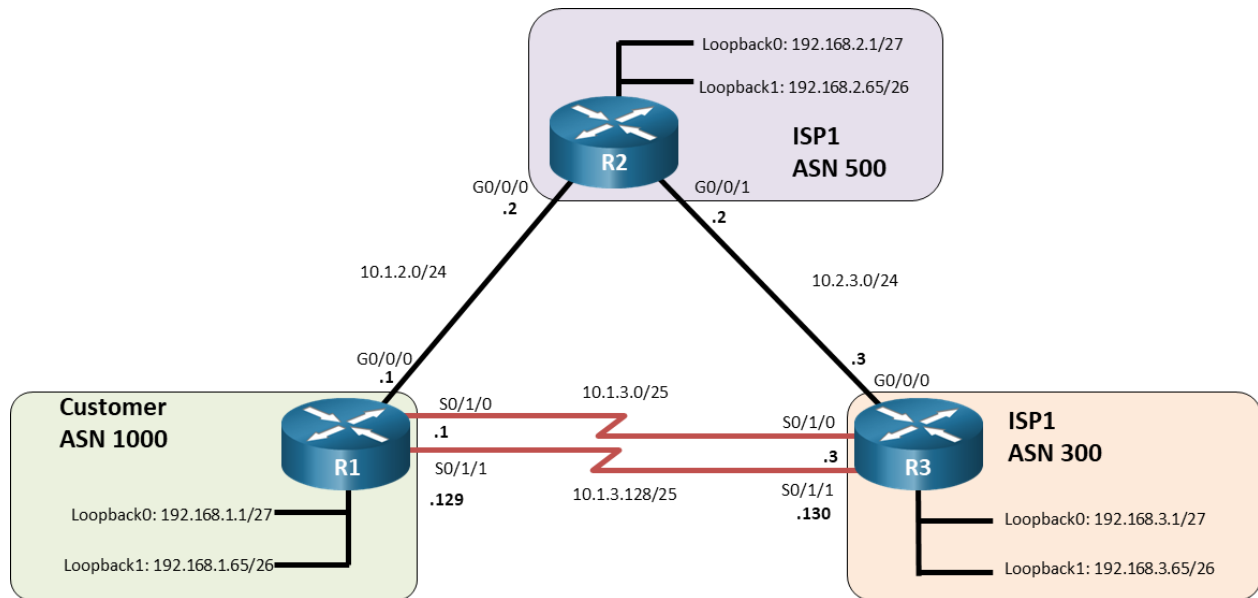


Lab - Implement eBGP for IPv4 (Instructor Version)

Instructor Note: Red font color or gray highlights indicate text that appears in the instructor copy only.

Answers: [11.1.2 Lab - Implement eBGP for IPv4](#)

Topology



Addressing Table

Device	Interface	IPv4 Address
R1	G0/0/0	10.1.2.1/24
	S0/1/0	10.1.3.1/25
	S0/1/1	10.1.3.129/25
	Loopback0	192.168.1.1/27
	Loopback1	192.168.1.65/26
R2	G0/0/0	10.1.2.2/24
	G0/0/1	10.2.3.2/24
	Loopback0	192.168.2.1/27
	Loopback1	192.168.2.65/26
R3	G0/0/0	10.2.3.3/24
	S0/1/0	10.1.3.3/25
	S0/1/1	10.1.3.130/25

Device	Interface	IPv4 Address
	Loopback0	192.168.3.1/27
	Loopback1	192.168.3.65/26

Objectives

Part 1: Build the Network and Configure Basic Device Settings and Interface Addressing

Part 2: Configure and Verify eBGP for IPv4 on all Routers

Part 3: Configure and Verify Route Summarization and Atomic Aggregate

Part 4: Configure and Verify Route Summarization with Atomic Aggregate and AS-Set

Part 5: Configure and Verify the Advertising of a Default Route

Background / Scenario

In this lab you will configure eBGP for IPv4.

Note: This lab is an exercise in developing, deploying, and verifying various path manipulation tools for BGP, and does not reflect networking best practices.

Note: The routers used with CCNP hands-on labs are Cisco 4221 with Cisco IOS XE Release 16.9.4 (universalk9 image). Other routers and Cisco IOS versions can be used. Depending on the model and Cisco IOS version, the commands available and the output produced might vary from what is shown in the labs.

Note: Ensure that the routers and switches have been erased and have no startup configurations. If you are unsure contact your instructor.

Instructor Note: Refer to the Instructor Lab Manual for the procedures to initialize and reload devices.

Required Resources

- 3 Routers (Cisco 4221 with Cisco IOS XE Release 16.9.4 universal image or comparable)
- 1 PC (Windows with a terminal emulation program, such as Tera Term)
- Console cables to configure the Cisco IOS devices via the console ports
- Ethernet and serial cables as shown in the topology

Instructions

Part 1: Build the Network and Configure Basic Device Settings and Interface Addressing

In Part 1, you will set up the network topology and configure basic settings and interface addressing on routers.

Step 1: Cable the network as shown in the topology.

Attach the devices as shown in the topology diagram, and cable as necessary.

Step 2: Configure basic settings for each router.

- a. Console into each router, enter global configuration mode, and apply the basic settings and interface addressing. A command list for each router is provided below.

