Configuring BGP on Cisco Routers

Lab Guide

Version 4.0

Part Number: 97-3541-01

Americas Headquarters	Asia Pacific Headquarters	Europe Headquarters
Cisco Systems, Inc. San Jose, CA	Cisco Systems (USA) Pte. Ltd. Singapore	Cisco Systems International BV Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Table of Contents

Lab Introduction

- For each challenge lab, read the lab Scenario and the Job Aids first, then proceed by solving challenges provided.
- Refer to the Command List if you need assistance with your lab exercise.
- If you need additional tips, refer to the Hints section, which is available in the CLL portal.
- The Validation section will help you verify whether you successfully completed the challenge.
- If you need additional guidance, refer to the Answer section, which contains a detailed, step-by-step solution for each challenge lab.

Challenge 1: Configure a Basic BGP Network

The company that you are working for is looking to expand its business into two new regions. The existing network will be used as a transit AS to connect the two new regions. To get prepared for this implementation, a proof of concept lab consisting of three devices has been established. All the IP addressing has been pre-configured and of the required interfaces has been enabled.

George, a Senior Network Engineer, has researched the proposed implementation. George has established the following requirements for the POC lab:

- Ensure all the routers show all of the routes that are announced into BGP.
- The BGP session between the R2 and R1 should be secured so that neighbor can be no more than 1 hop count away.
- The connection between the R2 and R3 should be secured to prevent DoS attacks using a clear text password **chall3ng3**.
- R3 should advertise only 1 route for 10.1.5.1/24,10.1.6.1/24 and 10.1.7.1/24 with a 16-bit mask.
- Redistribution is not allowed on any of the routers.

You should configure the device in the POC lab to meet these requirements. See the Job Aids section for additional information.

Topology

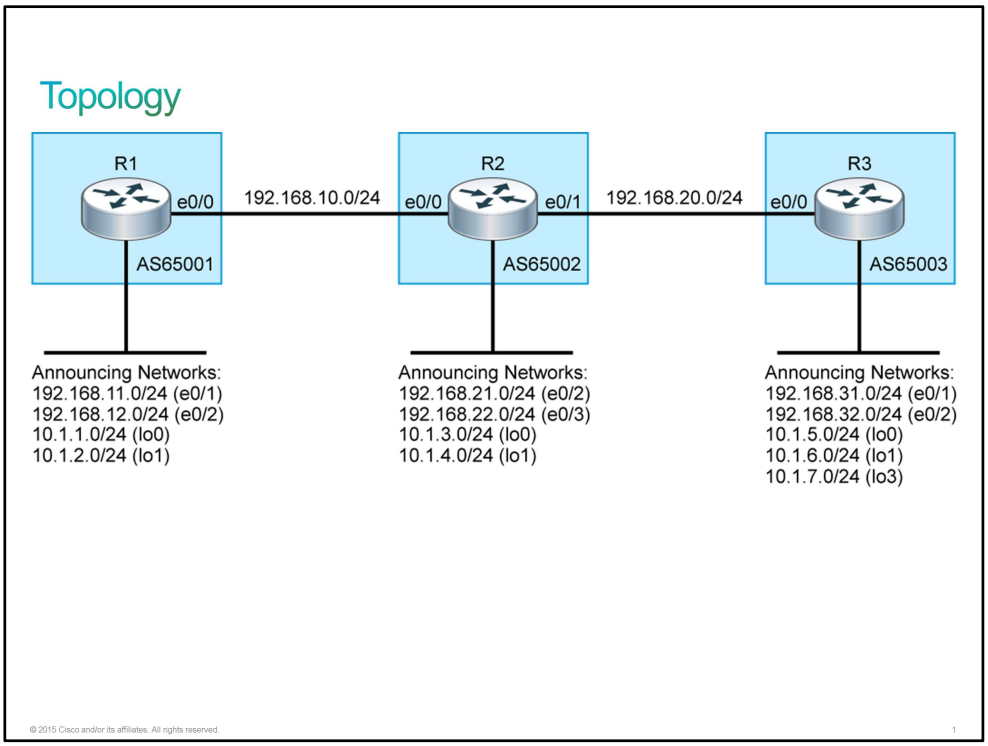

Job Aids

AS Information

Device	Assigned AS
R1	65001
R2	65002
R3	65003

BGP Advertised Networks

Device	Network
R1	10.1.1.0/24
R1	10.1.2.0/24
R1	192.168.11.0/24

Device	Network
R1	192.168.12.0/24
R2	10.1.3.0/24
R2	10.1.4.0/24
R2	192.168.21.0/24
R2	192.168.22.0/24
R3	10.1.5.0/24
R3	10.1.6.0/24
R3	10.1.7.0/24
R3	192.168.31.0/24
R3	192.168.32.0/24

IP Addressing Information

Device	Interface	IP address
R1	Ethernet 0/0	192.168.10.1/24
R1	Ethernet 0/1	192.168.11.1/24
R1	Ethernet 0/2	192.168.12.1/24
R1	Loopback0	10.1.1.1/24
R1	Loopback1	10.1.2.1/24
R2	Ethernet 0/0	192.168.10.2/24
R2	Ethernet 0/1	192.168.20.2/24
R2	Ethernet 0/2	192.168.21.1/24
R2	Ethernet 0/3	192.168.22.1/24
R2	Loopback0	10.1.3.1/24
R2	Loopback1	10.1.4.1/24
R3	Ethernet 0/0	192.168.20.3/24

Device	Interface	IP address
R3	Ethernet 0/1	192.168.31.1/24
R3	Ethernet 0/3	192.168.32.1/24
R3	Loopback0	10.1.5.1/24
R3	Loopback1	10.1.6.1/24
R3	Loopback3	10.1.7.1/24

Command List

The table describes the commands that are used in this activity. The commands are listed in alphabetical order so that you can easily locate the information that you need. Refer to this list if you need configuration command assistance during the lab activity.

Command	Description
aggregate-address *network mask* ***[summary-only]***	Creates Summary prefixes-use the **summary-only** keyword to suppress more-specific routes
interface type *number*	Configures an interface type and places you in interface configuration mode
ip address *ip-address mask*	Configures an IP address for the interface
neighbor *ip-address* **password** *password*	Enables/disables MD5 authentication on a TCP connection between two BGP peers
neighbor *ip-address* **remote-as** *as-number*	Establishes a BGP session with the specified AS number
neighbor *ip-address* **ttl-security hops** *hop-count*	Configures the maximum number of hops that separate two peers
network *network* **mask** *mask*	Announces IP prefixes in BGP
no shutdown	Administratively activates an interface
router bgp *as-number*	Places you in BGP configuration mode
show interfaces *interface-name*	Displays the statistical information specific to an interface
show ip bgp summary	Verifies whether the BGP session is up
show ip bgp neighbor	Enables you to view detailed information about the neighbor
show ip bgp	Enables you to inspect the contents of the BGP table
show ip bgp *network*	Enables you to view detailed information about prefixes (aggregates)
show ip interface brief	Displays the status of the Interfaces
show ip route	Displays the current state of the routing table

Task 1: Configuring a Basic BGP Network

Challenge Steps

Configure BGP on R1, R2, and R3. Make sure that the BGP session is established and routes are seen in the BGP routing table.

Step 1 Enable BGP on R1 based on the information that is provided in the graphics and Job Aids.

Step 2 Enable BGP on R2 based on the information that is provided in the graphics and Job Aids.

Step 3 Enable BGP on R3 based on the information that is provided in the graphics and Job Aids.

Verification for ILT Only

Use this section to verify your challenge results.

Step 1 Verify that the BGP configuration on R1 is correct.

1 Use the **show ip bgp summary** command to verify that R1 has established a neighbor relationship with R2.

You should get the following results:

```
R1# show ip bgp summary
BGP router identifier 10.1.2.1, local AS number 65001
BGP table version is 12, main routing table version 12
11 network entries using 1628 bytes of memory
11 path entries using 704 bytes of memory
4/4 BGP path/bestpath attribute entries using 544 bytes of memory
2 BGP AS-PATH entries using 48 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 2924 total bytes of memory
BGP activity 11/0 prefixes, 11/0 paths, scan interval 60 secs

Neighbor        V      AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down  State/PfxRcd
192.168.10.2    4    65002    9       8       12    0    0  00:03:07       7
```

Item Level Feedback: If you do not see the desired results:

- Verify that R1 has established a neighbor relationship with R2 (192.168.10.2). If it has not, verify that R1 is configured correctly. Take special note of the AS number (65001) and the neighbor statements.
- Verify that 7 prefixes have been received by R1. If this value is incorrect, verify that R2 and R3 are advertising the correct networks and that summarization are correctly configured on R3.

2 Use the **show ip bgp** command to verify the entries in the BGP routing table.

You should get the following results:


```
R1# show ip bgp
BGP table version is 12, local router ID is 10.1.2.1
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,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.1.0.0/16      192.168.10.2                           0 65002 65003 i
 *>  10.1.1.0/24      0.0.0.0                  0         32768 i
 *>  10.1.2.0/24      0.0.0.0                  0         32768 i
 *>  10.1.3.0/24      192.168.10.2             0             0 65002 i
 *>  10.1.4.0/24      192.168.10.2             0             0 65002 i
 *>  192.168.11.0     0.0.0.0                  0         32768 i
 *>  192.168.12.0     0.0.0.0                  0         32768 i
 *>  192.168.21.0     192.168.10.2             0             0 65002 i
 *>  192.168.22.0     192.168.10.2             0             0 65002 i
 *>  192.168.31.0     192.168.10.2                           0 65002 65003 i
 *>  192.168.32.0     192.168.10.2                           0 65002 65003 i
```

Item Level Feedback: If you do not see the desired results, verify that R2 and R3 are advertising the correct networks and that summarization are correctly configured on R3.

3 Use the **show ip bgp neighbor 192.168.10.2** to check the AS number and EBGP hop count.

You should get the following results:

```
R1# show ip bgp neighbor 192.168.10.2
BGP neighbor is 192.168.10.2,  remote AS 65002, external link
  BGP version 4, remote router ID 10.1.4.1
  BGP state = Established, up for 00:23:38

**** Output Omitted *****

  Address tracking is enabled, the RIB does have a route to 192.168.10.2
  Connections established 1; dropped 0
  Last reset never
  External BGP neighbor may be up to 1 hop away.
  Transport(tcp) path-mtu-discovery is enabled
  Graceful-Restart is disabled

***** Output Omitted ******
```

Item level Feedback: If you do not see the desired results, verify that the BGP configuration has the ttl value set to 1.

Step 2 Verify that the BGP configuration on R2 is correct.

1 Use **show ip bgp summary** command to verify that R2 has established a neighbor relationship with R1 and R3.

You should get the following results:

```
R2# show ip bgp summary
BGP router identifier 10.1.4.1, local AS number 65002
BGP table version is 12, main routing table version 12
11 network entries using 1628 bytes of memory
11 path entries using 704 bytes of memory
4/4 BGP path/bestpath attribute entries using 544 bytes of memory
2 BGP AS-PATH entries using 48 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 2924 total bytes of memory
BGP activity 11/0 prefixes, 11/0 paths, scan interval 60 secs

Neighbor        V       AS  MsgRcvd MsgSent TblVer InQ OutQ Up/Down  State/PfxRcd
192.168.10.1    4     65001    49      50      12     0   0   00:40:05       4
192.168.20.3    4     65003    47      50      12     0   0   00:37:31       3
```

Item Level Feedback: If you do not see the desired results:

- R2 should be in AS 65002. If it is not, verify R2 configuration paying special attention to the AS number.
- R2 should be receiving 4 networks from R1. If it is not, verify R2 configuration paying special attention to the AS number and the neighbor commands. If R2 is correct, verify R1 configuration paying special attention to the network commands.
- R2 should be receiving 3 networks from R3. If it is not, verify R2 configuration paying special attention to the AS number and the neighbor commands. If R2 is correct, verify R3 configuration paying special attention to the network command and summarization commands.

