BGP Operations and Security Lab

You will learn some of the BGP best practices to be implemented while peering with Transit, Customer and with in the Core Network.

- In this lab, each participant will be issued a pair of routers i.e. Border Router – BR and Core Router – CR.
- Border Router has a connectivity with Transit Router – TR, Customer Router – Cust and with Core Router – CR.
- The Border Router – BR and Core Router – CR are pre-configured.
- Connectivity between the BR – CR, BR – TR and BR – Cust is already established and eBGP/iBGP sessions are also Up.

Topology

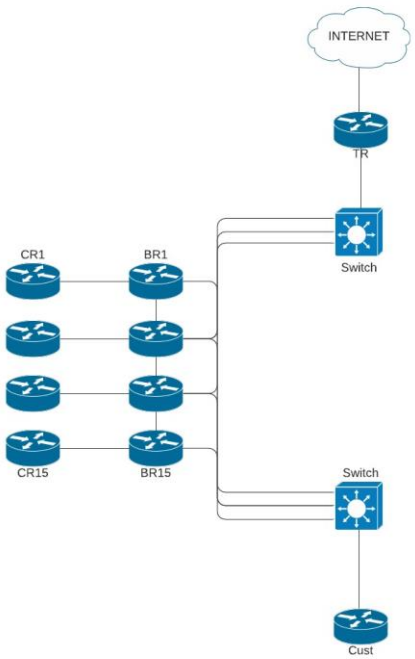

Router Assignment:

Group	Border Router - BR	Port No	Core Router - CR	Port No	AS No
1	BR1	2001	CR1	2021	1
2	BR2	2002	CR2	2022	2
3	BR3	2003	CR3	2023	3
4	BR4	2004	CR4	2024	4
5	BR5	2005	CR5	2025	5
6	BR6	2006	CR6	2026	6
7	BR7	2007	CR7	2027	7
8	BR8	2008	CR8	2028	8
9	BR9	2009	CR9	2029	9
10	BR10	2010	CR10	2030	10
11	BR11	2011	CR11	2031	11
12	BR12	2012	CR12	2032	12
13	BR13	2013	CR13	2033	13
14	BR14	2014	CR14	2034	14
15	BR15	2015	CR15	2035	15

Address Planning:

BR - TR

Group	BR	IP Address	TR	IP Address
1	BR1	192.168.1.1		
2	BR2	192.168.1.2		
3	BR3	192.168.1.3		
4	BR4	192.168.1.4		
5	BR5	192.168.1.5		
6	BR6	192.168.1.6		
7	BR7	192.168.1.7		
8	BR8	192.168.1.8	TR	192.168.1.31
9	BR9	192.168.1.9		
10	BR10	192.168.1.10		
11	BR11	192.168.1.11		
12	BR12	192.168.1.12		
13	BR13	192.168.1.13		
14	BR14	192.168.1.14		
15	BR15	192.168.1.15		

BR - Cust

Group	BR	IP Address	Cust	IP Address
1	BR1	192.168.3.1		
2	BR2	192.168.3.2		
3	BR3	192.168.3.3		
4	BR4	192.168.3.4		
5	BR5	192.168.3.5		
6	BR6	192.168.3.6		
7	BR7	192.168.3.7		
8	BR8	192.168.3.8	Cust	192.168.3.16
9	BR9	192.168.3.9		
10	BR10	192.168.3.10		
11	BR11	192.168.3.11		
12	BR12	192.168.3.12		
13	BR13	192.168.3.13		
14	BR14	192.168.3.14		
15	BR15	192.168.3.15		

BR – CR

Group	BR	IP Address	CR	IP Address
1	BR1	172.18.1.1	CR1	172.18.1.2
2	BR2	172.18.2.1	CR2	172.18.2.2
3	BR3	172.18.3.1	CR3	172.18.3.2
4	BR4	172.18.4.1	CR4	172.18.4.2
5	BR5	172.18.5.1	CR5	172.18.5.2
6	BR6	172.18.6.1	CR6	172.18.6.2
7	BR7	172.18.7.1	CR7	172.18.7.2
8	BR8	172.18.8.1	CR8	172.18.8.2
9	BR9	172.18.9.1	CR9	172.18.9.2
10	BR10	172.18.10.1	CR10	172.18.10.2
11	BR11	172.18.11.1	CR11	172.18.11.2
12	BR12	172.18.12.1	CR12	172.18.12.2
13	BR13	172.18.13.1	CR13	172.18.13.2
14	BR14	172.18.14.1	CR14	172.18.14.2
15	BR15	172.18.15.1	CR15	172.18.15.2

BR – BR (Peering)

- All the BRs are back to back connected via switch and a part of same broadcast domain, likely how IX – Internet Exchange Points interconnected via **192.168.2.0/24**.