Router R1

```
hostname R1
no ip domain lookup
line con 0
logging sync
exec-time 0 0
exit
interface Loopback0
 ip address 192.168.1.1 255.255.255.224
 no shut
 exit
interface Loopback1
 ip address 192.168.1.65 255.255.255.192
 no shut
 exit
interface GigabitEthernet0/0/0
 ip address 10.1.2.1 255.255.255.0
 no shut
 exit
interface Serial0/1/0
 ip address 10.1.3.1 255.255.255.128
 no shut
 exit
interface Serial0/1/1
 ip address 10.1.3.129 255.255.255.128
 no shut
 exit
```

Router R2

```
hostname R2
no ip domain lookup
line con 0
logging sync
exec-time 0 0
exit
interface Loopback0
 ip address 192.168.2.1 255.255.255.224
 no shut
 exit
interface Loopback1
 ip address 192.168.2.65 255.255.255.192
 no shut
 exit
interface GigabitEthernet0/0/0
 ip address 10.1.2.2 255.255.255.0
 no shut
 exit
```

```
interface GigabitEthernet0/0/1
 ip address 10.2.3.2 255.255.255.0
 no shut
 exit
```

Router R3

```
hostname R3
no ip domain lookup
line con 0
 logging sync
 exec-time 0 0
 exit
interface Loopback0
 ip address 192.168.3.1 255.255.255.224
 no shut
 exit
interface Loopback1
 ip address 192.168.3.65 255.255.255.192
 no shut
 exit
interface GigabitEthernet0/0/0
 ip address 10.2.3.3 255.255.255.0
 negotiation auto
 no shut
 exit
interface Serial0/1/0
 ip address 10.1.3.3 255.255.255.128
 no shut
 exit
interface Serial0/1/1
 ip address 10.1.3.130 255.255.255.128
 no shut
 exit
```

- b. Save the running configuration to startup-config.

Part 2: Configure and Verify eBGP for IPv4 on all Routers

Step 1: Implement BGP and neighbor relationships on R1.

- a. Enter BGP configuration mode from global configuration mode, specifying AS 1000.
R1(config)# **router bgp 1000**
- b. Configure the BGP router-id for R1.
R1(config-router)# **bgp router-id 1.1.1.1**
- c. Based on the topology diagram, configure all the designated neighbors for R1.
R1(config-router)# **neighbor 10.1.2.2 remote-as 500**
R1(config-router)# **neighbor 10.1.3.3 remote-as 300**

```
R1(config-router)# neighbor 10.1.3.130 remote-as 300
```

- d. Configure R1 to advertise the IPv4 prefixes local to ASN 1000.

```
R1(config-router)# network 192.168.1.0 mask 255.255.255.224
```

```
R1(config-router)# network 192.168.1.64 mask 255.255.255.192
```

Step 2: Implement BGP and neighbor relationships on R2.

- a. Enter BGP configuration mode from global configuration mode, specifying AS 500.

```
R2(config)# router bgp 500
```

- b. Configure the BGP router-id for R2.

```
R2(config-router)# bgp router-id 2.2.2.2
```

- c. Based on the topology diagram, configure all the designated neighbors for R2.

```
R2(config-router)# neighbor 10.1.2.1 remote-as 1000
```

```
R2(config-router)# neighbor 10.2.3.3 remote-as 300
```

- d. Configure R2 to advertise the IPv4 prefixes local to ASN 500.

```
R2(config-router)# network 192.168.2.0 mask 255.255.255.224
```

```
R2(config-router)# network 192.168.2.64 mask 255.255.255.192
```

Step 3: Implement BGP and neighbor relationships on R3.

- a. Enter BGP configuration mode from global configuration mode, specifying AS 300.

```
R3(config)# router bgp 300
```

- b. Configure the BGP router-id for R3.

```
R3(config-router)# bgp router-id 3.3.3.3
```

- c. Unlike the configuration on R1 and R2, disable the default IPv4 unicast behavior.

```
R3(config-router)# no bgp default ipv4-unicast
```

The default behavior in IOS is **bgp default ipv4-unicast**. Routers R1 and R2 were configured using this default behavior. The **bgp default ipv4-unicast** command enables the automatic exchange of IPv4 address family prefixes. When this command is disabled using **no bgp default ipv4-unicast**, bgp neighbors must be activated within IPv4 address family (AF) configuration mode. BGP **network** commands must also be configured within IPv4 AF mode.

- d. Based on the topology diagram, configure all the designated neighbors for R3.

```
R3(config-router)# neighbor 10.2.3.2 remote-as 500
```

```
R3(config-router)# neighbor 10.1.3.1 remote-as 1000
```

```
R3(config-router)# neighbor 10.1.3.129 remote-as 1000
```

Step 4: Verifying BGP neighbor relationships.