2 On R2, use the **show ip bgp** command to verify the R1 and R3 entries in the BGP routing table.

You should get the following results:

```
R2# show ip bgp
BGP table version is 20, local router ID is 10.1.4.1
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,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.1.0.0/16      192.168.20.3             0             0 65003 i
 *>  10.1.1.0/24      192.168.10.1             0             0 65001 i
 *>  10.1.2.0/24      192.168.10.1             0             0 65001 i
 *>  10.1.3.0/24      0.0.0.0                  0         32768 i
 *>  10.1.4.0/24      0.0.0.0                  0         32768 i
 *>  192.168.11.0     192.168.10.1             0             0 65001 i
 *>  192.168.12.0     192.168.10.1             0             0 65001 i
 *>  192.168.21.0     0.0.0.0                  0         32768 i
 *>  192.168.22.0     0.0.0.0                  0         32768 i
 *>  192.168.31.0     192.168.20.3             0             0 65003 i
 *>  192.168.32.0     192.168.20.3             0             0 65003 i
```

Item Level Feedback: If you do not see the desired results:

- R2 should be receiving 4 networks from R1. If it is not, verify R2 configuration paying special attention to the AS number and the neighbor commands. If R2 is correct, verify R1 configuration paying special attention to the network commands.

- R2 should be receiving 3 networks from R3. If it is not, verify R2 configuration paying special attention to the AS number and the neighbor commands. If R2 is correct, verify R3 configuration paying special attention to the network command and summarization commands.

3 Use the **show ip bgp neighbor 192.168.10.1** command to check the AS number, EBGP hop count.

You should get the following results:

```
R2# show ip bgp neighbor 192.168.10.1
BGP neighbor is 192.168.10.1,  remote AS 65001, external link
  BGP version 4, remote router ID 10.1.2.1
  BGP state = Established, up for 00:37:28

**** Output Omitted *****

  External BGP neighbor may be up to 1 hop away.

**** Output Omitted *****
```

Item Level Feedback: If you do not see the desired results, verify if the BGP session is established with R1 and has the ttl value set to 1.

4 Use the **show ip bgp neighbor 192.168.20.3** command to check the AS number.

You should get the following results:

```
R2# show ip bgp neighbor 192.168.20.3
BGP neighbor is 192.168.20.3,  remote AS 65003, external link
  BGP version 4, remote router ID 10.1.7.1
  BGP state = Established, up for 00:15:57

**** Output Omitted *****

  No active TCP connection
R2#
```

Item Level Feedback: If you do not see the desired results, verify R2 configuration paying special attention to the AS number and the neighbor commands. If R2 is correct, verify R3 configuration paying special attention to the network commands.

Step 3 Verify that the BGP configuration on R3 is correct.

1 Use the **show ip bgp summary** command to verify that R3 has established a neighbor relationship with R2.

You should get the following results:

```
R3# show ip bgp summary
BGP router identifier 10.1.7.1, local AS number 65003
BGP table version is 14, main routing table version 14
13 network entries using 1924 bytes of memory
13 path entries using 832 bytes of memory
3/3 BGP path/bestpath attribute entries using 408 bytes of memory
2 BGP AS-PATH entries using 48 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 3212 total bytes of memory
BGP activity 13/0 prefixes, 13/0 paths, scan interval 60 secs

Neighbor        V       AS  MsgRcvd MsgSent TblVer InQ OutQ Up/Down  State/PfxRcd
192.168.20.2    4     65002   10      6       14     0   0   00:02:39       8
```

Item Level Feedback: If you do not see the desired results:

- R3 should be in AS 65003. If it is not, verify R2 configuration special attention to the AS number.
- R3 should be receiving 8 networks from R2. If it is not, verify R3's configuration paying special attention to the AS number and the neighbor commands. If R3 is correct, verify R2's and R1's configurations paying special attention to the network commands.

2 Use the **show ip bgp** command to verify that the router has received entries from R1 and R2 into its BGP routing table. Also ensure that R3 is advertising only one route for the 10.1.5.1/24,10.1.6.1/24 and 10.1.7.1/24 networks with a 16-bit mask.

You should get the following results:

```
R3# show ip bgp
BGP table version is 22, local router ID is 10.1.7.1
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,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.1.0.0/16      0.0.0.0                             32768 i
 s>  10.1.1.0/24      192.168.20.2                            0 65002 65001 i
 s>  10.1.2.0/24      192.168.20.2                            0 65002 65001 i
 s>  10.1.3.0/24      192.168.20.2             0              0 65002 i
 s>  10.1.4.0/24      192.168.20.2             0              0 65002 i
 s>  10.1.5.0/24      0.0.0.0                  0          32768 i
 s>  10.1.6.0/24      0.0.0.0                  0          32768 i
 s>  10.1.7.0/24      0.0.0.0                  0          32768 i
 *>  192.168.11.0     192.168.20.2                            0 65002 65001 i
 *>  192.168.12.0     192.168.20.2                            0 65002 65001 i
 *>  192.168.21.0     192.168.20.2             0              0 65002 i
 *>  192.168.22.0     192.168.20.2             0              0 65002 i
 *>  192.168.31.0     0.0.0.0                  0          32768 i
 *>  192.168.32.0     0.0.0.0                  0          32768 i
```

Item Level Feedback: If you do not see the desired results, verify that all the networks are advertised into BGP correctly in R1 and R2. Also, verify that R3 is advertising only 1 summary route for 10.1.5.1/24,10.1.6.1/24 and 10.1.7.1/24 with a 16-bit mask.

Challenge 2: Configure a BGP Transit AS

Your ISP has made an agreement with several smaller ISPs to provide them with Internet access in 3 months. Your existing network will be used as a transit AS to provide the smaller ISPs this access. To get ready for this implementation, a proof of concept lab consisting of five devices has been established. The lab consists of three autonomous systems. AS 65001 contains one router (ISP1). AS 65002 contains three routers (R2, R3, and R4). AS 65003 contains one router (ISP2). All of the IP addressing has been pre-configured on all devices. The IGP in AS 65002 has been pre-configured and has full connectivity between the routers.

After researching the proposed implementation, Anish, a Senior Network Engineer has established the following requirements for the POC lab:

- When neighbors are directly connected, their physical interface addresses should be used to establish the neighbor relationship.
- When neighbors are logically connected, their loopback interface addresses should be used to establish the neighbor relationship.
- Redistribution is not allowed on any of the routers.
- The routers in AS 65001 and 65003 must advertise all locally originated networks except the interface that is connected to the external neighbor.
- The router in AS 65001 must be able to see all of routes that are advertised by AS 65003.
- The IGP routes in AS 65002 cannot be seen in AS 65001 or AS 65003.
- The router in AS 65003 must be able to see all of routes that are advertised by AS 65001.
- All routers in 65002 must see all of routes that are advertised by AS 65001 and 65003.

Please see Job Aids section for additional information.

Topology

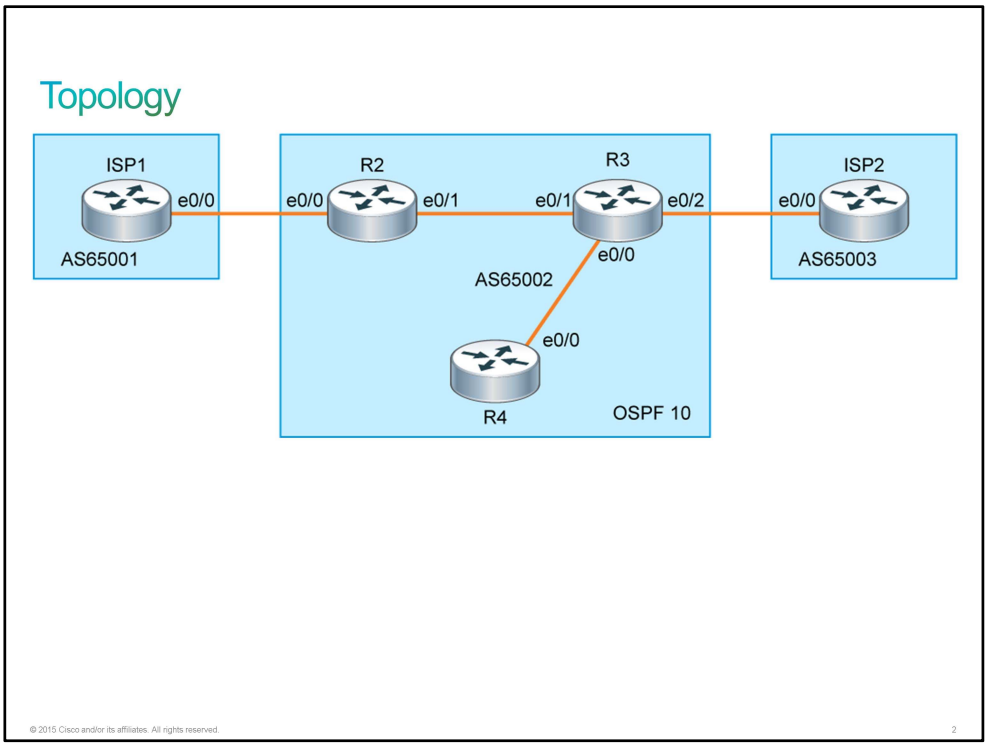

Job Aids

If you shut down an interface on a real router or switch, the connected device will see it as "down/down." Due to virtualization specifics, IOL behavior is slightly different. If you shut down an interface on a router or switch, the connected device will see it as "up/up." In IOL, the status of an interface can only be "up/up" or "administratively down/down."

Device Connections

Device	Interface	Connects To	Interface
ISP1	Ethernet 0/0	R2	Ethernet 0/0
R2	Ethernet 0/1	R3	Ethernet 0/1
R3	Ethernet 0/0	R4	Ethernet 0/0
R3	Ethernet 0/2	ISP2	Ethernet 0/0

Networks That Will Be Advertise by BGP

Device	IP address

ISP1	192.168.15.1/24
ISP1	192.168.16.1/24
ISP2	192.168.55.1/24
ISP2	192.168.56.1/24

Device Information

Device	Interface	IP Address
ISP1	Loopback0	10.1.1.1/24
ISP1	Ethernet 0/0	192.168.12.1/24
ISP1	Ethernet 0/1	192.168.15.1/24
ISP1	Ethernet 0/2	192.168.16.1/24
R2	Loopback0	10.1.2.1/24
R2	Ethernet 0/0	192.168.12.1/24
R2	Ethernet 0/1	192.168.23.2/24
R3	Loopback0	10.1.3.1/24
R3	Ethernet 0/0	192.168.34.3/24
R3	Ethernet 0/1	192.168.23.3/24
R3	Ethernet 0/2	192.168.35.3/24
R4	Loopback0	10.1.4.1/24
R4	Ethernet 0/0	192.168.34.4/24
ISP2	Loopback0	10.1.5.1/24
ISP2	Ethernet 0/0	192.168.35.5/24
ISP2	Ethernet 0/1	192.168.55.1/24
ISP2	Ethernet 0/2	192.168.56.1/24

Command List

The table describes the commands that are used in this activity. The commands are listed in alphabetical order so that you can easily locate the information that you need. Refer to this list if you need configuration command assistance during the lab activity.

Command	Description
neighbor *ip-address* **remote-as** *as-number*	Establishes a BGP session with the specified AS number
neighbor *ip-address* **update-source** *interface*	Uses the IP address of the specified interface as the source address for the BGP Session
network *network* **mask** *mask*	Announces IP prefixes in BGP
neighbor *network* **next-hop-self**	Configures the router as the next hop for a BGP speaking neighbor
router bgp *as-number*	Places you in BGP configuration mode
show ip interface brief	Displays the status of the Interfaces
show ip bgp summary	Verifies whether the BGP session is up
show ip bgp neighbor	Enables you to view detailed information about the neighbor
show ip bgp	Enables you to inspect the contents of the BGP table
show ip bgp *network*	Enables you to view detailed information about prefixes (aggregates)
show ip route	Displays the current state of the routing table

Task 1: Configure a Transit AS

Challenge Lab Steps

In this Challenge lab, you will be able to:

- Configure your ISP routers, R2, R3, and R4 to become a Transit AS.
- Also configure the Customer routers, ISP1 and ISP2 to establish an EBGP session with your ISP routers, R2 and R3.
- Ensure ISP1 and ISP2 use your ISP AS as a Transit AS to exchange the BGP routing updates.

