

BGP Basics

- *ASN is 2 bytes, private pool AS numbers, 64512 to 65535, Private AS cannot be leaked to the internet.*
- *BGP is path vector protocol and sits on top of TCP.* BGP updates are carried using TCP on port 179.
- *AS-path itself is an anti-loop mechanism. BGP is distance vector protocol ensuring path is loop free.*
- *BGP doesn’t have any transport mechanism and it depends on IGP for transport capability.*
- *RR and Confederation is used to get rid of IBGP full mesh topology and hence reduces IBGP sessions.*
- *Confederation we divide the local AS into sub-AS which gives more control but makes network complex.*
- *no advertise = no to IBGP peers/neighbours & no to EBGP peers/neighbours*
- *no export = no to EBGP peers/neighbours*
- Incoming route advertisements influence your outgoing traffic, and outgoing advertisements influence your incoming traffic.
- *BGP has a mechanism to "hold down" routes that appear to be flaky. Dampening will prevent BGP peers from listening to all routing updates from flapping peers.*
- *MP-BGP capability can be obtained by adding an extension called address family identifier AFI.*
- *BGP is control plane routing protocol because it allows for exchange of routes with a peer that is multiple hops away.*
- *Sessions established with an iBGP router that are in the same AS.*
- *Sessions established with a BGP router that are in different AS.*
- *TTL on eBGP packets is set to 1 by default. TTL on iBGP is set to 255, which allows multi-hop sessions*
- *AS = a logical domain under the control of a single entity*

BGP Packets Types & Notes

- **open:** *set up and establishes BGP adjacency.*
 - **update:** *advertises, updates, or withdraws routes.*
 - **notification:** *indicates an error condition to a BGP neighbour.*
 - **Keep alive:** *Ensures that BGP neighbours are still alive.*
- Redistributing routes from IGP into BGP is completely safe, however redistributing routes learnt from BGP should be done with caution.
- BGP route summarization is required to reduce the size of BGP table for internet routers. Dynamic route summarization is obtained using the command line.
- BGP communities provide capability for tagging routes and for modifying BGP routing policy on upstream and downstream routers. BGP communities are an optional transitive attribute that can traverse from AS to AS.
- The longest prefix length is always preferred, /28 is preferred more than /26.
- Multipath = Load balance the traffic between the EBGp peers
- Multi hop = indicates the number of hops away from the peer router

BGP Attributes Importance

Attributes	EBGP	IBGP
AS_Path	Well Known Mandatory	Well know Mandatory
Next_Hop	Well know Mandatory	Well know Mandatory
Origin	Well know Mandatory	Well know Mandatory
LOCAL_Pref	Not allowed	Well known Discretionary
Aggregator	Optional Transitive	Optional Transitive
Atomic_Agg	Well Known Discretionary	Well known Discretionary
Community	Optional Transitive	Optional Transitive
MED	Optional Non Transitive	Optional Non Transitive

BGP commands reference Cisco IOS-XE

- show ip bgp summary
- show ip bgp
- show ip bgp regexp ^65536\$
- clear ip bgp *
- show ip bgp 10.100.0.0
- show ip bgp vpnv4 vrf test1 192.168.13.1
- show ip bgp neighbors
- show ip community-list
- show ip bgp paths
- show ip bgp template peer-session
- clear bgp nsap {* | as-number | ip-address} [soft] [in | out]
- clear bgp nsap dampening [nsap-prefix]
- show bgp all community
- show ip bgp all dampening
- show ip bgp flap-statistics
- show ip bgp ipv4 unicast summary
- show ip bgp ipv4 multicast summary
- show ip bgp extcommuinty-list
- show ip bgp cidr-only
- show ip bgp all dampening dampened-paths
- show ip bgp dampening parameters
- show ip bgp l2vpn vpls all
- show ip bgp l2vpn vpls rd 45000:100
- show ip bgp neighbors 10.108.50.2
- show ip bgp neighbors 172.16.232.178 advertised-routes
- show ip bgp neighbors 192.168.20.72 received prefix-filter
- show ip bgp neighbors 192.168.1.2 policy
- show ip bgp peer-group internal
- show ip bgp replication
- show ip bgp rib-failure

BGP commands reference Cisco IOS-XR

- clear bgp dampening
- clear bgp
- clear bgp soft
- clear bgp performance-statistics
- clear bgp current-mode
- clear bgp shutdown
- show bgp all all summary
- show bgp advertised
- show bgp cidr-only
- show bgp neighbors
- show bgp sessions
- show bgp summary
- show bgp paths
- show bgp table
- show bgp vrf imported-routes
- show bgp nexthops
- show bgp flap-statistics
- show bgp community
- show bgp labels
- show bgp neighbor-group
- show bgp regexp
- show bgp process
- show bgp truncated-communities
- show bgp vpnv4 unicast rd 9500:1570008101 0.0.0.0
- show run router bgp 7777 vrf INT_MGMT
- show bgp vrf TEST ipv4 unicast advertised
- show bgp vrf TEST neighbors 203.118.143.9 advertised-routes
- show bgp vrf TEST all summary
- show bgp vpnv4 unicast vrf TEST neighbors 172.27.249.1 received-routes
- show bgp vpnv4 unicast vrf TEST neighbors 172.27.249.1 advertised-routes