- a. Examine the routing tables on each router. Notice that R1 and R2 are receiving BGP prefixes from each other but not receiving BGP prefixes from R3. And R3 is not receiving any prefixes from R1 or R2. This is because R3 was configured using **no bgp default ipv4-unicast** and the interfaces must be activated within IPv4 address configuration mode.

```
R1# show ip route bgp | begin Gateway
```

```
Gateway of last resort is not set
```

```
192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
```

```
B      192.168.2.0/27 [20/0] via 10.1.2.2, 00:28:40
B      192.168.2.64/26 [20/0] via 10.1.2.2, 00:28:40
```

```
R2# show ip route bgp | begin Gateway
```

```
Gateway of last resort is not set
```

```
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
B      192.168.1.0/27 [20/0] via 10.1.2.1, 00:29:41
B      192.168.1.64/26 [20/0] via 10.1.2.1, 00:29:41
```

```
R3# show ip route bgp | begin Gateway
```

```
Gateway of last resort is not set
```

- b. This can be further verified by examining the BGP neighbor adjacencies on R2. Notice the BGP state between R2 and R1 is **established**, while the BGP state between R2 and R3 is **idle**.

```
R2# show ip bgp neighbors
```

```
BGP neighbor is 10.1.2.1, remote AS 1000, external link
```

```
BGP version 4, remote router ID 1.1.1.1
```

```
BGP state = Established, up for 00:35:34
```

```
Last read 00:00:28, last write 00:00:35, hold time is 180, keepalive interval is 60 seconds
```

```
Neighbor sessions:
```

```
1 active, is not multisession capable (disabled)
```

```
<output omitted>
```

```
BGP neighbor is 10.2.3.3, remote AS 300, external link
```

```
BGP version 4, remote router ID 0.0.0.0
```

```
BGP state = Idle, down for never
```

```
Neighbor sessions:
```

```
0 active, is not multisession capable (disabled)
```

```
<output omitted>
```

- c. The interfaces on R3 need to be activated in IPv4 AF configuration mode. The **neighbor activate** command in IPv4 AF configuration mode is required to enable the exchange of BGP information between neighbors. This will enable R3 to form an established neighbor adjacency with both R1 and R2. Additionally, because **bgp default ipv4-unicast** is disabled, **network** commands must be configured in IPv4 AF configuration mode.

```
R3(config-router)# address-family ipv4
```

```
R3(config-router-af)# neighbor 10.1.3.1 activate
```

```
R3(config-router-af)# neighbor 10.1.3.129 activate
```

```
R3(config-router-af)# neighbor 10.2.3.2 activate
```

```
R3(config-router-af)# network 192.168.3.0 mask 255.255.255.224
```

```
R3(config-router-af)# network 192.168.3.64 mask 255.255.255.192
```

- d. Verify that all BGP speakers are receiving prefixes from their neighbors. The prefixes from R3 are highlighted in the routing tables of R1 and R2.

Note: The prefixes in the lab are for example purposes only. Most service providers do not accept prefixes larger than /24 for IPv4 (/25 through /32).

```
R1# show ip route bgp | begin Gateway
```

```
Gateway of last resort is not set
```

```
192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
B      192.168.2.0/27 [20/0] via 10.1.2.2, 00:51:09
B      192.168.2.64/26 [20/0] via 10.1.2.2, 00:51:09
192.168.3.0/24 is variably subnetted, 2 subnets, 2 masks
B      192.168.3.0/27 [20/0] via 10.1.3.3, 00:01:43
B      192.168.3.64/26 [20/0] via 10.1.3.3, 00:01:43
```

R2# **show ip route bgp | begin Gateway**

Gateway of last resort is not set

```
192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
B      192.168.1.0/27 [20/0] via 10.1.2.1, 00:51:17
B      192.168.1.64/26 [20/0] via 10.1.2.1, 00:51:17
192.168.3.0/24 is variably subnetted, 2 subnets, 2 masks
B      192.168.3.0/27 [20/0] via 10.2.3.3, 00:01:51
B      192.168.3.64/26 [20/0] via 10.2.3.3, 00:01:51
```

R3# **show ip route bgp | begin Gateway**

Gateway of last resort is not set

```
192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
B      192.168.1.0/27 [20/0] via 10.1.3.1, 00:02:11
B      192.168.1.64/26 [20/0] via 10.1.3.1, 00:02:11
192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
B      192.168.2.0/27 [20/0] via 10.2.3.2, 00:02:11
B      192.168.2.64/26 [20/0] via 10.2.3.2, 00:02:11
```

- e. Verify that the BGP state between R2 and R3 has now been **established**.

R2# **show ip bgp neighbors | begin BGP neighbor is 10.2.3.3**

BGP neighbor is 10.2.3.3, remote AS 300, external link

BGP version 4, **remote router ID 3.3.3.3**

BGP state = Established, up for 00:12:16

Last read 00:00:37, last write 00:00:52, hold time is 180, keepalive interval is 60 seconds

Neighbor sessions:

1 active, is not multisession capable (disabled)

<output omitted>

Step 5: Examining the running-configs.

Examine the running-configs on all three routers. Because router R3 was configured using **no bgp default ipv4-unicast** command, notice that the network commands were automatically entered under the IPv4 AF. This is the same configuration mode where the neighbors were activated to exchange BGP information.

R1# **show running-config | section bgp**

router bgp 1000

bgp router-id 1.1.1.1

```
bgp log-neighbor-changes
network 192.168.1.0 mask 255.255.255.224
network 192.168.1.64 mask 255.255.255.192
neighbor 10.1.2.2 remote-as 500
neighbor 10.1.3.3 remote-as 300
neighbor 10.1.3.130 remote-as 300
```

```
R2# show running-config | section bgp
router bgp 500
  bgp router-id 2.2.2.2
  bgp log-neighbor-changes
  network 192.168.2.0 mask 255.255.255.224
  network 192.168.2.64 mask 255.255.255.192
  neighbor 10.1.2.1 remote-as 1000
  neighbor 10.2.3.3 remote-as 300
```

```
R3# show running-config | section bgp
router bgp 300
  bgp log-neighbor-changes
  no bgp default ipv4-unicast
  neighbor 10.1.3.1 remote-as 1000
  neighbor 10.1.3.129 remote-as 1000
  neighbor 10.2.3.2 remote-as 500
  !
  address-family ipv4
    network 192.168.3.0 mask 255.255.255.224
    network 192.168.3.64 mask 255.255.255.192
    neighbor 10.1.3.1 activate
    neighbor 10.1.3.129 activate
    neighbor 10.2.3.2 activate
  exit-address-family
```

Step 6: Verifying BGP operations.