Step 1 Configure BGP on R2 using the information that is contained the scenario and Job Aids.

Step 2 Configure BGP on R3 using the information that is contained the scenario and Job Aids.

Step 3 Configure BGP on R4 using the information that is contained the scenario and Job Aids.

Step 4 Configure BGP on ISP1 using the information that is contained the scenario and Job Aids.

Step 5 Configure BGP on ISP2 using the information that is contained the scenario and Job Aids.

Verification-ILT

Use this section to verify your challenge results.

Step 1 Verify router (R2)

1 Use the **show ip bgp summary** command to verify that R2 has established an EBGP session with ISP1 and IBGP session with R3, and R4.

You should get the following results.

```
R2# show ip bgp summary
BGP router identifier 10.1.2.1, local AS number 65002
BGP table version is 7, main routing table version 7
6 network entries using 888 bytes of memory
6 path entries using 384 bytes of memory
2/2 BGP path/bestpath attribute entries using 272 bytes of memory
2 BGP AS-PATH entries using 48 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 1592 total bytes of memory
BGP activity 6/0 prefixes, 6/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down
State/PfxRcd
10.1.4.1        4        65002     157     160        7    0    0 02:22:10        0
192.168.12.1    4        65001     159     159        7    0    0 02:20:02        3
192.168.23.3    4        65002     161     164        7    0    0 02:24:48        3
R2#
```

FEEDBACK: If you do not see the desired results, verify that:

- R2 is in AS 65002. If it is not, verify the AS number that is configured on R2.
- R2 has established an IBGP session with R4 (**10.1.4.1**) that is in AS 65002. If it is not, verify the configuration on R2; paying special attention to the AS number and the neighbor commands. If the configuration on R2 is correctly, verify the BGP configuration on R4; paying special attention to the AS number and the neighbor commands.
- R2 has established an EBGP session with ISP1 (**192.168.12.1**) that is in AS 65001. If it is not, verify the configuration on R2; paying special attention to the AS number and the neighbor commands. If the configuration on R2 is correctly, verify the BGP configuration on ISP1; paying special attention to the AS number and the neighbor commands.
- R2 has established an IBGP session with R3 (**192.168.23.3)** that is in AS 65002. If it is not, verify the configuration on R2; paying special attention to the AS number and the neighbor commands. If the configuration on R2 is correctly, verify the BGP configuration on R3; paying special attention to the AS number and the neighbor commands.

2 Use the **show ip bgp** command on R3 to verify that R2 advertising itself as the next hop for routes from AS 65001.

You should get the following results.

```
R3# show ip bgp
BGP table version is 7, local router ID is 10.1.3.1
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,
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 10.1.1.0/24      192.168.23.2             0    100      0 65001 i
 *>  10.1.5.0/24      192.168.35.5             0             0 65003 i
 *>i 192.168.15.0     192.168.23.2             0    100      0 65001 i
 *>i 192.168.16.0     192.168.23.2             0    100      0 65001 i
 *>  192.168.55.0     192.168.35.5             0             0 65003 i
 *>  192.168.56.0     192.168.35.5             0             0 65003 i
R3#
```

FEEDBACK: If you do not see the desired results, verify the configuration on R2; paying special attention to the next-hop-self command.

Step 2 Verify router (R3)

1 Use the **show ip bgp summary** command to verify that R3 has established an EBGP session with ISP2 and IBGP sessions with R2 and R4.

You should get the following results:

```
R3# show ip bgp summary
BGP router identifier 10.1.3.1, local AS number 65002
BGP table version is 7, main routing table version 7
6 network entries using 888 bytes of memory
6 path entries using 384 bytes of memory
2/2 BGP path/bestpath attribute entries using 272 bytes of memory
2 BGP AS-PATH entries using 48 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 1592 total bytes of memory
BGP activity 6/0 prefixes, 6/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down
State/PfxRcd
192.168.23.2    4        65002     166     163        7    0    0 02:26:22        3
192.168.34.4    4        65002     161     162        7    0    0 02:23:40        0
192.168.35.5    4        65003       7       7        7    0    0 00:02:05        3
R3#
```

FEEDBACK: If you do not see the desired results, verify that:

- Verify if R3 is in AS 65002. If it is not, verify the AS configuration on R3.
- R3 has established an IBGP session with R2 (192.168.23.2) that is in AS 65002. If it is not, verify the configuration on R3; paying special attention to the AS number and the neighbor commands. If the configuration on R3 is correctly, verify the configuration on R2; paying special attention to the AS number and the neighbor commands.
- R3 has established an IBGP session with R4 (192.168.34.4) that is in AS 65002. If it is not, verify the configuration on R3; paying special attention to the AS number and the neighbor commands. If the configuration on R3 is correctly, verify the configuration on R4; paying special attention to the AS number and the neighbor commands.

- R3 has established an EBGP session with ISP2 (192.168.35.5) that is in AS 65003. If it is not, verify the configuration on R3; paying special attention to the AS number and the neighbor commands. If the configuration on R3 is correctly, verify the configuration on IPS2; paying special attention to the AS number and the neighbor commands.

2 Use the **show ip bgp** command on R2 to verify that R3 advertising itself as the next hop for routes from AS 65003.

You should get the following results.

```
R2# show ip bgp
BGP table version is 7, local router ID is 10.1.2.1
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,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.1.1.0/24      192.168.12.1             0             0 65001 i
 *>i 10.1.5.0/24      192.168.23.3             0    100      0 65003 i
 *>  192.168.15.0     192.168.12.1             0             0 65001 i
 *>  192.168.16.0     192.168.12.1             0             0 65001 i
 *>i 192.168.55.0     192.168.23.3             0    100      0 65003 i
 *>i 192.168.56.0     192.168.23.3             0    100      0 65003 i
R2#
```

FEEDBACK: If you do not see the desired results, verify the configuration on R3; paying special attention to the next-hop-self command.

Step 3 Verify router (R4)

1 Use the **show ip bgp summary** command to verify that router (R4) has established BGP sessions with router (R2) and router (R3).

You should get the following results.

```
R4# show ip bgp summary
BGP router identifier 10.1.4.1, local AS number 65002
BGP table version is 4, main routing table version 4
6 network entries using 888 bytes of memory
6 path entries using 384 bytes of memory
2/1 BGP path/bestpath attribute entries using 272 bytes of memory
2 BGP AS-PATH entries using 48 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 1592 total bytes of memory
BGP activity 6/0 prefixes, 6/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down
State/PfxRcd
10.1.2.1        4        65002     166     162        4    0    0 02:27:27        3
192.168.34.3    4        65002     166     165        4    0    0 02:27:23        3
R4#
```

FEEDBACK: If you do not see the desired results, verify that:

- R4 is in AS 65002. If it is not, verify the AS configuration on R4.

- R4 has established an IBGP session with R2 (10.1.2.1) that is in AS 65002. If it is not, verify the configuration on R4; paying special attention to the AS number and the neighbor commands. If the configuration on R4 is correctly, verify the configuration on R2; paying special attention to the AS number and the neighbor commands.
- R4 has established an IBGP session with R3 (192.168.34.3) that is in AS 65002. If it is not, verify the configuration on R4; paying special attention to the AS number and the neighbor commands. If the configuration on R4 is correctly, verify the configuration on R3; paying special attention to the AS number and the neighbor commands.

Step 4 Verify router (ISP1)

1 Use the **show ip bgp summary** command to verify that router (ISP1) has established an EBGP session with router (R2) and advertise all the networks into BGP.

You should get the following results.

```
ISP1# show ip bgp summary
BGP router identifier 10.1.1.1, local AS number 65001
BGP table version is 7, main routing table version 7
6 network entries using 888 bytes of memory
6 path entries using 384 bytes of memory
2/2 BGP path/bestpath attribute entries using 272 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 1568 total bytes of memory
BGP activity 6/0 prefixes, 6/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down
State/PfxRcd
192.168.12.2    4        65002     168     168        7    0    0 02:28:05        3
ISP1#
```

FEEDBACK: If you do not see the desired results, verify that:

- ISP1 is in AS 65001. If it is not, verify ISP1 configuration.
- ISP1 has established an EBGP session with R2 (192.168.12.2) that is in AS 65001. If it is not, verify the configuration on ISP1; paying special attention to the AS number and the neighbor commands. If the configuration on ISP1 is correctly, verify the configuration on R2; paying special attention to the AS number and the neighbor commands.

2 Use the **show ip bgp** command to verify that the ISP1 is advertising the correct routes.

You should get the following results.


```
ISP1# show ip bgp
BGP table version is 7, local router ID is 10.1.1.1
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,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.1.1.0/24      0.0.0.0                  0         32768 i
 *>  10.1.5.0/24      192.168.12.2                           0 65002 65003 i
 *>  192.168.15.0     0.0.0.0                  0         32768 i
 *>  192.168.16.0     0.0.0.0                  0         32768 i
 *>  192.168.55.0     192.168.12.2                           0 65002 65003 i
 *>  192.168.56.0     192.168.12.2                           0 65002 65003 i
ISP1#
```

FEEDBACK: If you do not see the desired results, verify the configuration on ISP1; paying special attention to the network commands.

Step 5 Verify router (ISP2)

1 Use the **show ip bgp summary** command to verify that router (ISP2) has established an EBGP session with router (R3) and advertise all the networks into BGP.

You should get the following results.

```
ISP2# show ip bgp summary
BGP router identifier 10.1.5.1, local AS number 65003
BGP table version is 13, main routing table version 13
6 network entries using 888 bytes of memory
6 path entries using 384 bytes of memory
2/2 BGP path/bestpath attribute entries using 272 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 1568 total bytes of memory
BGP activity 9/3 prefixes, 9/3 paths, scan interval 60 secs

Neighbor        V           AS  MsgRcvd MsgSent TblVer InQ OutQ Up/Down  State/PfxRcd
192.168.35.3    4        65002    81      82      13    0    0   01:09:49        3
ISP2#
```

FEEDBACK: If you do not see the desired results,

- Verify if router (ISP2) is in AS 65003. If it is not, verify ISP2 configuration.
- ISP2 has established an EBGP session with R3 (192.168.35.3) that is in AS 65003. If it is not, verify the configuration on ISP2; paying special attention to the AS number and the neighbor commands. If the configuration on ISP2 is correctly, verify the configuration on R3; paying special attention to the AS number and the neighbor commands.

2 Use the **show ip bgp** command to verify that the ISP2 has advertised the correct routes in the BGP routing table.

You should get the following results.

```
ISP2# show ip bgp
BGP table version is 13, local router ID is 10.1.5.1
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,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.1.1.0/24     192.168.35.3                           0 65002 65001 i
 *>  10.1.5.0/24     0.0.0.0                  0         32768 i
 *>  192.168.15.0    192.168.35.3                           0 65002 65001 i
 *>  192.168.16.0    192.168.35.3                           0 65002 65001 i
 *>  192.168.55.0    0.0.0.0                  0         32768 i
 *>  192.168.56.0    0.0.0.0                  0         32768 i
ISP2#
```

- Router (ISP2) is advertising 3 routes. If it is not, check router
- ISP2 should be receiving 3 routes. If it is not, check router (ISP2) BGP configuration paying special attention to the network and the neighbor commands. if router (ISP2) is correct, check the router (R3) configuration.

FEEDBACK: If you do not see the desired results, verify the BGP configuration on ISP2; paying special attention to the network.

Challenge 3: Configure BGP Using BGP Filtering

A large company has bought another large company. With the merger of 2 companies, the Network Operations team are advised to upgrade the overall Network to have more control on the BGP routing updates. The team has found 6 major verticals that are based on the Services and created 6 Autonomous Systems. The company is growing at a rapid speed and every thing was smooth until recently where the Support team has started getting complaints regarding the over utilization of the CPU, unwanted BGP entries in the routers.