ISP Prefixes

Group	ISP Prefixes	IP Address
1	BR1	10.0.1.0/24
2	BR2	10.0.2.0/24
3	BR3	10.0.3.0/24
4	BR4	10.0.4.0/24
5	BR5	10.0.5.0/24
6	BR6	10.0.6.0/24
7	BR7	10.0.7.0/24
8	BR8	10.0.8.0/24
9	BR9	10.0.9.0/24
10	BR10	10.0.10.0/24
11	BR11	10.0.11.0/24
12	BR12	10.0.12.0/24
13	BR13	10.0.13.0/24
14	BR14	10.0.14.0/24
15	BR15	10.0.15.0/24

Customer Prefixes

Group	Customer Prefix	Customer Loopback
1	172.16.0.0/19	172.16.0.1
2	172.16.32.0 /19	172.16.32.1
3	172.16.64.0 /19	172.16.64.1
4	172.16.96.0 /19	172.16.96.1
5	172.16.128.0	172.16.128.1
6	172.16.160.0/19	172.16.160.1
7	172.16.192.0/19	172.16.192.1
8	172.16.224.0/19	172.16.224.0
9	172.17.0.0 2/19	172.17.0.1
10	172.17.32.0 /19	172.17.32.1
11	172.17.64.0 /19	172.17.64.1
12	172.17.96.0 /19	172.17.96.1
13	172.17.128.0/19	172.17.128.1
14	172.17.160.0/19	172.17.160.1
15	172.17.192.0/19	172.17.192.1

Note: Customer is already advertising assigned Prefixes to every Border Router.

Access the Lab

- The Routers are accessible via Jump Host.
- Jump Host details and assignment will be shared by Instructor.
- From the Jump Host you need to telnet your assigned Routers.

  telnet 192.168.30.254 20XX

  where XX = Port Number

Lab Tasks

1. Protecting BGP Sessions
2. Prefix filtering Inbound/Outbound – Transit Network
3. Prefix filtering Inbound/Outbound – Customer
4. BGP next-hop-self
5. Protecting BGP Sessions
6. Max Prefix Limit

Note: Do not copy and paste the configuration, this is just the template for reference only whereas the actual configuration may vary depending on the which Router you are going to apply.

Task 1: Protecting BGP Sessions

- MD5 authentication (RFC2385) – To protect the BGP TCP session between peers
- Login to your BR and set the password <**bgpops**> against the neighbour - TR

```
BR1(config-router)#neighbor 192.168.1.31 password bgpops
```

Verification:

```
show ip bgp summary

192.168.1.31    4       100     26     23     319    0    0 00:00:06      51
```

Note: Password is already set on the TR, BGP so until the password authentication

Task 2: Prefix filtering Inbound/Outbound – Transit

Inbound:

- Do not accept Bogons.
- Do not accept your own Prefixes
- Legitimate traffic to be allowed le /24.
- Apply prefix filter to your BGP neighbour.

Create the Prefix List:

BR1(config)ip prefix-list transit-in seq 10 deny 10.0.0.0/8 le 24

BR1(config)ip prefix-list transit-in seq 20 deny 100.64.0.0/14 le 24

BR1(config)ip prefix-list transit-in seq 30 deny 127.0.0.0/24 le 32

BR1(config)ip prefix-list transit-in seq 40 deny 172.16.0.0/12 le 24

BR1(config)ip prefix-list transit-in seq 50 deny 192.168.0.0/16 le 24

BR1(config)ip prefix-list transit-in seq 60 deny 10.1.0.0/24 le 32

BR1(config)ip prefix-list transit-in seq 70 permit 0.0.0.0/0 le 24

!

Apply the Prefix List:

BR1(config-router)#address-family ipv4 unicast

BR1(config-router-af)#neighbor 192.168.1.31 prefix-list transit-in in

Verification:

Check if the prefix-list is applied to the BGP neighbour

BR1#show running-config | section bgp

neighbor 192.168.1.31 prefix-list transit-in in

Outbound:

- Your and Your Customer Prefixes to be advertised to your Transit.
- Only prefixes advertised from you and your Customer AS would be allowed.