- a. To verify the BGP operation on R2, issue the **show ip bgp** command.

```
R2# show ip bgp
BGP table version is 11, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
               t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*	192.168.1.0/27	10.2.3.3			0 300	1000 i
*>		10.1.2.1	0		0 1000	i


```
*      192.168.1.64/26  10.2.3.3                0 300 1000 i
*>      10.1.2.1                0          0 1000 i
*>      192.168.2.0/27  0.0.0.0                0      32768 i
*>      192.168.2.64/26  0.0.0.0                0      32768 i
*>      192.168.3.0/27  10.2.3.3                0          0 300 i
*      10.1.2.1                0 1000 300 i
*>      192.168.3.64/26  10.2.3.3                0          0 300 i
*      10.1.2.1                0 1000 300 i
```

What does the * at the beginning of an entry indicate?

The entry is valid, reachable.

What does the angle bracket (>) in an entry indicate?

An angle bracket (>) indicates that the route has been selected as the best route.

What is the address of the preferred next hop router to reach the 192.168.1.0/27 network? Explain.

10.1.2.1. It is preferred because it has shorter AS path.

How can you verify that 10.1.2.1 is the next hop router used to reach 192.168.1.0/27?

Display the routing table on R2.

B 192.168.1.0/27 [20/0] via 10.1.2.1, 17:31:46

What does a next hop of 0.0.0.0 indicate?

This advertised prefix is a directly connected network.

- b. Use the **show ip bgp ip-prefix** command to display all the paths for a specific route and the BGP path attributes for that route.

```
R2# show ip bgp 192.168.1.0
```

```
BGP routing table entry for 192.168.1.0/27, version 14
```

```
Paths: (2 available, best #2, table default)
```

```
Advertised to update-groups:
```

```
1
```

```
Refresh Epoch 1
```

```
300 1000
```

```
10.2.3.3 from 10.2.3.3 (3.3.3.3)
```

```
Origin IGP, localpref 100, valid, external
```

```
rx pathid: 0, tx pathid: 0
```

```
Refresh Epoch 2
```

```
1000
```

```
10.1.2.1 from 10.1.2.1 (1.1.1.1)
```

```
Origin IGP, metric 0, localpref 100, valid, external, best
rx pathid: 0, tx pathid: 0x0
```

What is the IPv4 address of the next hop router with the best path?

10.1.2.1

- c. Examine the BGP neighbor relationships on R2 using the **show ip bgp neighbors** command.

```
R2# show ip bgp neighbors
```

```
BGP neighbor is 10.1.2.1, remote AS 1000, external link
  BGP version 4, remote router ID 1.1.1.1
  BGP state = Established, up for 00:00:51
  Last read 00:00:00, last write 00:00:51, hold time is 180, keepalive interval is 60
seconds
```

```
Neighbor sessions:
```

```
  1 active, is not multisession capable (disabled)
```

```
Neighbor capabilities:
```

```
  Route refresh: advertised and received(new)
```

```
  Four-octets ASN Capability: advertised and received
```

```
  Address family IPv4 Unicast: advertised and received
```

```
  Enhanced Refresh Capability: advertised and received
```

```
  Multisession Capability:
```

```
  Stateful switchover support enabled: NO for session 1
```

```
Message statistics:
```

```
  InQ depth is 0
```

```
  OutQ depth is 0
```

	Sent	Rcvd
Opens:	1	1
Notifications:	0	0
Updates:	5	5
Keepalives:	2	3
Route Refresh:	0	0
Total:	10	11

```
<output omitted>
```

```
BGP neighbor is 10.2.3.3, remote AS 300, external link
```

```
  BGP version 4, remote router ID 3.3.3.3
```

```
  BGP state = Established, up for 16:23:45
```

```
  Last read 00:00:29, last write 00:00:51, hold time is 180, keepalive interval is 60
seconds
```

```
Neighbor sessions:
```

```
  1 active, is not multisession capable (disabled)
```

```
Neighbor capabilities:
```

```
  Route refresh: advertised and received(new)
```

```
  Four-octets ASN Capability: advertised and received
```

```
  Address family IPv4 Unicast: advertised and received
```

```
  Enhanced Refresh Capability: advertised and received
```

```
  Multisession Capability:
```

```
Stateful switchover support enabled: NO for session 1
Message statistics:
  InQ depth is 0
  OutQ depth is 0