Philip, a Technical Leader, after going through the Network has given the expert advise for the Team to use the following BGP route filtering mechanisms (as-path filter, prefix-lists, and Outbound Route Filtering) to save the bandwidth and also to reduce the IP/BGP table sizes.

Network Operations team has been asked to do the following tasks:

- Ensure AS 65002 is not a transit AS.
- AS 65006 should not advertise any routes from AS 65005 to AS 65001.
- AS 65001 should not receive any routes from AS 65002.
- In AS 65004, filter all the /32 prefixes of 10.0.0.0/8 network and allow all other prefixes coming from AS 65003.
- In AS 65004, filter the 172.16.0.0/16 network with a string expression that is based on the length value of *ge* /26 and *le* /28. Make sure, all other prefixes are allowed.
- Use default sequence numbering while configuring prefix-lists.
- From local AS 65005, negotiate with remote AS 65004 to receive only 172.16.5.0/29 route.

As per the company policy, route-maps and distribute lists should not be used.

Device IP addressing and BGP neighbor relationship have been pre-configured.

Please see Job Aids section for additional information.

Topology

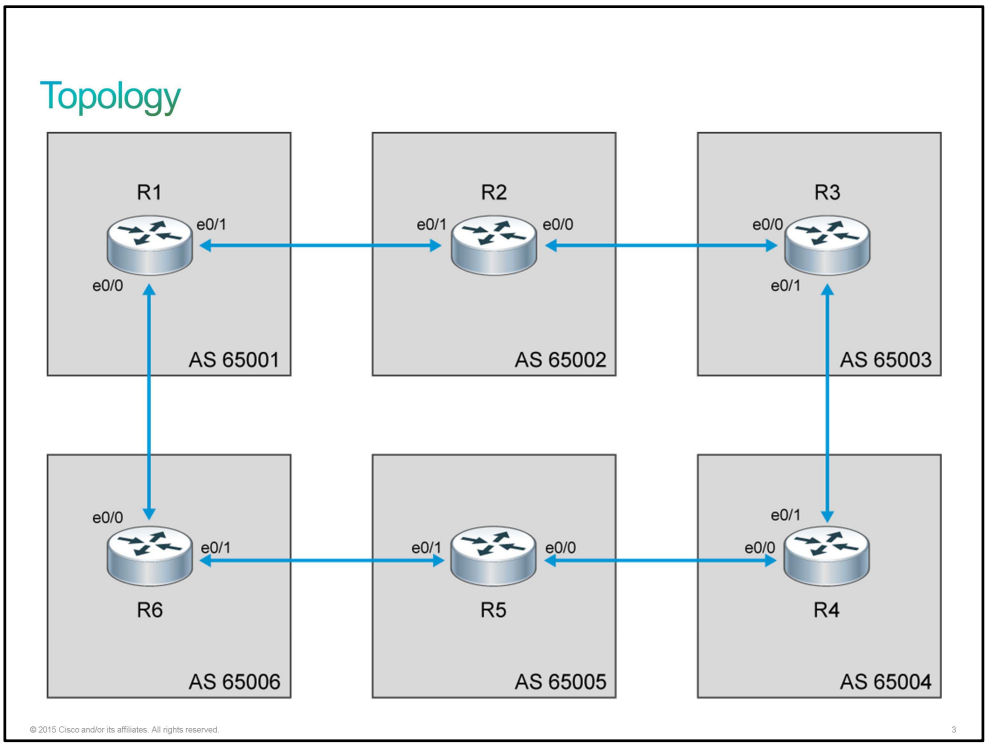

Job Aids

Device Connections

Device	Interface	Connects To	Interface
R1	Ethernet 0/0	R6	Ethernet 0/0
R1	Ethernet 0/1	R2	Ethernet 0/1
R2	Ethernet 0/0	R3	Ethernet 0/0
R3	Ethernet 0/1	R4	Ethernet 0/1
R4	Ethernet 0/0	R5	Ethernet 0/0
R5	Ethernet 0/1	R6	Ethernet 0/1

Device IP Addressing

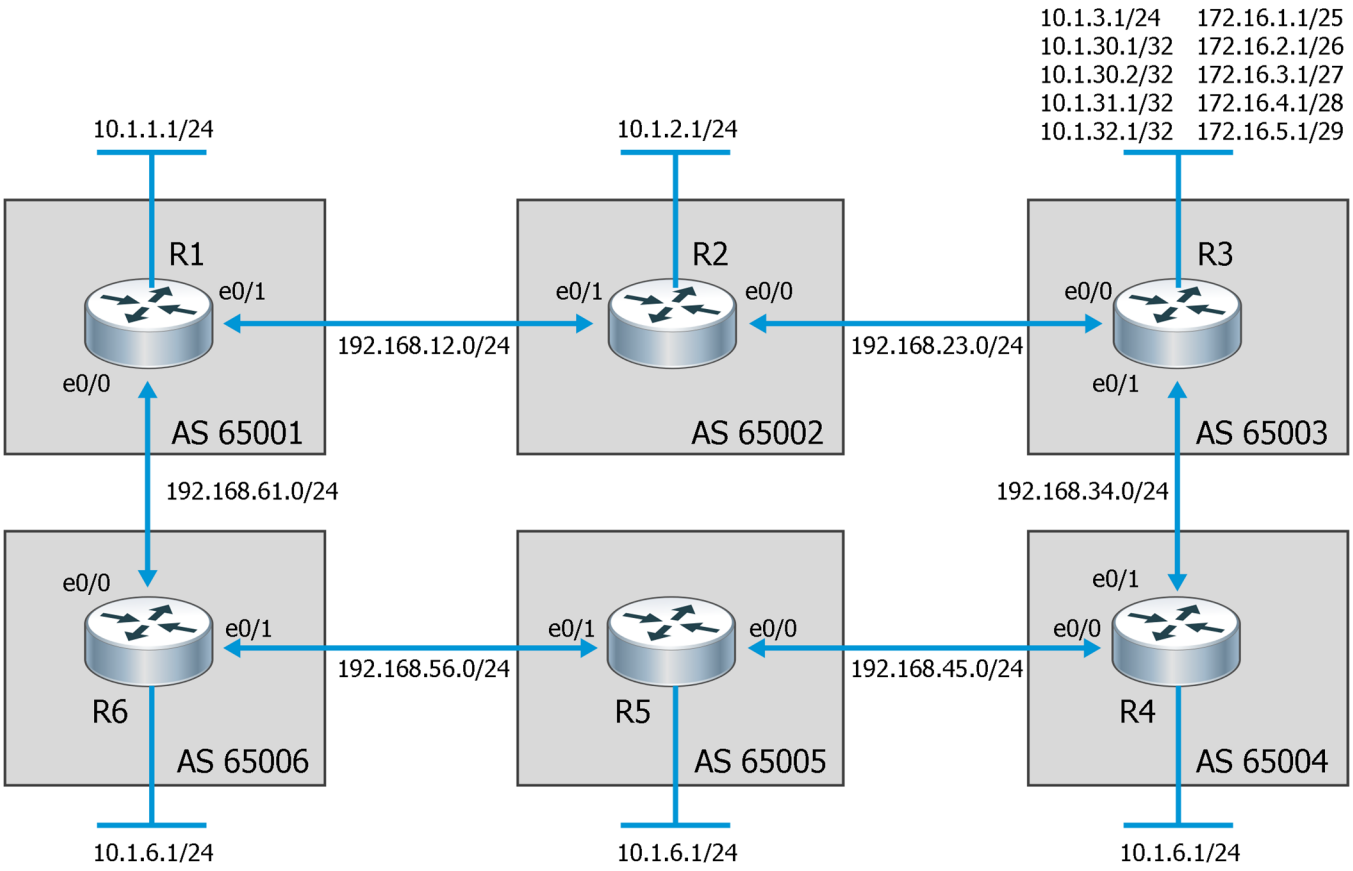

Device	Interface	Neighbor	IP address
R1	Loopback0	None	10.1.1.1/24
R1	Ethernet 0/0	R6	192.168.61.1/24
R1	Ethernet 0/1	R2	192.168.12.1/24
R2	Loopback0	None	10.1.2.1/24
R2	Ethernet 0/0	R3	192.168.23.2/24
R2	Ethernet 0/1	R1	192.168.12.2/24
R3	Loopback0	None	10.1.3.1/24
R3	Loopback1	None	10.1.30.1/32
R3	Loopback2	None	10.1.31.1/32
R3	Loopback3	None	10.1.32.1/32

Device	Interface	Neighbor	IP address
R3	Loopback11	None	172.16.1.1/25
R3	Loopback12	None	172.16.2.1/26
R3	Loopback13	None	172.16.3.1/27
R3	Loopback14	None	172.16.4.1/28
R3	Loopback15	None	172.16.5.1/29
R3	Ethernet 0/0	R2	192.168.23.3/24
R3	Ethernet 0/1	R4	192.168.34.3/24
R4	Loopback0	None	10.1.4.1/24
R4	Ethernet 0/0	R5	192.168.45.4/24
R4	Ethernet 0/1	R3	192.168.34.4/24
R5	Loopback0	None	10.1.5.1/24
R5	Ethernet 0/0	R4	192.168.45.5/24
R5	Ethernet 0/1	R6	192.168.56.5/24
R6	Loopback0	None	10.1.6.1/24
R6	Ethernet 0/0	R1	192.168.61.6/24
R6	Ethernet 0/1	R5	192.168.56.6/24

Access-lists Information

Device	access-list number
R1	1
R2	2
R6	6

Prefix-list Information

Device	prefix-list name
R4	PREFIXFILTER
R5	ORF

Command List

The table describes the commands that are used in this activity. The commands are listed in alphabetical order so that you can easily locate the information that you need. Refer to this list if you need configuration command assistance during the lab activity.

Command	Description
clear ip bgp * soft	Triggers both inbound and outbound soft reset allowing routing tables to be reconfigured and activated without clearing the BGP session.
ip as-path access-list *number* **permit** *regexp*	Specifies the AS-path filter
ip prefix-list *name* [**seq** *seq*] {**permit** | **deny**} *network/length* [**ge** *length* | **le** *length*]	Enables you to create an entry in a prefix-list
neighbor *ip-address* **filter-list** *as-path-filter* {**in** | **out**}	Filters incoming or outgoing updates with the specified AS-path filter
neighbor *ip-address* **prefix-list** *prefix-list* **in**	Filters incoming prefixes
neighbor *ip-address* **remote-as** *as-number*	Establishes a BGP session with the specified AS number
network *network* **mask** *mask*	Announces IP prefixes in BGP
router bgp *as-number*	Places you in BGP configuration mode
show ip interface brief	Displays the status of the Interfaces
show ip bgp	Enables you to inspect the contents of the BGP table
show ip bgp filter-list *access-list-number*	Enables you to check a filter
show ip bgp prefix-list *prefix-list-name*	Displays a BGP routing table that is based on the contents of the prefix-list
show ip bgp *network*	Enables you to view detailed information about prefixes (aggregates)
show ip bgp summary	Verifies whether the BGP session is up
show ip bgp neighbors *ip-address*	Enables you to view detailed information about the neighbor
show ip bgp neighbors *ip-address* **advertised-routes**	Display the routes that are advertised to a BGP neighbor

Command	Description
show ip bgp neighbors *ip-address* **received-routes**	Display the received routes from neighbor
show ip bgp neighbors *ip-address* **routes**	Display routes that are learned from neighbor
show ip bgp regexp *regexp*	Enables you to use a regular expression to filter the output of the **show ip bgp** command
show ip route	Displays the current state of the routing table

Task 1: Configuring BGP Using BGP Filtering Techniques

Steps

Step 1 Configure router R2 and make AS 65002, a Non-transit AS.

Step 2 Configure router R6 and make sure AS 65006 will not advertise any routes from AS 65005 to AS 65001.

Step 3 Configure router R1 and make sure AS 65001 should not receive any routes from AS 65002.

Step 4 Configure router R4 and make sure the AS 65004 filter that are based on the prefix length and has the BGP ORF receive capability towards AS 65005.

