



Implementing Basic BGP Routing

Implement BGP in the Service Provider Network



Objectives


- Describe planning for BGP deployments
- Describe the basic configuration steps for EBGp
- Describe how networks are advertised in BGP networks
- Describe the basic configuration steps for IBGP
- Describe full-mesh IBGP networks
- Describe BG support for IPv6
- Describe how to shut down a BGP neighbor
- Describe next-hop behavior in BGP
- Describe the BGP next hop self feature
- Describe configuration template support in Cisco IOS XR Software
- Describe the BGP neighbor states
- Describe BGP neighbor authentication
- Describe clearing the BGP session

Objectives (Cont.)


- Describe how to monitor BGP routes
- Describe the importance of the BGP path attributes in the path selection
- Describe the BGP weight attribute
- Describe the BGP local preference attribute
- Describe the BGP AS path attribute
- Describe the BGP multi-exit discriminator attribute

Planning for BGP

- Define network requirements.
- Define internal connectivity.
- Define external connectivity to service provider.
- Gather required parameters.



Configure Basic EBG



RP/0/RSP0/CPU0:PE1#show bgp summary

BGP router identifier 10.1.1.1, local AS number 64500


< text omitted >

Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
Speaker	5	5	5	5	5	5


Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
192.168.101.11	0	64501	4302	3909	5	0	0	2d16h	0

Advertise Networks

- Two options
 - Configure the local networks to be advertised, and include them in BGP.
 - Use redistribution from IGP to BGP.



Configure Basic IBGP



```
RP/0/RSP0/CPU0:PE1#show bgp summary
BGP router identifier 10.1.1.1, local AS number 64500
< text omitted >
```


Process	RcvTblVer	bRIB/RIB	LabelVer	ImportVer	SendTblVer	StandbyVer
Speaker	7	7	7	7	7	7

Neighbor	Spk	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	St/PfxRcd
10.0.1.1	0	64500	2	2	0	0	0	00:00:04	0
192.168.101.11	0	64501	4377	3977	7	0	0	2d17h	1

Full-Mesh IBGP


```
router bgp 64501
neighbor 10.0.1.2
remote-as 64501
update-source Loopback0
address-family ipv4 unicast
neighbor 10.0.1.3
remote-as 64501
update-source Loopback0
address-family ipv4 unicast
```

```
router bgp 64501
neighbor 10.0.1.1
remote-as 64501
update-source Loopback0
address-family ipv4 unicast
neighbor 10.0.1.3
remote-as 64501
update-source Loopback0
address-family ipv4 unicast
```




```
router bgp 64501
neighbor 10.0.1.1 remote-as 64501
neighbor 10.0.1.1 update-source Loopback0
neighbor 10.0.1.2 remote-as 64501
neighbor 10.0.1.2 update-source Loopback0
```


BGP Support for IPv6




Shutting Down a BGP Neighbor



Next-Hop Behavior



BGP Next Hop Self



Configuration Templates

Cisco IOS XR

The **af-group** is used to group address family-specific neighbor commands within an IPv4, IPv6, or VPNv4, address family.

```
router bgp 1
af-group afmcast1 address-family ipv4
multicast
```

Define address family group.

The **session-group** allows you to create a session group from which neighbors can inherit address family-independent configuration.

```
router bgp 1
session-group session1
```

Define session group.

The **neighbor-group** helps you apply the same configuration to one or more neighbors.

```
router bgp 1
neighbor-group nbrgroup1
!
router bgp 1
neighbor-group nbrgroup1
address-family ipv4 unicast
```

Define neighbor group.

Use neighbor group.

Cisco IOS/IOS XE

The **BGP peer group** groups BGP neighbors who share the same policies.

```
router bgp 1
neighbor peer-group-name peer-group
```

Define peer group.

BGP States

When establishing a BGP session, BGP goes through the following states:

1. **Idle**: The router is searching the routing table to see whether a route exists to reach the neighbor.
2. **Connect**: The router found a route to the neighbor and has completed the three-way TCP handshake.
3. **Open sent**: The open message is sent, with the parameters for the BGP session.
4. **Open confirm**: The router received an agreement on the parameters for establishing a session.
 - Alternatively, the router goes into the **active** state if there is no response to the open message.
5. **Established**: Peering is established; routing begins.

BGP Established and Idle States

- **Idle:** The router cannot find the address of the neighbor in the routing table.
 - **Solution:** Check for an IGP problem. Is the neighbor announcing the route?
- **Established:** This is the proper state for BGP operations

```
RP/0/RSP0/CPU0:PE1#show bgp summary
BGP router identifier 10.1.1.1, local AS number 64500
< text omitted >
Process          RcvTblVer    bRIB/RIB    LabelVer    ImportVer    SendTblVer    Standby
Speaker          7            7           7           7            7            7

Neighbor        Spk     AS MsgRcvd MsgSent   TblVer   InQ  OutQ  Up/Down  St/PfxRcd
10.0.1.1         0 64500   355    343      36      0    0 00:16:28      4
192.168.101.11   0 64501  4444   4051      36      0    0 00:16:18      1
```

A number in the state column indicating the number of routes learned from this neighbor.

```
RP/0/RSP0/CPU0:PE1#show bgp neighbor
< text omitted >
BGP neighbor is 192.168.101.11
Remote AS 64501, local AS 64500, external link
Remote router ID 10.1.10.1
BGP state = Established, up for 00:16:18
Last read 00:00:49, Last read before reset 00:16:52
Hold time is 180, keepalive interval is 60 seconds
Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
< text omitted >
```

BGP Active State


Active: The router has sent an open packet and is waiting for a response.