BGP command reference Huawei

- display bgp peer
- display bgp vpnv4 all peer
- display bgp peer statistics
- display bgp routing-table
- display bgp routing-table flap-info
- display bgp vpnv4 brief
- display bgp vpnv4 routing-table
- display bgp routing-table dampened
- display bgp group
- display bgp resource
- display bgp ipv6 routing-table statistics
- display bgp link-state unicast peer
- display bgp network
- display bgp vpnv4 all routing-table peer 10.168.32.24 received-routes
- display bgp vpnv4 all routing-table peer 10.168.32.24 advertised-routes
- display bgp routing-table peer 10.168.0.24 received-routes
- display bgp routing-table peer 10.168.0.24 advertised-routes
- display bgp error
- display bgp error discard
- display bgp mdt routing-table
- display bgp paths
- display bgp network
- display bgp multicast peer
- display bgp routing-table label
- reset bgp
- reset bgp dampening
- reset bgp flap-info
- reset bgp flapping-count
- reset bgp multicast dampening
- reset bgp multicast external

BGP command reference Juniper

- show bgp summary
- show bgp summary rib-sharding main
- show bgp summary group Group2
- show route receive-protocol bgp 10.4.12.11 detail
- show route advertising-protocol bgp neighbor-address X.X.X.X
- show route forwarding-table
- show route protocol bgp
- show route table bgp.l3vpn-inet6.0
- show route table TV-Ingest.inet.4
- show route table TV-CDN.inet.0
- show bgp summary [no-more
- show bgp summary logical-system TEST | match Establ | no-more
- show bgp summary logical-system TEST | match Establ | count
- show log bgp
- clear bgp neighbor
- show bgp neighbor 10.0.3.X
- show route protocol bgp 192.168.70.1 detail
- show route protocol bgp 192/8 | match 192.168 | count
- show route protocol bgp 192.168.60/24
- show route advertising-protocol bgp 10.0.3.3 192.168.100/24 detail
- show protocols bgp
- show protocols bgp group 65000
- show route damping suppressed
- show route damping suppressed 62.109.0.0/19 detail
- show route source-gateway 172.16.0.14 detail
- show route resolution unresolved detail
- show route protocol aggregate detail

BGP Attributes

- **Weight** is 16 bit value (0 to 65535) assigned locally on the router. Used when 2 ISP’s are connected to the same router. The path with the higher weight is preferred. Applied on inbound and controls outgoing traffic from AS.
- **LP** is 32 bit value (0 to 4,294,967,295), used when 2 ISP’s are connected to different routers, applied on inbound direction while controls outgoing traffic. The route which has higher value LP is preferred.
- **Locally originated** routes is given preference, networks that have been aggregated locally and routes received by BGP peers.
- **A shorter AS path** is preferred over a longer AS path. The attribute is applied on out policy and affects incoming traffic.
- **Origin** IGP = I, EGP and then incomplete origin which is redistributed networks.
- **MED** 32 bit value, a lower MED is preferred over a higher MED. The attribute is applied on out policy and affects incoming traffic.

BGP neighbour ship states

- **Idle:** *BGP detects s start event and tries to initiate a TCP connection to the BGP peer and also listens for a new connections from the peer router.*
- **Connect:** *BGP initiates the TCP connection. If the TCP three way handshake is completed, the established BGP session process resets the ConnectRetryTimer and sends the open message to the neighbour, it then changes to OpenSent state.*
- **Active:** *BGP starts a new three way handshake, If a connection is established, an open message is sent.*
- **OpenSent:** *An open message has been sent from the originating router and is awaiting an Open message from the other router, when the other router sends Open message then following errors are checked. BGP version, the source IP address with neighbour IP address, ASN, BGP identifier must be unique (RIDs), Security parameters TTL must be set. If there is error, a notification has been sent out and the state move back to Idle.*
- **OpenConfirm:** *BGP awaits for a Keep alive or notification message. Upon receipt from neighbor’s keep alive message, the state moves to established. If the hold timer expires, a stop event occurs , the state is moved to Idle.*
- **Established:** *BGP neighbours exchange routing using update messages.*

BGP Route Selection Process

Attribute	Description	Preference
Weight	Administrative Weight	Highest
Local Preference	Communicated between peers within an AS	Highest
Self-Originated	Prefer path originated locally	next hop 0.0.0.0
AS Path	AS hop required to reach destination	Shortes
Origin Code	IGP (i) < EGP (e) < Redistributed / Incomplete (?)	Lowest (IGP)
MED	Used by externally peers to enter an AS	Lowest
BGP AD	Prefer eBGP routes over iBGP	eBG
IGP Cost	IGP metric to reach next hop or neighbor	Lowest
eBGP Peering	Compared the age of the route	Oldest
Router ID	Compares router ID of neighbor	Lowest
Neighbor IP	Compares IP address of the neighbor	Lowest