- Filter all the /32 prefixes of 10.0.0.0/8 network coming from AS 65003.
- Filters 172.16.0.0/16 based on the length value *ge* /26 and *le* /28.
- Make sure, all other prefixes are allowed.
- Ensure that the router has the BGP ORF receive capability towards AS 65005.

Step 5 Configure router R5 and make sure AS 65005 negotiate with AS 65004 while receiving update message and has the BGP ORF send capability towards AS 65004.

- Negotiate with R4 to send only 172.16.5.0/29 prefix.
- Ensure that the router has the BGP ORF send capability towards AS 65004.

Verification—for ILT

Use this section to verify your challenge results.

Step 1 Verify the Non-Transit AS configuration on router R2.

1 Use the **show ip bgp neighbors 192.168.12.1 advertised-routes** command to ensure that R2 only advertises prefixes from its own autonomous system to router R1.

You should get the following results.

```
R2# show ip bgp neighbors 192.168.12.1 advertised-routes
BGP table version is 15, local router ID is 10.1.2.1
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,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.1.2.0/24      0.0.0.0                  0         32768 i

Total number of prefixes 1
```

FEEDBACK: If you do not see the desired results, verify that:

- You cleared the BGP process or wait sufficient time for automatic BGP update.
- AS-path access-list is configured correctly with correct regular expression which allows R2 only to advertise routes that are originated its own AS, AS 65002.
- AS-path filter is applied to the EBGP neighbors in the outgoing direction.

2 Use the **show ip bgp neighbors 192.168.23.3 advertised-routes** command to ensure that R2 only advertises prefixes from its own autonomous system to router R3.

You should get the following results.

```
R2# show ip bgp neighbors 192.168.23.3 advertised-routes
BGP table version is 15, local router ID is 10.1.2.1
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,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.1.2.0/24      0.0.0.0                  0         32768 i

Total number of prefixes 1
```

FEEDBACK: If you do not see the desired results, verify that:

- You cleared the BGP process or wait sufficient time for automatic BGP update.
- AS-path access-list is configured correctly with correct regular expression which allows R2 only to advertise routes that are originated its own AS, AS 65002.
- AS-path filter is applied to the EBGP neighbors in the outgoing direction.

Step 2 Verify router R6 which denied the prefixes that are originated behind AS 65005 to AS 65001.

1 Use the **show ip bgp neighbors 192.168.61.1 advertised-routes** command to ensure that R6 only advertises prefixes from its own autonomous system to router R1 and denies prefixes that are originated behind AS 65005.

You should get the following results.


```
R6# show ip bgp neighbors 192.168.61.1 advertised-routes
BGP table version is 24, local router ID is 10.1.6.1
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,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.1.6.0/24      0.0.0.0                  0         32768 i

Total number of prefixes 1
```

FEEDBACK: If you do not see the desired results, verify that:

- You cleared the BGP process or wait sufficient time for automatic BGP update.
- AS-path access-list is configured correctly with correct regular expression which allows R6 to deny prefixes that are originated behind AS 65005 and permit other prefixes.
- AS-path filter is applied to the AS 65001 EBGP neighbor in the outgoing direction.

Step 3 Verify router R1 which should not receive any prefixes from AS 65002.

1. Use the **show ip bgp** command to ensure that R1 is not receiving any prefixes from AS 65002.

   You should get the following results.

```
R1# show ip bgp
BGP table version is 46, local router ID is 10.1.1.1
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,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.1.1.0/24      0.0.0.0                  0         32768 i
 *>  10.1.6.0/24      192.168.61.6             0             0 65006 i
```

FEEDBACK: If you do not see the desired results, verify that:

- You cleared the BGP process or wait sufficient time for automatic BGP update.
- AS-path access-list is configured correctly with correct regular expression which allows R1 to deny prefixes from AS 65002 and permit other prefixes.
- AS-path filter is applied to the AS 65002 EBGP neighbor in the inbound direction.

Step 4 Verify router R4 which filters prefixes that are based on the prefix length.

1. Use the **show ip bgp** command to ensure that router R4 has filtered the prefixes that are based on the prefix-list configured.

   You should get the following results.

```
R4# show ip bgp
BGP table version is 21, local router ID is 10.1.4.1
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,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.1.1.0/24      192.168.45.5                           0 65005 65006 65001 i
 *>  10.1.2.0/24      192.168.34.3                           0 65003 65002 i
 *>  10.1.3.0/24      192.168.34.3             0             0 65003 i
 *>  10.1.4.0/24      0.0.0.0                  0         32768 i
 *>  10.1.5.0/24      192.168.45.5             0             0 65005 i
 *>  10.1.6.0/24      192.168.45.5                           0 65005 65006 i
 *>  172.16.1.0/25    192.168.34.3             0             0 65003 i
 *>  172.16.5.0/29    192.168.34.3             0             0 65003 i
```

FEEDBACK: If you do not see the desired results, verify that:

- You cleared the BGP process or wait sufficient time for automatic BGP update.
- Prefix-list is configured correctly with correct deny and permit statements which allow R4 to filter prefixes that are based on the prefix length.
- Prefix-list is applied to the AS 65003 EBGP neighbor in the inbound direction.

2 Use the **show run | sec router bgp** command to ensure that router R4 has the BGP ORF receive capability towards AS 65005.

You should get the following results.

```
R4#show run | section router bgp
router bgp 65004
 bgp log-neighbor-changes
 network 10.1.4.0 mask 255.255.255.0
 neighbor 192.168.34.3 remote-as 65003
 neighbor 192.168.34.3 soft-reconfiguration inbound
 neighbor 192.168.34.3 prefix-list PREFIXFILTER in
 neighbor 192.168.45.5 remote-as 65005
 neighbor 192.168.45.5 capability orf prefix-list receive
 neighbor 192.168.45.5 soft-reconfiguration inbound
```

FEEDBACK: If you do not see the desired results, verify that:

- You cleared the BGP process or wait sufficient time for automatic BGP update.
- Prefix-list is configured correctly with correct deny and permit statements which allow R4 to filter prefixes that are based on the prefix length.
- Prefix-list is applied to the AS 65003 EBGP neighbor in the inbound direction.

Step 5 Verify router R5 which filters prefixes that are based on the Outbound Route Filtering technique.

1 Use the **show ip bgp** command to ensure that router R5 has received only 172.16.5.0/29 prefix that is based on the prefix-list and ORF configured.

You should get the following results.


```
R5# show ip bgp
BGP table version is 13, local router ID is 10.1.5.1
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,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.1.1.0/24      192.168.56.6                           0 65006 65001 i
 *>  10.1.5.0/24      0.0.0.0                  0         32768 i
 *>  10.1.6.0/24      192.168.56.6             0             0 65006 i
 *>  172.16.5.0/29    192.168.45.4                           0 65004 65003 i
```

FEEDBACK: If you do not see the desired results, verify that:

- You cleared the BGP process or wait sufficient time for automatic BGP update.
- Prefix-list is configured correctly with correct statement which allows R5 to permit only 172.16.5.0/29 prefix.
- Prefix-list is applied to the AS 65004 EBGP neighbor in the inbound direction.
- AS 65004 EBGP neighbor has the ORF send capability that is configured.

Challenge 4: Configure BGP Route Selection Using BGP Attributes

You work as a network admin and there is a long weekend approaching. The network implementation team has planned for the migration of some of the devices in their lab to provide more critical services to the customers. The monitoring and support team manager, has created service requests to accommodate the change. Below are the tasks details.

- With weight attribute, configure router R1 so that traffic to 10.1.6.0/24 keeps router R2 as the next hop even if someone configures other attributes.
- With the partially transitive attribute, Local Preference, configure router R3 and ensure all networks that are originated from AS65008 have a local preference of 700. This solution should be implemented using an as-path access-list.
- In router R2, link towards AS 65005 has planned for the migration event and they want to have a different path for their 10.1.6.0/24 traffic by prepend AS's 65020 65025 65030. This solution should be implemented using a prefix-list.
- Using MED attribute, configure router R6 so that the traffic for 10.1.1.0/24 network reaches via router R7. This solution should be implemented using a standard access-list.

You are allowed to use only route-maps to complete the lab.

Topology

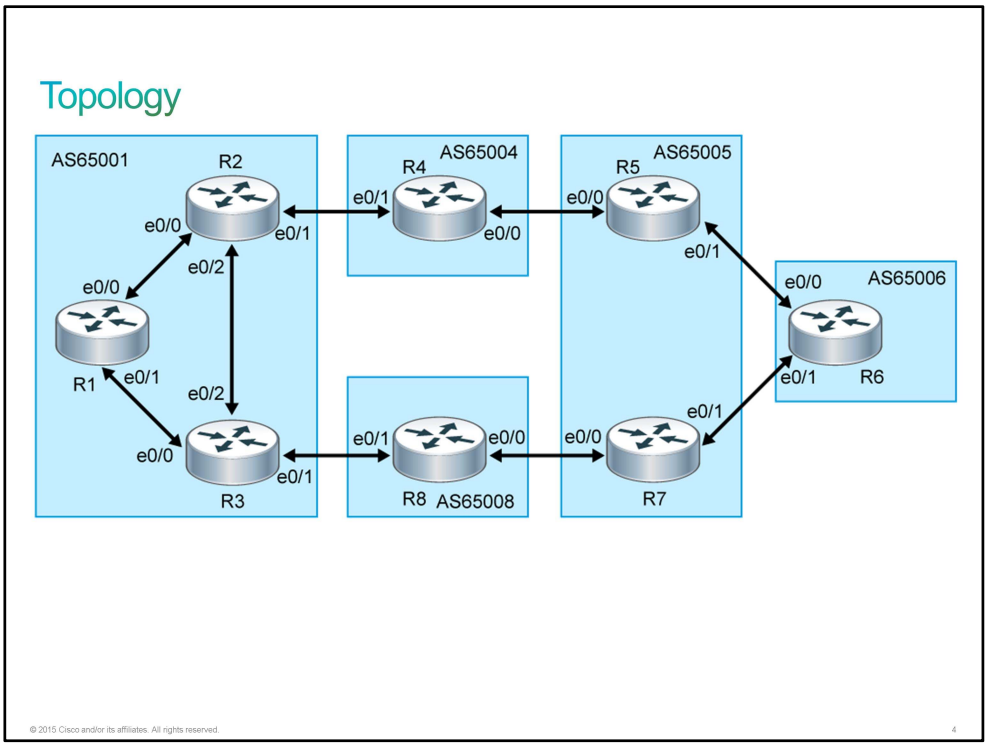

Job Aids

Device Connections

Device	Interface	Connects To	Interface
R1	Ethernet 0/0	R2	Ethernet 0/0
R1	Ethernet 0/1	R3	Ethernet 0/0
R2	Ethernet 0/1	R4	Ethernet 0/1
R2	Ethernet 0/2	R3	Ethernet 0/2
R3	Ethernet 0/1	R8	Ethernet 0/1
R4	Ethernet 0/0	R5	Ethernet 0/0
R5	Ethernet 0/1	R6	Ethernet 0/0
R6	Ethernet 0/1	R7	Ethernet 0/1
R7	Ethernet 0/0	R8	Ethernet 0/0

AS Information

Device	Assigned AS
R1	65001
R2	65001
R3	65001
R4	65004
R5	65005
R6	65006
R7	65005
R8	65008

Device IP Addressing

Device	Interface	IP Address
R1	Loopback0	10.1.1.1/24
R1	Ethernet0/0	192.168.12.1/24
R1	Ethernet0/1	192.168.13.1/24
R2	Ethernet0/0	192.168.12.2/24
R2	Ethernet0/1	192.168.24.2/24
R2	Ethernet0/2	192.168.23.2/24
R3	Ethernet0/0	192.168.13.3/24
R3	Ethernet0/1	192.168.38.3/24
R3	Ethernet0/2	192.168.23.3/24
R4	Loopback0	10.1.4.1/24
R4	Ethernet0/0	192.168.45.4/24
R4	Ethernet0/1	192.168.24.4/24
R5	Ethernet0/0	192.168.45.5/24

Device	Interface	IP Address
R5	Ethernet0/1	192.168.56.5/24
R6	Loopback0	10.1.6.1/24
R6	Ethernet0/0	192.168.56.6/24
R6	Ethernet0/1	192.168.67.6/24
R7	Ethernet0/0	192.168.78.7/24
R7	Ethernet0/1	192.168.67.7/24
R8	Ethernet0/0	192.168.78.8/24
R8	Ethernet0/1	192.168.38.8/24

Route-map Name

Attribute	Route-map Name
Weight	WEIGHT-ATTRIBUTE
Local Preference	LOCALPREF
AS-PATH prepending	ASPATHPREPEND
MED	MED

Parameter and Value Information

Device	Parameter	Name/Value
R1	*weight*	3500
R1	*access-list*	1
R3	*local preference*	700
R3	*as-path access-list*	3
R2	*prefix-list*	LIST
R2	*as-path prepend*	65020 65025 65030

Device	Parameter	Name/Value
R6	*access-list*	6
R6	*metric*	50

Command List

The table describes the commands that are used in this activity. The commands are listed in alphabetical order so that you can easily locate the information that you need. Refer to this list if you need configuration command assistance during the lab activity.

