type: external, path is valid, not best reason: Router Id, no labeled nexthop
AS-Path: 65535 65534 65533 , path sourced external to AS
172.20.255.255 (metric 0) from 172.20.255.255 (172.20.255.255)
Origin incomplete, MED 0, localpref 100, weight 0
```

-> As Pa
-> Next

```
Path type: external, path is valid, not best reason: newer EBGp path, no labeled nexthop
AS-Path: 65535 65534 65533 , path sourced external to AS
172.19.255.255 (metric 0) from 172.19.255.255 (172.19.255.255)
```

Origin incomplete, MED 0, localpref 100, weight 0

Advertised path-id 1

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

-> Path s

AS-Path: 65535 65534 65533 , path sourced external to AS


172.18.255.255 (metric 0) from 172.18.255.255 (172.18.255.255)

Origin incomplete, MED 0, localpref 100, weight 0

Path-id 1 advertised to peers:

-> BGP pee

10.1.1.2

 **Note:** *AS-Path represents the ASN traversed to reach the location where the prefix originated.
**The AS-Path reads from right to left.

To review the Best Path Selection Process in BGP, see [BGP Best Path Selection](#).

Related Information

- [Cisco Technical Support & Downloads](#)