              Sent          Rcvd
Opens:              1            1
Notifications:      0            0
Updates:            9            5
Keepalives:        1082         1088
Route Refresh:       0            0
Total:             1096         1096
Do log neighbor state changes (via global configuration)
Default minimum time between advertisement runs is 30 seconds
<output omitted>
```

How many neighbors does R2 have and what are their router IDs?

Two neighbors: 1.1.1.1 and 3.3.3.3

What is the BGP state of both neighbors?

established

What are the keepalive and hold time value for both neighbors?

Keepalive is 60 seconds, hold time is 180 seconds

Part 3: Configure and Verify Route Summarization and Atomic Aggregate

Step 1: Configure route summarization using atomic aggregate.

Summarizing prefixes conserves router resources and accelerates best-path calculation by reducing the size of the table. Summarization can be configured either for prefixes originated by the AS or prefixes received from downstream providers. Summarization also provides the benefits of stability by hiding flapping routes or having to install new prefixes when they are contained within a summary.

Although AS 1000 only has two prefixes 192.168.1.0/27 and 192.168.1.64/26, this customer has been allocated the entire 192.168.1.0/24 prefix. R3 in AS 300 has two prefixes 192.168.3.0/27 and 192.168.3.64/26 but has been allocated the entire 192.168.3.0/24 prefix.

Configure R1 and R3 to advertise a summary or aggregate route using the **aggregate-address** command. The **summary-only** option suppresses the specific prefixes that are summarized from also being advertised. Notice that this command is configured in **address-family ipv4** configuration mode on R3.

```
R1(config)# router bgp 1000
R1(config-router)# aggregate-address 192.168.1.0 255.255.255.0 summary-only

R3(config)# router bgp 300
R3(config-router)# address-family ipv4
```

```
R3(config-router-af)# aggregate-address 192.168.3.0 255.255.255.0 summary-only
```

Step 2: Verify route summarization using atomic aggregate.

- a. Examine the routing tables on each router to verify the route summarization for the two prefixes. Verify that R1 and R3 are each receiving the summary route from the other router. Verify that R2 is receiving aggregate routes from both R1 and R3.

```
R1# show ip route bgp | begin Gateway
```

```
Gateway of last resort is not set
```

```
      192.168.1.0/24 is variably subnetted, 5 subnets, 4 masks
B       192.168.1.0/24 [200/0], 00:27:47, Null0
      192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
B       192.168.2.0/27 [20/0] via 10.1.2.2, 13:34:31
B       192.168.2.64/26 [20/0] via 10.1.2.2, 13:34:31
B       192.168.3.0/24 [20/0] via 10.1.3.3, 00:26:01
```

```
R2# show ip route bgp | begin Gateway
```

```
Gateway of last resort is not set
```

```
B       192.168.1.0/24 [20/0] via 10.1.2.1, 00:33:53
B       192.168.3.0/24 [20/0] via 10.2.3.3, 00:32:08
```

```
R3# show ip route bgp | begin Gateway
```

```
Gateway of last resort is not set
```

```
B       192.168.1.0/24 [20/0] via 10.1.3.1, 00:36:52
      192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
B       192.168.2.0/27 [20/0] via 10.2.3.2, 02:10:48
B       192.168.2.64/26 [20/0] via 10.2.3.2, 02:10:48
      192.168.3.0/24 is variably subnetted, 5 subnets, 4 masks
B       192.168.3.0/24 [200/0], 00:35:07, Null0
```

Why do R1 and R3 contain an entry with a next hop address of Null0? What is the result of having this Null0 route in the routing table?

This is the same prefix as the aggregate route advertised by this router. The router's routing table contains more specific routes. The Null0 route is a discard route for loop prevention. Any packets with a destination IPv4 address matching the aggregated prefix with the Null0 egress interface, but do not match a more specific route, will be dropped.

- b. Examine the BGP table on router R2 to verify the route summarization. When a prefix has the default classful mask, the subnet mask is not displayed. Both 192.168.1.0 and 192.168.3.0 prefixes have a /24 prefix length which would be the default mask for a Class C address.

```
R2# show ip bgp
```

```
BGP table version is 69, local router ID is 1.1.1.1
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
<output omitted>
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*	192.168.1.0	10.2.3.3				0 300 1000 i
*>		10.1.2.1	0			0 1000 i
*>	192.168.2.0/27	0.0.0.0	0		32768	i
*>	192.168.2.64/26	0.0.0.0	0		32768	i
*	192.168.3.0	10.1.2.1				0 1000 300 i
*>		10.2.3.3	0			0 300 i

- c. Examine the BGP table on routers R2 and R3 and verify that each router is receiving the summary route from the other router.

R1# **show ip bgp**

BGP table version is 69, local router ID is 1.1.1.1

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
<output omitted>

	Network	Next Hop	Metric	LocPrf	Weight	Path
s>	192.168.1.0/27	0.0.0.0	0		32768	i
*>	192.168.1.0	0.0.0.0				32768 i
s>	192.168.1.64/26	0.0.0.0	0		32768	i
*	192.168.2.0/27	10.1.3.130				0 300 500 i
*		10.1.3.3				0 300 500 i
*>		10.1.2.2	0			0 500 i
*	192.168.2.64/26	10.1.3.130				0 300 500 i
*		10.1.3.3				0 300 500 i
*>		10.1.2.2	0			0 500 i
*	192.168.3.0	10.1.2.2				0 500 300 i
*		10.1.3.130	0			0 300 i
*>		10.1.3.3	0			0 300 i

R3# **show ip bgp**

BGP table version is 22, local router ID is 3.3.3.3

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
x best-external, a additional-path, c RIB-compressed,
t secondary path, L long-lived-stale,

Origin codes: i - IGP, e - EGP, ? - incomplete

RPKI validation codes: V valid, I invalid, N Not found

	Network	Next Hop	Metric	LocPrf	Weight	Path
*	192.168.1.0	10.2.3.2				0 500 1000 i
*>		10.1.3.1	0			0 1000 i
*		10.1.3.129	0			0 1000 i
*	192.168.2.0/27	10.1.3.1				0 1000 500 i
*		10.1.3.129				0 1000 500 i
*>		10.2.3.2	0			0 500 i
*	192.168.2.64/26	10.1.3.1				0 1000 500 i
*		10.1.3.129				0 1000 500 i
*>		10.2.3.2	0			0 500 i
s>	192.168.3.0/27	0.0.0.0	0		32768	i

```
*> 192.168.3.0 0.0.0.0 32768 i
s> 192.168.3.64/26 0.0.0.0 0 32768 i
```

Why do two of the entries have the status code of "s"? Specifically, this is the result of what command or option that was configured on these two routers?