Commands

Command	Description
access-list [*acl-id*] **permit** *ip-address* wildcard-mask	Creates a standard access-list.
clear ip bgp *	Clears the BGP session.
ip as-path access-list *number* **permit** *regexp*	Specifies the AS-path filter.
ip prefix-list *prefix-list name* **permit** *ip-address*	Creates the prefix-list.
match as-path *list-number*	Specifies route-map matching criteria to match an as-path access-list.
match ip address [*acl_id*] [*prefix-list*]	Matches any routes that have a destination network that matches a standard ACL or prefix list.
neighbor *ip-address* **route-map** *map-name* {**in** | **out**}	Applying the route-map to the incoming or outgoing routes.
router bgp *as-number*	Places you in the BGP configuration mode.
route-map *name* {**permit** | **deny**} [*sequence_number*]	Creates the route map entry. Enters route-map configuration mode.
set as-path prepend *as-prepend*	Sets the AS prepend value.
set local-preference *local-preference value*	Sets the local-preference value in route-map.
set metric *metric*	Sets the metric value.
set weight *weight*	Sets the weight in route-map.
show ip bgp	Enables you to inspect the contents of the BGP table.
show ip bgp regexp *regexp*	Enables you to use a regular expression to filter the output of the **show ip bgp** command.

Task 1: Configuring BGP Route Selection Using BGP Attributes

Configuring BGP Route Selection Using BGP Attributes

Step 1 Use standard ip access-list and configure router R1 and make sure traffic to 10.1.6.0/24 keeps router R2 as the next hop.

Step 2 Use the as-path access-list and configure router R3 and ensure all networks that are originated from AS65008 have a local preference value of 700.

Step 3 Use prefix-list and configure router R2 so that traffic to 10.1.6.0/24 traffic will be diverted using AS prepending attribute by adding AS's 65020 65025 65030.

Step 4 With MED attribute, use normal access-list and configure router R6, so that the traffic for 10.1.1.0/24 reaches via router R7.

Verification—for ILT

Use this section to verify your challenge results.

Step 1 Verify that router R1 has set the weight 3500 to 10.1.6.0/24 network and keeps router R2 as the next hop.

1 Use the **show ip bgp** command to ensure that R1 has set a weight value of 3500 for networks that are received from 192.168.12.2.

You should get the following results.

```
R1# show ip bgp
BGP table version is 5, local router ID is 10.1.1.1
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,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.1.1.0/24      0.0.0.0                  0         32768 i
 *>i 10.1.4.0/24      192.168.12.2             0    100      0 65004 i
 * i 10.1.6.0/24      192.168.13.3             0    100      0 65008 65005 65006 i
 *>i                  192.168.12.2             0    100   3500 65004 65005 65006 i
 *>i 10.1.8.0/24      192.168.13.3             0    700      0 65008 i
```

FEEDBACK: If you do not see the desired results, verify that:

- Access-list that is configured is 1 and has permitted the 10.1.6.0/24 network.
- The configuration of Route-map **WEIGHT-ATTRIBUTE** is correct.
- Route-map is applied to the EBGP neighbor 192.168.12.2/24 under the BGP process in the inbound direction.

Step 2 Verify that router R3 has set the local preference value 700 to the networks originated from AS 65008.

1. Use the **show ip bgp** command to ensure that R3 has set the local preference value 700 to the networks originated from AS 65008.

   You should get the following results.

```
R3# show ip bgp
BGP table version is 5, local router ID is 192.168.38.3
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,
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 10.1.1.0/24      192.168.13.1             0    100      0 i
 *>i 10.1.4.0/24      192.168.23.2             0    100      0 65004 i
 * i 10.1.6.0/24      192.168.23.2             0    100      0 65004 65005 65006 i
 *>                   192.168.38.8                          0 65008 65005 65006 i
 *>  10.1.8.0/24      192.168.38.8             0    700      0 65008 i
```

   FEEDBACK: If you do not see the desired results, verify that:

   - AS-PATH access-list that is configured is 3 and has a regular expression for networks that are originated from AS 65008.
   - The configuration of Route-map **LOCALPREF** is correct.
   - Route-map is applied to the EBGP neighbor 192.168.38.8/24 under the BGP process in the inbound direction.

Step 3 Verify that router R2 has diverted the traffic for 10.1.6.0/24 network using AS prepending attribute.

1. Use the **show ip bgp** command to ensure that router R2 has diverted the traffic for 10.1.6.0/24 network using AS prepending attribute.

   You should get the following results.

```
R2# show ip bgp
BGP table version is 5, local router ID is 192.168.24.2
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 10.1.1.0/24      192.168.12.1             0    100      0 i
 *>  10.1.4.0/24      192.168.24.4             0             0 65004 i
 *   10.1.6.0/24      192.168.24.4                           0 65020 65025 65030
65004 65005 65006 i
 *>i                  192.168.23.3             0    100      0 65008 65005 65006 i
 *>i 10.1.8.0/24      192.168.23.3             0    700      0 65008 i
```

   FEEDBACK: If you do not see the desired results, verify that:

   - Prefix-list configured is LIST and has a permit statement for 10.1.6.0/24 network.
   - The configuration of Route-map **ASPATHPREPEND** is correct.
   - Route-map is applied to the EBGP neighbor 192.168.24.4/24 under the BGP process in the inbound direction.

Step 4 Verify that router R6 has altered the traffic for 10.1.1.0/24 network using MED attribute.

1 Use the **show ip bgp** command to ensure that router R6 has altered the traffic for 10.1.1.0/24 network using MED attribute.

You should get the following results.

```
R6# show ip bgp
BGP table version is 26, local router ID is 10.1.6.1
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,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
*    10.1.1.0/24      192.168.56.5             50             0 65005 65004 65001 i
*>                    192.168.67.7                            0 65005 65008 65001 i
*    10.1.4.0/24      192.168.67.7                            0 65005 65008 65001
65004 i
*>                    192.168.56.5                            0 65005 65004 i
*>   10.1.6.0/24      0.0.0.0                  0         32768 i
*>   10.1.8.0/24      192.168.56.5                            0 65005 65004 65001
65008 i
*                     192.168.67.7                            0 65020 65025 65030
65005 65008 i
```

FEEDBACK: If you do not see the desired results, verify that:

- Access-list that is configured is 6 and has a permit statement for 10.1.1.0/24 network.
- The configuration of Route-map **MED** is correct.
- Route-map is applied to the EBGP neighbor 192.168.56.5/24 under the BGP process in the inbound direction.

Challenge 5: Configure BGP Route Reflectors

One of the middle range Service Provider in the city has a good customer base and is expected to increase in the coming years. In order to support their new customers, company has bought many new routers to be included in their core network. As of now, the Network Team has maintained the internal BGP network with the IBGP full mesh sessions between their routers.

To have less IBGP sessions and more centralized control of the network, the migration team has decided to insert Route Reflectors in the Autonomous System. The team got the approval from the management for the new design and you are asked to implement the routers R6, R7 and R8 in router hierarchy mode. Below are the task details:

- Router R6 has routers R2, R3, R7, and R8 as its route reflector clients.
- Routers R7 and R8 have router R4 as a route reflector client.
- Routers R7, R8 belong to a cluster with global cluster ID 124.
- Router R6 belongs to a cluster with global cluster ID 23.
- Use only the physical IP address for IBGP neighborship.
- Use loopback 0 IP address as its router-id.
- All client routers: R2, R3, and R4 and EBGP routers—R1 and R5 should receive all of the advertised routes that are originated through IBGP and EBGP.

Route reflectors R6, R7 and R8 have basic IP addressing and IGP preconfigured. See Job Aids for more information.

All other routers R1, R2, R3, R4 and R5 have IP addressing, IGP, IBGP, and EBGP preconfigured.

Topology

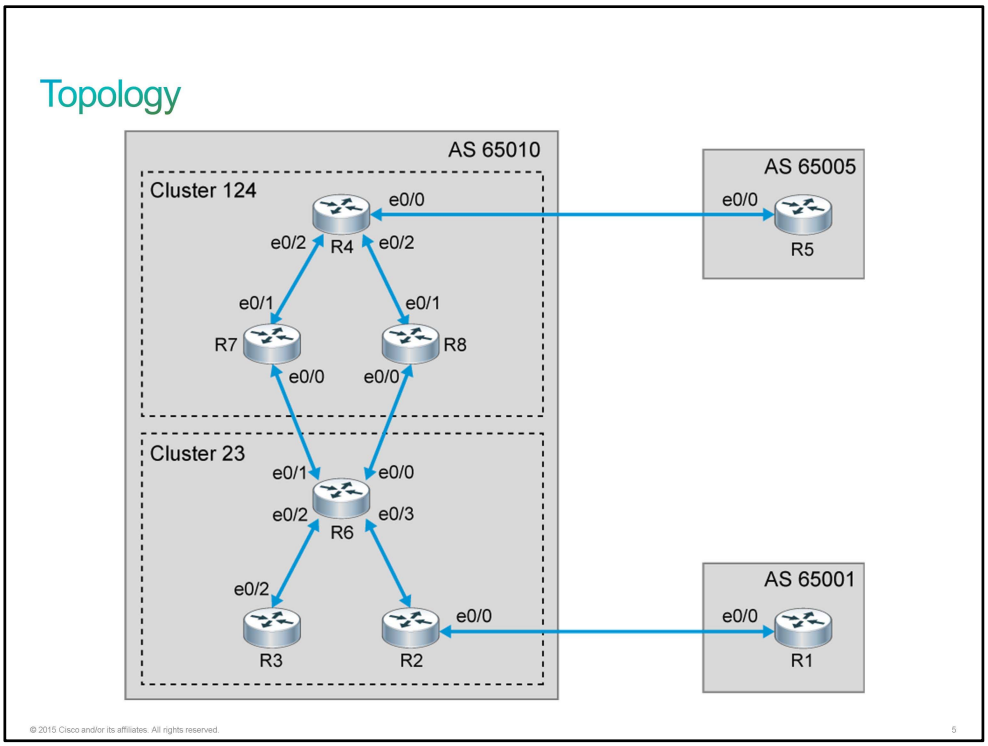

Job Aids

AS Information

Device	Assigned AS
R1	65001
R2	65010
R3	65010
R4	65010
R5	65005
R6	65010
R7	65010
R8	65010

Cluster ID Information

Device	Cluster ID
R6	23
R7	124
R8	124

IP Addressing Information

Device	Interface	IP address
R1	Loopback0	10.1.1.1/24
R1	Ethernet 0/0	192.168.12.1/24
R2	Loopback0	10.1.2.1/24
R2	Ethernet 0/0	192.168.12..2/24
R2	Ethernet 0/1	192.168.26.2/24
R3	Loopback0	10.1.3.1/24
R3	Ethernet 0/0	192.168.36.3/24
R4	Loopback1	10.1.4.1/24
R4	Ethernet 0/0	192.168.45.4/24
R4	Ethernet 0/1	192.168.48.4/24
R5	Loopback0	10.1.5.1/24
R5	Ethernet 0/0	192.168.45.5/24
R6	Loopback0	10.1.6.1/24
R6	Ethernet 0/0	192.168.68.6/24
R6	Ethernet 0/1	192.168.67.6/24
R6	Ethernet 0/2	192.168.36.6/24
R6	Ethernet 0/3	192.168.26.6/24

Device	Interface	IP address
R7	Loopback0	10.1.7.1/24
R7	Ethernet 0/0	192.168.67.7/24
R7	Ethernet 0/1	192.168.47.7/24
R8	Loopback0	10.1.8.1/24
R8	Ethernet 0/0	192.168.68.8/24
R8	Ethernet 0/1	192.168.48.8/24

Command List

The table describes the commands that are used in this activity. The commands are listed in alphabetical order so that you can easily locate the information that you need. Refer to this list if you need configuration command assistance during the lab activity.

