

BGP Foundations Observation Lab – eBGP Directly Connected

Objective

The purpose of this lab is to observe how routers establish directly connected eBGP sessions and how routes propagate between autonomous systems. This is an observation-focused lab. The learner should focus on **what changes in the routing and BGP tables** as sessions are formed. No troubleshooting is required in this exercise

Topology Overview

Three routers exist in separate autonomous systems. Each router will form directly connected eBGP sessions with neighboring routers.




Figure 1 – Lab Topology

Task 1 - Establish eBGP Session Between RTR1 and RTR2

Configure an eBGP session between RTR1 and RTR2 using the directly connected interfaces.

RTR1 Configuration

```
router bgp 65001
 neighbor 12.1.1.2 remote-as 65002
```

RTR2 Configuration

```
router bgp 65002
 neighbor 12.1.1.1 remote-as 65001
```

Expected Behavior

- The session should transition through Idle → Connect → OpenSent → OpenConfirm → Established.
- Each router should recognize the neighbor as external (eBGP).

Verification

```
show bgp ipv4 unicast summary
```

```

BGP router identifier 12.1.1.1, local AS number 65001
BGP table version is 2, main routing table version 2
1 network entries using 144 bytes of memory
1 path entries using 84 bytes of memory
1/1 BGP path/bestpath attribute entries using 160 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 388 total bytes of memory
BGP activity 1/0 prefixes, 1/0 paths, scan interval 60 secs

Neighbor      V      AS MsgRcvd MsgSent  TblVer  InQ  OutQ Up/Down  State/PfxRcd
12.1.1.2      4      65002    19     19       2     0     0 00:13:47      0
RTR1#

```

Figure 2 – RTR1 Show BGP ipv4 Unicast Summary Output

Observation Questions

- Which value in the output indicates this is an eBGP neighbor?
- How would this output differ if the neighbor were iBGP?
- Which fields in the output confirm the session is Established?

Task 2 – Advertise RTR1’s Loopback into BGP

Inject RTR1’s Loopback0 network into BGP using a network statement.

```
router bgp 65001
 network 1.1.1.0 mask 255.255.255.0
```

Expected Behavior

- RTR1 should originate the prefix.
- RTR2 should learn the route via eBGP.
- The AS_PATH should contain RTR1’s ASN.

Verification

```
show ip bgp
```

```
RTR1#show ip bgp
BGP table version is 2, local router ID is 12.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
               t secondary path,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network        Next Hop           Metric LocPrf Weight Path
*>  1.1.1.0/24     0.0.0.0                0         32768 i
RTR1#
```

Figure 3 – RTR1 Show IP BGP Output

In *Figure 4*, it reveals the prefix with RTR1's ASN. This is because external peers append their ASN as the route is advertised.

```
RTR2#show ip bgp
BGP table version is 2, local router ID is 23.1.1.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
               t secondary path,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network        Next Hop           Metric LocPrf Weight Path
  *> 1.1.1.0/24    12.1.1.1                0             0 65001 i
RTR2#
```

Figure 4 – RTR2 Show IP BGP Output

Engineer Insight

The AS_PATH is one of several attributes tied to a prefix. It serves as a method for determining the best path between multiple for the same prefix as well as a loop prevention mechanism. It is appended to a prefix as it is advertised to external peers.

Observation Questions

- How does RTR2 identify the route as external?
- What attribute changed compared to locally originated routes?
- Why was the AS_PATH added.

Task 3 – Establish eBGP Session Between RTR2 and RTR3

Configure an eBGP session between RTR2 and RTR3.

Expected Behavior

- RTR3 should learn RTR1's prefix through RTR2.
- The AS_PATH should now contain two ASNs.

```
RTR3#show ip bgp summary
BGP router identifier 23.1.1.3, local AS number 65003
BGP table version is 2, main routing table version 2
1 network entries using 144 bytes of memory
1 path entries using 84 bytes of memory
1/1 BGP path/bestpath attribute entries using 160 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 412 total bytes of memory
BGP activity 1/0 prefixes, 1/0 paths, scan interval 60 secs

Neighbor      V      AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
23.1.1.2      4      65002    12     10       2     0     0 00:05:46      1
RTR3#
```

Figure 5 – RTR3 Show IP BGP Summary Output

```
RTR3#show ip bgp
BGP table version is 2, local router ID is 23.1.1.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
               t secondary path,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network        Next Hop           Metric LocPrf Weight Path
  *> 1.1.1.0/24    23.1.1.2              0  65002 65001 i
RTR3#
```

Figure 6 – RTR3 Show IP BGP Output

Observation Questions

- How did the AS_PATH change?
- Why does the AS number grow as the route moves between ASNs?

Summary

This lab demonstrated:

- Directly connected eBGP peering between different ASNs.
 - Route advertisement between autonomous systems.
 - Growth of the AS_PATH attribute.
-