- The state may cycle between active and idle.
- The neighbor may not know how to get back to this router due to the following reasons:
 - There is no route to the source IP address of the BGP open packet.
 - The neighbor is peering with the wrong address.
 - There is no neighbor statement for this router.
 - The AS number is misconfigured.

Active does not mean that BGP peering is established.

```
RP/0/RSP0/CPU0:PE1#show bgp summary
BGP router identifier 10.1.1.1, local AS number 64500
< text omitted >
Neighbor      Spk      AS  MsgRcvd  MsgSent   TblVer   InQ  OutQ  Up/Down  St/PfxRcd
10.0.1.1      0 64500      0     343       0      0    0  00:00:00  Active
```

BGP Neighbor Authentication



- BGP authentication uses MD5.
- Configure a key—password; the router generates a message digest (is sent), or hash, of the key (is not sent) and the message.
- The router generates and checks the MD5 digest of every segment that is sent on the TCP connection.
- The router authenticates the source of each routing update packet that it receives.

Clearing the BGP Session

- When policies change, the change takes effect immediately.
- The next time that a prefix or path is advertised or received, the new policy is used. This can take a long time for all networks.
- You must trigger an update for immediate action.

Hard reset

```
clear bgp *
```

Soft reset

```
clear bgp ipv4 unicast 192.168.101.11 soft out
```

```
clear bgp ipv4 unicast 192.168.101.11 soft in
```

Soft reconfiguration feature or
ROUTE_REFRESH capability
is not required.

Soft reconfiguration feature or
ROUTE_REFRESH capability
is required.




EBGP




```
RP/0/RSP0/CPU0:PE1#show bgp neighbor 192.168.101.11 | include Route refresh
Route refresh: advertised and received
Route refresh request: received 2, sent 1
```

Monitoring BGP Routes



BGP Path Selection

- The BGP table can have several paths for each network to choose from.
- BGP is not designed to perform load balancing:
 - Paths are chosen because of policy.
 - Paths are not chosen based upon bandwidth.
- The BGP selection process eliminates any multiple paths until a single best path remains.




Route Selection Decision Process

Consider only (synchronized) routes with no AS loops and a valid next hop. The next steps in the evaluation process are:

1.	Prefer the highest weight (local to router).
2.	Prefer the highest local preference (global within AS).
3.	Prefer the route originated by the local router (next hop = 0.0.0.0).
4.	Prefer the shortest AS path.
5.	Prefer the lowest origin code (IGP < EGP < incomplete).
6.	Prefer the lowest MED (exchanged between autonomous systems).
7.	Prefer the EBGP path over the IBGP path.
8.	Prefer the path through the closest IGP neighbor.
9.	Prefer the oldest route for EBGP paths.
10.	Prefer the path with the lowest neighbor BGP router ID.
11.	Prefer the path with the lowest neighbor IP address.


BGP Weight

- Weight is an attribute that is proprietary to Cisco.
- Weight is not sent to any BGP neighbors.
- It is local to the router only.
- Paths with the highest weight value are preferred.



BGP Local Preference

- Used to select the outbound EBGP path
- Sent to IBGP neighbors only (and only within the AS)
- Stripped in the outgoing EBGP updates except in the EBGP updates with confederation peers
- Local preference attribute is **well known** and **discretionary**
- Default value = 100
- Paths with highest local preference value are preferred




BGP AS Path

- Fourth BGP path selection criteria
- Prefer shorter AS paths (only length is compared)
- Influences the inbound path selection in a multihomed AS
- Manual manipulation of AS path length—AS path prepending
- AS path prepending specified per neighbor by complex criteria

BGP Multi-Exit Discriminator

- The paths with the lowest MED (also called the metric) value are the most desirable.
- MED is used to advertise an exit path to be used by EBGW neighbors to reach networks owned by this AS.
- The MED attribute is **optional** and **nontransitive**.



Summary

- For a BGP configuration, the following must be defined: BGP requirements, BGP parameters, and connectivity.
- Basic EBGP configuration requires three main steps: define the BGP process, establish the neighbor relationship, and advertise the networks into BGP.
- Networks can be advertised into BGP using the network command or by redistribution.
- It is recommended to use Loopback interfaces when establishing IBGP sessions.
- Full-mesh IBGP is required because of the split horizon rule.
- To exchange IPv6 networks over BGP, you have to activate a neighbor for IPv6 address family.
- You can use the shutdown command to manually shut down a BGP neighbor.

Summary (Cont.)

- By default, next-hop over EBGP session is IP address of a router that is sending the update. Next-hop over IBGP session is as advertised by EBGP, and should not change.
- You can use the next-hop-self command to change the default next-hop behavior.
- You can use configuration templates to group configuration that can be applied to several BGP neighbors.
- When establishing a BGP session, the BGP goes through the following states: idle, connect, open sent, open confirm, and established.
- BGP supports MD5 authentication to authenticate each received routing packet.
- You should trigger a BGP update by resetting a BGP session when you change routing policy.

Summary (Cont.)

- Use various show commands to verify BGP operations.
- After BGP receives updates about multiple destinations from different autonomous systems, it follows a multiple-step process for selecting the best route to reach a destination; the best route is a candidate for the routing table.
- Weight is Cisco's proprietary attribute that is configured locally on a router and is not propagated to any other routers.
- Local preference is a well-known discretionary attribute that provides information to routers in the AS about the path that is preferred for exiting the AS.
- AS path length is the fourth path selection criteria. The shortest AS path is preferred.
- MED is an indication to EBGP neighbors about the preferred path into an AS.