Create the Prefix List:

BR1(config)ip prefix-list transit-out seq 10 permit 10.0.1.0/24

BR1(config)ip prefix-list transit-out seq 20 permit 172.16.0.0/19 le 24

BR1(config)ip prefix-list transit-out seq 30 deny 0.0.0.0/0 le 32

Create the AS Path filter allow only your and prefixes originated by Customer AS

BR1(config)ip as-path access-list 10 permit ^(1_)+([0-9]+)

Apply the Prefix List:

BR1(config-router)#address-family ipv4 unicast

BR1(config-router-af)#neighbor 192.168.1.31 prefix-list transit-out out

Verification:

Check if the prefix-list is applied to the BGP neighbour

BR1#show running-config | section bgp

neighbor 192.168.1.31 prefix-list transit-out out

Task 3: Prefix filtering Inbound/Outbound – Customer

Inbound:

- Only Customer Prefixes to be allowed le /24.
- Deny everything else from Customer.

Create the Prefix List:

BR2(config)#ip prefix-list cust-in seq 10 permit 172.16.32.0/19 le 24

BR2(config)#ip prefix-list cust-in seq 20 deny 0.0.0.0/0 le 32

Apply the Prefix List:

BR2(config-router)#address-family ipv4 unicast

BR2(config-router-af)#neighbor 192.168.3.16 prefix-list cust-in in

Verification:

Check if the prefix-list is applied to the BGP neighbour

BR2#show running-config | section bgp

neighbor 192.168.3.16 prefix-list cust-in in

Outbound:

- Internet Full Feed.
- Default Route (As per Customer discretion)
- Your Network Prefixes.

Create the Prefix List:

BR2(config)#ip prefix-list cust-out <deny all bogons>

BR2(config)#ip prefix-list cust-out seq 10 permit 10.0.2.0/24

BR2(config)#ip prefix-list cust-out seq 20 permit 0.0.0.0/0 le 24

Apply the Prefix List:

BR2(config-router)#address-family ipv4 unicast

BR2(config-router-af)#neighbor 192.168.3.16 prefix-list cust-out out

Verification:

Check if the prefix-list is applied to the BGP neighbour

BR2#show running-config | section bgp

neighbor 192.168.3.16 prefix-list cust-out out

Task 4: BGP next-hop-self

Scenario-1: Without next hop self, iBGP (CR) should not be able to reach the Internet (Ex-loopback on TR) since the route will not be installed given next hop is not reachable.

```
CR1#sh ip bgp neighbor 172.18.1.1 routes
BGP table version is 2, local router ID is 172.18.1.2
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,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 * i 0.0.0.0          192.168.3.10             0    100      0 200 10 i
 * i 1.0.0.0/24       192.168.1.31             0    100      0 100 i
```

```
* i 8.0.0.0/24      192.168.1.31          0    100      0 100 i
* i 10.0.0.0        192.168.3.2           0    100      0 200 2 100 ?
```

Note: Verify whether the next-hop IP Address for all the routes received from Border Router is reachable or not. For verification you can either ping or trace the destinations.

Scenario-2: Apply next-hop-self and now your iBGP peers (CR) should be able to reach the Internet (TR loopback) because the route is now installed since next hop is reachable.

To make the routes reachable you need to use the command next-hop-self and apply under the BGP against the iBGP peer.

BR2(config-router-af)#neighbor 172.18.1.2 next-hop-self

Verification: Now, check again on the iBGP neighbor.

```
CR1#sh ip bgp neighbor 172.18.1.1 routes
BGP table version is 53, local router ID is 172.18.1.2
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,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
*>i 0.0.0.0          172.18.1.1               0    100      0 200 10 i
*>i 1.0.0.0/24       172.18.1.1               0    100      0 100 i
*>i 8.0.0.0/24       172.18.1.1               0    100      0 100 i
*>i 10.0.0.0         172.18.1.1               0    100      0 200 2 100 ?
*>i 10.0.2.0/24      172.18.1.1               0    100      0 200 2 i
*>i 10.0.3.0/24      172.18.1.1               0    100      0 200 3 i
*>i 10.0.4.0/24      172.18.1.1               0    100      0 200 4 i
*>i 10.0.5.0/24      172.18.1.1               0    100      0 200 5 i
*>i 10.0.6.0/24      172.18.1.1               0    100      0 200 6 i
```

Task 5: Max Prefix Limit

Set the max prefix limit on for your Customers and Peering to allow unexpected number of prefixes, prevent any route leaks and set the limit in consideration of the future growth. You can set the threshold limit, set the warning, and also have an option to drop the prefixes if the received routes are more than of set value.

BR15(config-router-af)#neighbor 192.168.3.16 maximum-prefix 40

OR

BR15(config-router-af)#neighbor 192.168.3.16 maximum-prefix 36 warning-only

OR

BR15(config-router-af)#neighbor 192.168.3.16 maximum-prefix 40 90

Where it will triggered the warning when number of prefixes reaches to 90% for threshold value and drops if more than 40 is received.