Command	Description
bgp cluster-id *id*	Configures an interface type and places you in interface configuration mode
bgp router-id *id*	Configures the router identifier manually
neighbor *ip-address* **remote-as** *as-number*	Establishes a BGP session with the specified AS number
neighbor *ip-address* **route-reflector-client**	Specifies the neighbor as a route-reflector client
router bgp *as-number*	Places you in BGP configuration mode
show ip bgp summary	Verifies whether the BGP session is up
show ip bgp neighbor	Enables you to view detailed information about the neighbor
show ip bgp	Enables you to inspect the contents of the BGP table
show ip bgp *network*	Enables you to view detailed information about prefixes (aggregates)
show ip interface brief	Displays the status of the Interfaces

Task 1: Configuring Route Reflectors

Steps

Step 1 Configure router R7 based on the information from Scenario and Job Aids.

Step 2 Configure router R8 based on the information from Scenario and Job Aids.

Step 3 Configure router R6 based on the information from Scenario and Job Aids.

Verification—for ILT

Use this section to verify your challenge results.

Step 1 Verify that the BGP configuration on router R7 is correct.

1 Use the **show ip bgp summary** command to ensure that R7 has created IBGP sessions with router R6 and router R4.

You should get the following results.

```
R7# show ip bgp summary
BGP router identifier 10.1.7.1, local AS number 65010
BGP table version is 3, main routing table version 3
2 network entries using 296 bytes of memory
2 path entries using 128 bytes of memory
2/2 BGP path/bestpath attribute entries using 272 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 720 total bytes of memory
BGP activity 4/2 prefixes, 4/2 paths, scan interval 60 secs

Neighbor        V           AS   MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.47.4    4        65010     20      17       3     0    0    00:13:10        2
192.168.67.6    4        65010      0       0       1     0    0     never      Idle
```

Item Level Feedback: If you do not see the desired results, verify that

- Router R7 is configured correctly. Take special note of the AS number, bgp router-id, cluster id, and the neighbor commands.
- Make sure, Router R7 is receiving 2 prefixes from router R4.

2 Use the **show ip bgp 10.1.4.1** command to ensure that R7 has received routes from the RR-client which is router R4.

You should get the following results.

```
R7# show ip bgp 10.1.4.1
BGP routing table entry for 10.1.4.0/24, version 2
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 1
  Local, (Received from a RR-client)
    192.168.47.4 from 192.168.47.4 (10.1.4.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
```

Item Level Feedback: If you do not see the desired results, verify that Router R7 is configured with the ***route-reflector client*** command for router R4 neighbor 192.168.47.4/24.

3 Use the **show ip bgp 10.1.5.1** command to ensure that R7 has received routes from the RR-client which is router R4.

You should get the following results.

```
R7# show ip bgp 10.1.5.1
BGP routing table entry for 10.1.5.0/24, version 3
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 1
  65005, (Received from a RR-client)
    192.168.47.4 from 192.168.47.4 (10.1.4.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
```

Item Level Feedback: If you do not see the desired results, verify that Router R7 is configured with the ***route-reflector client*** command for router R4 neighbor 192.168.47.4/24.

Step 2 Verify that the BGP configuration on router R8 is correct.

1. Use the **show ip bgp summary** command to ensure that R8 has created IBGP sessions with router R6 and router R4.

   You should get the following results.

```
R8# show ip bgp summary
BGP router identifier 10.1.8.1, local AS number 65010
BGP table version is 3, main routing table version 3
2 network entries using 296 bytes of memory
2 path entries using 128 bytes of memory
2/2 BGP path/bestpath attribute entries using 272 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 720 total bytes of memory
BGP activity 2/0 prefixes, 2/0 paths, scan interval 60 secs

Neighbor        V           AS   MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.48.4    4        65010      10       5        3     0    0 00:02:56        2
192.168.68.6    4        65010       0       0        1     0    0      never    Idle
R8#
```

   Item Level Feedback: If you do not see the desired results, verify that

   - Router R8 is configured correctly. Take special note of the AS number, bgp router-id, cluster id, and the neighbor commands.
   - Make sure, Router R8 is receiving 2 prefixes from router R4.

2. Use the **show ip bgp 10.1.4.1** command to ensure that R8 has received routes from the RR-client which is router R4.

   You should get the following results.

```
R8# show ip bgp 10.1.4.1
BGP routing table entry for 10.1.4.0/24, version 2
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 2
  Local, (Received from a RR-client)
    192.168.48.4 from 192.168.48.4 (10.1.4.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
```

   Item Level Feedback: If you do not see the desired results, verify that Router R8 is configured with the ***route-reflector client*** command for router R4 neighbor 192.168.48.4/24.

3. Use the **show ip bgp 10.1.5.1** command to ensure that R8 has received routes from the RR-client which is router R4.

   You should get the following results.

```
R8# show ip bgp 10.1.5.1
BGP routing table entry for 10.1.5.0/24, version 3
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 2
  65005, (Received from a RR-client)
    192.168.48.4 from 192.168.48.4 (10.1.4.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
```

Item Level Feedback: If you do not see the desired results, verify that Router R8 is configured with the ***route-reflector client*** command for router R4 neighbor 192.168.48.4/24.

Step 3 Verify that the BGP configuration on router R6 is correct.

1 Use the **show ip bgp summary** command to ensure that RR has created IBGP sessions with routers R2, R3, R7, and R8.

You should get the following results.

```
R6# show ip bgp summary
BGP router identifier 10.1.6.1, local AS number 65010
BGP table version is 6, main routing table version 6
5 network entries using 740 bytes of memory
7 path entries using 448 bytes of memory
3/3 BGP path/bestpath attribute entries using 408 bytes of memory
1 BGP rrinfo entries using 24 bytes of memory
2 BGP AS-PATH entries using 48 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 1668 total bytes of memory
BGP activity 5/0 prefixes, 7/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down
State/PfxRcd
192.168.26.2    4        65010       9      12        6    0    0 00:03:26        2
192.168.36.3    4        65010       8      12        6    0    0 00:03:24        1
192.168.67.7    4        65010       9      12        6    0    0 00:03:31        2
192.168.68.8    4        65010      10      12        6    0    0 00:03:33        2
```

Item Level Feedback: If you do not see the desired results, verify that

- Router R6 is configured correctly. Take special note of the AS number, bgp router-id, cluster id, and the neighbor commands.
- Router R6 is receiving 2 prefixes from router R2.
- Router R6 is receiving 1 prefixes from router R3.
- Router R6 is receiving 2 prefixes from router R7.
- Router R6 is receiving 2 prefixes from router R8.

2 Use the **show ip bgp 10.1.1.1** command to ensure that R6 has received routes from the RR-client which is router R2.

You should get the following results.

```
R6# show ip bgp 10.1.1.1
BGP routing table entry for 10.1.1.0/24, version 2
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  65001, (Received from a RR-client)
    192.168.26.2 from 192.168.26.2 (10.1.2.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
```

Item Level Feedback: If you do not see the desired results, verify that Router R6 is configured with the ***route-reflector client*** command for router R2 neighbor 192.168.26.2/24.

3 Use the **show ip bgp 10.1.2.1** command to ensure that R6 has received routes from the RR-client which is router R2.

You should get the following results.

```
R6# show ip bgp 10.1.2.1
BGP routing table entry for 10.1.2.0/24, version 3
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local, (Received from a RR-client)
    192.168.26.2 from 192.168.26.2 (10.1.2.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
```

Item Level Feedback: If you do not see the desired results, verify that Router R6 is configured with the ***route-reflector client*** command for router R2 neighbor 192.168.26.2/24.

4 Use the **show ip bgp 10.1.3.1** command to ensure that R6 has received routes from the RR-client which is router R3.

You should get the following results.

```
R6# show ip bgp 10.1.3.1
BGP routing table entry for 10.1.3.0/24, version 4
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local, (Received from a RR-client)
    192.168.36.3 from 192.168.36.3 (10.1.3.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
```

Item Level Feedback: If you do not see the desired results, verify that Router R6 is configured with the ***route-reflector client*** command for router R3 neighbor 192.168.36.3/24.

5 Use the **show ip bgp 10.1.4.1** command to ensure that R6 has received routes from the RR-client which is router R7 and R8.

You should get the following results.

```
R6# show ip bgp 10.1.4.1
BGP routing table entry for 10.1.4.0/24, version 5
Paths: (2 available, best #1, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local, (Received from a RR-client)
    192.168.47.4 (metric 20) from 192.168.67.7 (10.1.7.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Originator: 10.1.4.1, Cluster list: 0.0.0.124
  Refresh Epoch 1
  Local, (Received from a RR-client)
    192.168.48.4 (metric 20) from 192.168.68.8 (10.1.8.1)
      Origin IGP, metric 0, localpref 100, valid, internal
      Originator: 10.1.4.1, Cluster list: 0.0.0.124
```

Item Level Feedback: If you do not see the desired results, verify that Router R6 is configured with the ***route-reflector client*** command for routers R7 neighbor

192.168.67.7/24 and R8 neighbor 192.168.68.8/24.

6 Use the **show ip bgp 10.1.5.1** command to ensure that R6 has received routes from the RR-clients which are router R7 and R8.

You should get the following results.

```
R6# show ip bgp 10.1.5.1
BGP routing table entry for 10.1.5.0/24, version 6
Paths: (2 available, best #1, table default)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  65005, (Received from a RR-client)
    192.168.47.4 (metric 20) from 192.168.67.7 (10.1.7.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Originator: 10.1.4.1, Cluster list: 0.0.0.124
  Refresh Epoch 1
  65005, (Received from a RR-client)
    192.168.48.4 (metric 20) from 192.168.68.8 (10.1.8.1)
      Origin IGP, metric 0, localpref 100, valid, internal
      Originator: 10.1.4.1, Cluster list: 0.0.0.124
```

Item Level Feedback: If you do not see the desired results, verify that Router R6 is configured with the ***route-reflector client*** command for routers R7 neighbor 192.168.67.7/24 and R8 neighbor 192.168.68.8/24.