This indicates that the route was suppressed. These entries were suppressed because a summary route ("*" valid) was advertised instead. Prefixes with a next hop of 0.0.0.0 are directly connected networks on this router. The summary-only option in the aggregate-address command suppressed these more specific prefixes from being advertised.

- d. Examine the explicit 192.168.1.0 prefix entry in R2's BGP table. The route's NLRI information indicates that the route was aggregated in AS 1000 by router with the RID 1.1.1.1.

```
R2# show ip bgp 192.168.1.0
BGP routing table entry for 192.168.1.0/24, version 45
Paths: (2 available, best #2, table default)
  Advertised to update-groups:
    1
  Refresh Epoch 1
  300 1000, (aggregated by 1000 1.1.1.1)
    10.2.3.3 from 10.2.3.3 (3.3.3.3)
      Origin IGP, localpref 100, valid, external, atomic-aggregate
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 2
  1000, (aggregated by 1000 1.1.1.1)
    10.1.2.1 from 10.1.2.1 (1.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, external, atomic-aggregate, best
      rx pathid: 0, tx pathid: 0x0
```

Part 4: Configure and Verify Route Summarization with Atomic Aggregate and AS-Set

Step 1: Configure route summarization using atomic aggregate and AS-Set.

- a. Shut down both serial interfaces on R1. This will create a single path from R1 (AS 1000) to R2 (AS 500) to R3 (AS 300).

```
R1(config)# interface s0/1/0
R1(config-if)# shutdown
R1(config-if)# exit
R1(config)# interface s0/1/1
R1(config-if)# shutdown
```

- b. Remove route aggregation previously configured on R1.

```
R1(config)# router bgp 1000
R1(config-router)# no aggregate-address 192.168.1.0 255.255.255.0 summary-only
```

- c. Verify that R3 is now receiving the non-summarized prefixes 192.168.1.0/27 and 192.168.1.64/26.

```
R3# show ip route 192.168.1.0
Routing entry for 192.168.1.0/24, 2 known subnets
  Variably subnetted with 2 masks
```

```
B      192.168.1.0/27 [20/0] via 10.2.3.2, 00:01:26
B      192.168.1.64/26 [20/0] via 10.2.3.2, 00:01:26
```

- d. On R2, summarize the prefixes 192.168.1.0/27 and 192.168.1.64/26 received from R1 as 192.168.1.0/24.

```
R2(config)# router bgp 500
```

```
R2(config-router)# aggregate-address 192.168.1.0 255.255.255.0 summary-only
```

Step 2: Verify route summarization using atomic aggregate and AS-Set.

- a. Verify that R3 is receiving the aggregated prefix 192.168.1.0/24.

```
R3# show ip route bgp | begin Gateway
```

```
Gateway of last resort is not set
```

```
B      192.168.1.0/24 [20/0] via 10.2.3.2, 00:00:51
      192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
B      192.168.2.0/27 [20/0] via 10.2.3.2, 08:46:37
B      192.168.2.64/26 [20/0] via 10.2.3.2, 08:46:37
      192.168.3.0/24 is variably subnetted, 5 subnets, 4 masks
B      192.168.3.0/24 [200/0], 08:46:07, Null0
```

- b. Examine R3's BGP table. Notice that the AS path only includes the AS that summarized the route, AS 500, router R2.

```
R3# show ip bgp
```

```
<output omitted>
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	192.168.1.0	10.2.3.2	0		0	500 i
*>	192.168.2.0/27	10.2.3.2	0		0	500 i
*>	192.168.2.64/26	10.2.3.2	0		0	500 i
s>	192.168.3.0/27	0.0.0.0	0		32768	i
*>	192.168.3.0	0.0.0.0			32768	i
s>	192.168.3.64/26	0.0.0.0	0		32768	i

- c. On R2, remove the current route aggregation for the 192.168.1.0/24 prefix and configure it again, this time using the **as-set** option.

```
R2(config)# router bgp 500
```

```
R2(config-router)# no aggregate-address 192.168.1.0 255.255.255.0 summary-only
```

```
R2(config-router)# aggregate-address 192.168.1.0 255.255.255.0 as-set summary-only
```

- d. Verify that R3 is receiving the aggregated prefix 192.168.1.0/24.