Answer Key

Challenge 1: Configure a Basic BGP Network

Task 1: Configuring a Basic BGP Network

Step 1

On R1, enable the BGP session with R2, and advertise all the networks by entering the following commands.

```
R1(config)# router bgp 65001
R1(config-router)# network 10.1.1.0 mask 255.255.255.0
R1(config-router)# network 10.1.2.0 mask 255.255.255.0
R1(config-router)# network 192.168.11.0 mask 255.255.255.0
R1(config-router)# network 192.168.12.0 mask 255.255.255.0
R1(config-router)# neighbor 192.168.10.2 remote-as 65002
R1(config-router)# neighbor 192.168.10.2 ttl-security hops 1
```

Step 2

On R2, enable the BGP session with R1 and R3. Also, advertise all the networks in the BGP configuration

```
R2(config)# router bgp 65002
R2(config-router)# network 10.1.3.0 mask 255.255.255.0
R2(config-router)# network 10.1.4.0 mask 255.255.255.0
R2(config-router)# network 192.168.21.0 mask 255.255.255.0
R2(config-router)# network 192.168.22.0 mask 255.255.255.0
R2(config-router)# neighbor 192.168.10.1 remote-as 65001
R2(config-router)# neighbor 192.168.10.1 ttl-security hops 1
R2(config-router)# neighbor 192.168.20.3 remote-as 65003
R2(config-router)# neighbor 192.168.20.3 password chall3ng3
```

Step 3 On R3, define the interfaces and enable the BGP session. Also advertise the networks into BGP

Enable the BGP session between with R2, and advertise all the networks. Also make sure that R3 is advertising only one route for the 10.1.5.1/24,10.1.6.1/24 and 10.1.7.1/24 networks with a 16-bit mask

```
R3(config)# router bgp 65003
R3(config-router)# network 10.1.5.0 mask 255.255.255.0
R3(config-router)# network 10.1.6.0 mask 255.255.255.0
R3(config-router)# network 10.1.7.0 mask 255.255.255.0
R3(config-router)# network 192.168.31.0 mask 255.255.255.0
R3(config-router)# network 192.168.32.0 mask 255.255.255.0
R3(config-router)# neighbor 192.168.20.2 remote-as 65002
R3(config-router)# neighbor 192.168.20.2 password ch4ll3ng3
R3(config-router)# aggregate-address 10.1.0.0 255.255.0.0 summary-only
```

Challenge 2: Configure a BGP Transit AS

Task 1: Configure a Transit AS

Step 1 On R2, enable EBGP session with ISP1, and IBGP sessions with R3, and R4.

Enter the following commands on R2.

```
R2(config)# router bgp 65002
R2(config-router)# neighbor 10.1.4.1 remote-as 65002
R2(config-router)# neighbor 10.1.4.1 update-source Loopback0
R2(config-router)# neighbor 10.1.4.1 next-hop-self
R2(config-router)# neighbor 192.168.12.1 remote-as 65001
R2(config-router)# neighbor 192.168.23.3 remote-as 65002
R2(config-router)# neighbor 192.168.23.3 next-hop-self
```

Step 2 On R3, enable EBGP session with ISP2, and IBGP sessions with R2 and R4.

Enter the following commands on R3.

```
R3(config)# router bgp 65002
R3(config-router)# neighbor 192.168.23.2 remote-as 65002
R3(config-router)# neighbor 192.168.23.2 next-hop-self
R3(config-router)# neighbor 192.168.34.4 remote-as 65002
R3(config-router)# neighbor 192.168.35.5 remote-as 65003
```

Step 3 On R4, enable IBGP sessions with R2 and R3

Enter the following commands on R4.

```
R4(config)# router bgp 65002
R4(config-router)# neighbor 10.1.2.1 remote-as 65002
R4(config-router)# neighbor 10.1.2.1 update-source Loopback0
R4(config-router)# neighbor 10.1.2.1 next-hop-self
R4(config-router)# neighbor 192.168.34.3 remote-as 65002
```

Step 4 On ISP1, enable EBGP session with R2 and advertise all the networks into BGP.

Enter the following commands on ISP1.


```
ISP1(config)# router bgp 65001
ISP1(config-router)# network 10.1.1.0 mask 255.255.255.0
ISP1(config-router)# network 192.168.15.0 mask 255.255.255.0
ISP1(config-router)# network 192.168.16.0 mask 255.255.255.0
ISP1(config-router)# neighbor 192.168.12.2 remote-as 65002
ISP1(config-router)# neighbor 192.168.12.2 next-hop-self
```

Step 5 On ISP2, enable EBGP session with R3, and advertise all the networks into BGP.

Enter the following commands on ISP2.

```
ISP2(config)# router bgp 65003
ISP2(config-router)# network 10.1.5.0 mask 255.255.255.0
ISP2(config-router)# network 192.168.55.0 mask 255.255.255.0
ISP2(config-router)# network 192.168.56.0 mask 255.255.255.0
ISP2(config-router)# neighbor 192.168.35.3 remote-as 65002
ISP2(config-router)# neighbor 192.168.35.3 next-hop-self
```

Challenge 3: Configure BGP Using BGP Filtering

Task 1: Configuring BGP Using BGP Filtering Techniques

Step 1 On R2, create an AS-path access-list, which advertises only prefixes originated from AS 65002 and also apply the AS-path filter to the EBGP neighbors.

Enter the following commands on router R2.

```
R2(config)# ip as-path access-list 2 permit ^$
R2(config)#
R2(config)# router bgp 65002
R2(config-router)# neighbor 192.168.12.1 filter-list 2 out
R2(config-router)# neighbor 192.168.23.3 filter-list 2 out
```

Step 2 On R6, create an AS-path access-list, which denies prefixes behind AS 65005 to be advertising it to AS 65001 and also apply the AS-path filter to the EBGP neighbor connected to AS 65001.

Enter the following commands on router R6.

```
R6(config)# ip as-path access-list 6 deny ^65005_
R6(config)# ip as-path access-list 6 permit .*
R6(config)#
R6(config)# router bgp 65006
R6(config-router)# neighbor 192.168.61.1 filter-list 6 out
```

Step 3 On R1, create an AS-path access-list, which will deny prefixes from AS65002 to be advertising it to AS 65001 and also apply the AS-path filter to the R2 EBGP neighbor 192.168.12.2 that is connected to AS 65001.

Enter the following commands on router R1.

```
R1(config)# ip as-path access-list 1 deny _65002$
R1(config)# ip as-path access-list 1 permit .*
R1(config)#
R1(config)# router bgp 65001
R1(config-router)# neighbor 192.168.12.2 filter-list 1 in
```

Step 4 On R4, create a prefix-list which filters all /32 prefixes of 10.0.0.0/8 network, /26 to /28 prefix length of 172.16.0.0/16 network and allow other prefixes coming from AS 65003.

Enter the following commands on router R4.

```
R4(config)# ip prefix-list PREFIXFILTER seq 5 deny 10.0.0.0/8 ge 32
R4(config)# ip prefix-list PREFIXFILTER seq 10 deny 172.16.0.0/16 ge 26 le 28
R4(config)# ip prefix-list PREFIXFILTER seq 15 permit 0.0.0.0/0 le 32
R4(config)#
R4(config)# router bgp 65004
R4(config-router)# neighbor 192.168.34.3 prefix-list PREFIXFILTER in
R4(config-router)# neighbor 192.168.45.5 capability orf prefix-list receive
```

Step 5 On R5, create a prefix-list which permits only 172.16.5.0/29 prefix from AS65004 and also apply the prefix-list to the R4 EBGP neighbor 192.168.45.4 that is connected to AS65004. .

Enter the following commands on router R5.


```
R5(config)# ip prefix-list ORF seq 5 permit 172.16.5.0/29
R5(config)# router bgp 65005
R5(config-router)# neighbor 192.168.45.4 prefix-list ORF in
R5(config-router)# neighbor 192.168.45.4 capability orf prefix-list send
```

Challenge 4: Configure BGP Route Selection Using BGP Attributes

Task 1: Configuring BGP Route Selection Using BGP Attributes

Step 1 On router R1, create a standard access-list and route-map, set the weight and match the network with an access-list created. Also apply the route-map to the EBGP neighbor 192.168.12.2/24 in the Inbound direction.

Enter the following commands on router R1.

```
R1(config)# route-map WEIGHT-ATTRIBUTE permit 10
R1(config-route-map)# match ip address 1
R1(config-route-map)# set weight 3500
R1(config-route-map)# route-map WEIGHT-ATTRIBUTE permit 20
R1(config-route-map)# exit
R1(config)# access-list 1 permit 10.1.6.0 0.0.0.255
R1(config)# router bgp 65001
R1(config-router)# neighbor 192.168.12.2 route-map WEIGHT-ATTRIBUTE in
```

Step 2 On router R3, create an as-path access-list and route-map, match the as-path access-list and set the local preference value. Also, apply the route-map to the EBGP neighbor 192.168.38.8 in the Inbound direction.

Enter the following commands on router R3.

```
R3(config)# route-map LOCALPREF permit 10
R3(config-route-map)# match as-path 3
R3(config-route-map)# set local-preference 700
R3(config-route-map)# route-map LOCALPREF permit 20
R3(config-route-map)# exit
R3(config)# ip as-path access-list 3 permit _65008$
R3(config)# router bgp 65001
R3(config-router)# neighbor 192.168.38.8 route-map LOCALPREF in
```

Step 3 On router R2, create a prefix-list and route-map, match the prefix--list and set the as-path prepend. Also, apply the route-map to the EBGP neighbor 192.168.24.4/24 in the Inbound direction.

Enter the following commands on router R2.

```
R2(config)# route-map ASPATHPREPEND permit 10
R2(config-route-map)# match ip address prefix-list LIST
R2(config-route-map)# set as-path prepend 65020 65025 65030
R2(config-route-map)# route-map ASPATHPREPEND permit 20
R2(config-route-map)# exit
R2(config)# ip prefix-list LIST permit 10.1.6.0/24
R2(config)# router bgp 65001
R2(config-router)# neighbor 192.168.24.4 route-map ASPATHPREPEND in
```

Step 4 On router R6, create an access-list and route-map, match the access--list and set the metric value. Also, apply the route-map to the EBGP neighbor 192.168.56.5/24 in the Inbound direction.

Enter the following commands on router R6.


```
R6(config)# route-map MED permit 10
R6(config-route-map)# match ip address 6
R6(config-route-map)# set metric 50
R6(config-route-map)# route-map MED permit 20
R6(config-route-map)# exit
R6(config)# access-list 6 permit 10.1.1.0 0.0.0.255
R6(config)# router bgp 65006
R6(config-router)# neighbor 192.168.56.5 route-map MED in
```

Challenge 5: Configure BGP Route Reflectors

Task 1: Configuring Route Reflectors

Step 1 On router R7, create IBGP sessions with router R6 and router R4. Also, ensure that router R4 as its router reflector client. Make sure, the bgp router-id is the loopback 0 ip address and the cluster id is configured under the bgp process.

Enter the following commands on router R7.

```
R7(config)# router bgp 65010
R7(config-router)# bgp router-id 10.1.7.1
R7(config-router)# bgp cluster-id 124
R7(config-router)# neighbor 192.168.47.4 remote-as 65010
R7(config-router)# neighbor 192.168.47.4 route-reflector-client
R7(config-router)# neighbor 192.168.67.6 remote-as 65010
R7(config-router)# exit
```

Step 2 On router R8, create IBGP sessions with router R6 and router R4. Also, ensure that router R4 as the router reflector client. Make sure, the bgp router-id is the loopback 0 ip address and the cluster id is configured under the bgp process.

Enter the following commands on router R8.

```
R8(config)# router bgp 65010
R8(config-router)# bgp router-id 10.1.8.1
R8(config-router)# bgp cluster-id 124
R8(config-router)# neighbor 192.168.48.4 remote-as 65010
R8(config-router)# neighbor 192.168.48.4 route-reflector-client
R8(config-router)# neighbor 192.168.68.6 remote-as 65010
R8(config-router)# exit
```

Step 3 On router R6, create IBGP sessions with routers R2, R3, R7, and R8. Also, ensure all these routers as router reflector clients. Make sure, the bgp router-id is the loopback 0 ip address and the cluster id is configured under the bgp process.

Enter the following commands on router RR.

```
RR(config)# router bgp 65010
RR(config-router)# bgp router-id 10.1.6.1
RR(config-router)# bgp cluster-id 23
RR(config-router)# neighbor 192.168.26.2 remote-as 65010
RR(config-router)# neighbor 192.168.26.2 route-reflector-client
RR(config-router)# neighbor 192.168.36.3 remote-as 65010
RR(config-router)# neighbor 192.168.36.3 route-reflector-client
RR(config-router)# neighbor 192.168.67.7 remote-as 65010
RR(config-router)# neighbor 192.168.67.7 route-reflector-client
RR(config-router)# neighbor 192.168.68.8 remote-as 65010
RR(config-router)# neighbor 192.168.68.8 route-reflector-client
RR(config-router)# exit
```