```
R3# show ip route bgp | begin Gateway
```

```
Gateway of last resort is not set
```

```
B      192.168.1.0/24 [20/0] via 10.2.3.2, 00:01:35
      192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
B      192.168.2.0/27 [20/0] via 10.2.3.2, 08:50:02
B      192.168.2.64/26 [20/0] via 10.2.3.2, 08:50:02
      192.168.3.0/24 is variably subnetted, 5 subnets, 4 masks
B      192.168.3.0/24 [200/0], 08:49:32, Null0
```

- e. Examine R3's BGP table again. Notice that the entry for 192.168.1.0 this time includes the entire AS path. The output from the **show ip bgp 192.168.1.0** command displays both AS numbers and identifies that R2 (2.2.2.2) aggregated the route.

R3# **show ip bgp**

<output omitted>

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	192.168.1.0	10.2.3.2	0	0	500 1000	i
*>	192.168.2.0/27	10.2.3.2	0	0	500	i
*>	192.168.2.64/26	10.2.3.2	0	0	500	i
s>	192.168.3.0/27	0.0.0.0	0	32768		i
*>	192.168.3.0	0.0.0.0		32768		i
s>	192.168.3.64/26	0.0.0.0	0	32768		i

R3# **show ip bgp 192.168.1.0 | begin Refresh**

Refresh Epoch 7

500 1000, (aggregated by 500 2.2.2.2)

10.2.3.2 from 10.2.3.2 (2.2.2.2)

Origin IGP, metric 0, localpref 100, valid, external, best

rx pathid: 0, tx pathid: 0x0

Part 5: Configure and Verify the Advertising of a Default Route

Step 1: Configure default route advertisement on R2.

Configure R2 to advertise a default route to R1. R2 does not necessarily have to have a default route of its own. Core internet routers that have full internet routing tables and do not require a default route are referred to as being in a default-free zone (DFZ).

R2(config)# **router bgp 500**

R2(config-router)# **neighbor 10.1.2.1 default-originate**

Step 2: Verify default route advertisement on R1.

- Examine R1's routing table to verify that it has received a default route.

R1# **show ip route bgp | begin Gateway**

Gateway of last resort is 10.1.2.2 to network 0.0.0.0

B* 0.0.0.0/0 [20/0] via 10.1.2.2, 00:00:37

192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks

B 192.168.2.0/27 [20/0] via 10.1.2.2, 21:24:43

B 192.168.2.64/26 [20/0] via 10.1.2.2, 21:24:43

B 192.168.3.0/24 [20/0] via 10.1.2.2, 12:41:58

- Examine R1's BGP table to verify that it has received a default route.

R1# **show ip bgp**

<output omitted>

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	0.0.0.0	10.1.2.2		0	500	i
*>	192.168.1.0/27	0.0.0.0	0	32768		i
*>	192.168.1.64/26	0.0.0.0	0	32768		i
*>	192.168.2.0/27	10.1.2.2	0	0	500	i
*>	192.168.2.64/26	10.1.2.2	0	0	500	i
*>	192.168.3.0	10.1.2.2		0	500 300	i

Router Interface Summary Table

Router Model	Ethernet Interface #1	Ethernet Interface #2	Serial Interface #1	Serial Interface #2
1800	Fast Ethernet 0/0 (F0/0)	Fast Ethernet 0/1 (F0/1)	Serial 0/0/0 (S0/0/0)	Serial 0/0/1 (S0/0/1)
1900	Gigabit Ethernet 0/0 (G0/0)	Gigabit Ethernet 0/1 (G0/1)	Serial 0/0/0 (S0/0/0)	Serial 0/0/1 (S0/0/1)
2801	Fast Ethernet 0/0 (F0/0)	Fast Ethernet 0/1 (F0/1)	Serial 0/1/0 (S0/1/0)	Serial 0/1/1 (S0/1/1)
2811	Fast Ethernet 0/0 (F0/0)	Fast Ethernet 0/1 (F0/1)	Serial 0/0/0 (S0/0/0)	Serial 0/0/1 (S0/0/1)
2900	Gigabit Ethernet 0/0 (G0/0)	Gigabit Ethernet 0/1 (G0/1)	Serial 0/0/0 (S0/0/0)	Serial 0/0/1 (S0/0/1)
4221	Gigabit Ethernet 0/0/0 (G0/0/0)	Gigabit Ethernet 0/0/1 (G0/0/1)	Serial 0/1/0 (S0/1/0)	Serial 0/1/1 (S0/1/1)
4300	Gigabit Ethernet 0/0/0 (G0/0/0)	Gigabit Ethernet 0/0/1 (G0/0/1)	Serial 0/1/0 (S0/1/0)	Serial 0/1/1 (S0/1/1)

Note: To find out how the router is configured, look at the interfaces to identify the type of router and how many interfaces the router has. There is no way to effectively list all the combinations of configurations for each router class. This table includes identifiers for the possible combinations of Ethernet and Serial interfaces in the device. The table does not include any other type of interface, even though a specific router may contain one. An example of this might be an ISDN BRI interface. The string in parenthesis is the legal abbreviation that can be used in Cisco IOS commands to represent the interface.

Device Configs - Final

Router R1

```
R1# show running-config
Building configuration...
```

```
Current configuration : 1576 bytes
!
! Last configuration change at 21:51:24 UTC Thu Dec 19 2019
!
version 16.9
service timestamps debug datetime msec
service timestamps log datetime msec
platform qfp utilization monitor load 80
no platform punt-keepalive disable-kernel-core
!
hostname R1
!
boot-start-marker
boot-end-marker
!
```

Lab - Implement eBGP for IPv4

```
no aaa new-model
!
no ip domain lookup
!
login on-success log
!
subscriber templating
!
multilink bundle-name authenticated
!
license udi pid ISR4221/K9 sn FGL23313183
no license smart enable
diagnostic bootup level minimal
!
spanning-tree extend system-id
!
redundancy
mode none
!
interface Loopback0
 ip address 192.168.1.1 255.255.255.224
!
interface Loopback1
 ip address 192.168.1.65 255.255.255.192
!
interface GigabitEthernet0/0/0
 ip address 10.1.2.1 255.255.255.0
 negotiation auto
!
interface GigabitEthernet0/0/1
 no ip address
 negotiation auto
!
interface Serial0/1/0
 ip address 10.1.3.1 255.255.255.128
 shutdown
!
interface Serial0/1/1
 ip address 10.1.3.129 255.255.255.128
 shutdown
!
router bgp 1000
 bgp router-id 1.1.1.1
 bgp log-neighbor-changes
 network 192.168.1.0 mask 255.255.255.224
 network 192.168.1.64 mask 255.255.255.192
 neighbor 10.1.2.2 remote-as 500
 neighbor 10.1.3.3 remote-as 300
 neighbor 10.1.3.130 remote-as 300
!
```

```
ip forward-protocol nd
no ip http server
ip http secure-server
!
control-plane
!
!
line con 0
  exec-timeout 0 0
  logging synchronous
  transport input none
  stopbits 1
line aux 0
  stopbits 1
line vty 0 4
  login
!
end
```

Router R2

```
R2# show running-config
Building configuration...
```

```
Current configuration : 1519 bytes
!
! Last configuration change at 15:45:14 UTC Sun Dec 22 2019
!
version 16.9
service timestamps debug datetime msec
service timestamps log datetime msec
platform qfp utilization monitor load 80
no platform punt-keepalive disable-kernel-core
!
hostname R2
!
boot-start-marker
boot-end-marker
!
no aaa new-model
!
no ip domain lookup
!
login on-success log
!
subscriber templating
!
multilink bundle-name authenticated
!
license udi pid ISR4221/K9 sn FGL23313182
```

Lab - Implement eBGP for IPv4

```
no license smart enable
diagnostic bootup level minimal
!
spanning-tree extend system-id
!
redundancy
mode none
!
interface Loopback0
ip address 192.168.2.1 255.255.255.224
!
interface Loopback1
ip address 192.168.2.65 255.255.255.192
!
interface GigabitEthernet0/0/0
ip address 10.1.2.2 255.255.255.0
negotiation auto
!
interface GigabitEthernet0/0/1
ip address 10.2.3.2 255.255.255.0
negotiation auto
!
router bgp 500
bgp router-id 2.2.2.2
bgp log-neighbor-changes
network 192.168.2.0 mask 255.255.255.224
network 192.168.2.64 mask 255.255.255.192
aggregate-address 192.168.1.0 255.255.255.0 as-set summary-only
neighbor 10.1.2.1 remote-as 1000
neighbor 10.1.2.1 default-originate
neighbor 10.2.3.3 remote-as 300
!
ip forward-protocol nd
no ip http server
ip http secure-server
!
control-plane
!
line con 0
exec-timeout 0 0
logging synchronous
transport input none
stopbits 1
line aux 0
stopbits 1
line vty 0 4
login
!
end
```

Router R3

```
R3# show running-config
```

```
Building configuration...
```

```
Current configuration : 1781 bytes
```

```
!  
! Last configuration change at 18:11:15 UTC Sat Dec 21 2019  
!  
version 16.9  
service timestamps debug datetime msec  
service timestamps log datetime msec  
platform qfp utilization monitor load 80  
no platform punt-keepalive disable-kernel-core  
!  
hostname R3  
!  
boot-start-marker  
boot-end-marker  
!  
no aaa new-model  
!  
no ip domain lookup  
!  
login on-success log  
!  
subscriber templating  
!  
multilink bundle-name authenticated  
!  
license udi pid ISR4221/K9 sn FGL23313186  
no license smart enable  
diagnostic bootup level minimal  
!  
spanning-tree extend system-id  
!  
redundancy  
mode none  
!  
interface Loopback0  
ip address 192.168.3.1 255.255.255.224  
!  
interface Loopback1  
ip address 192.168.3.65 255.255.255.192  
!  
interface GigabitEthernet0/0/0  
ip address 10.2.3.3 255.255.255.0  
negotiation auto  
!  
interface GigabitEthernet0/0/1
```

Lab - Implement eBGP for IPv4

```
no ip address
negotiation auto
!
interface Serial0/1/0
ip address 10.1.3.3 255.255.255.128
!
interface Serial0/1/1
ip address 10.1.3.130 255.255.255.128
!
router bgp 300
bgp router-id 3.3.3.3
bgp log-neighbor-changes
no bgp default ipv4-unicast
neighbor 10.1.3.1 remote-as 1000
neighbor 10.1.3.129 remote-as 1000
neighbor 10.2.3.2 remote-as 500
!
address-family ipv4
network 192.168.3.0 mask 255.255.255.224
network 192.168.3.64 mask 255.255.255.192
aggregate-address 192.168.3.0 255.255.255.0 summary-only
neighbor 10.1.3.1 activate
neighbor 10.1.3.129 activate
neighbor 10.2.3.2 activate
exit-address-family
!
ip forward-protocol nd
no ip http server
ip http secure-server
!
control-plane
!
line con 0
exec-timeout 0 0
logging synchronous
transport input none
stopbits 1
line aux 0
stopbits 1
line vty 0 4
login
!
end
```