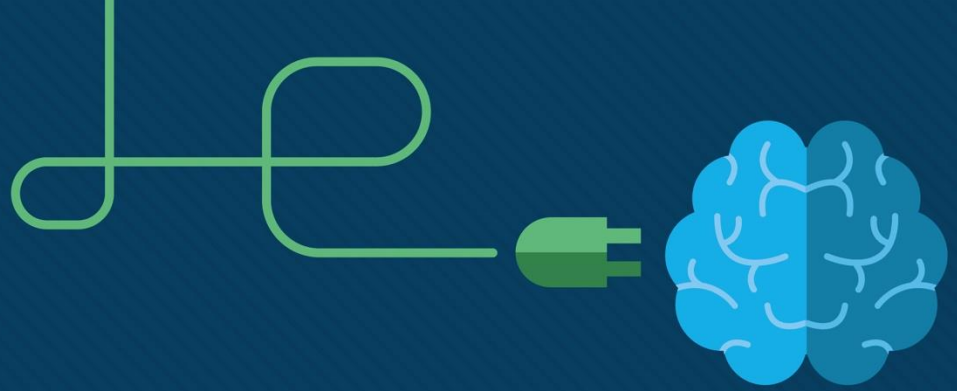

Chapter 11: BGP

Instructor Materials

CCNP Enterprise: Core Networking

Chapter 11 Content

This chapter covers the following content:

- **BGP Fundamentals -** This section provides an overview of the fundamentals of the BGP routing protocol.
- **Basic BGP Configuration -** This section walks through the process of configuring BGP to establish a neighbor session and how routes are exchanged between peers.
- **Route Summarization -** This section provides an overview of how route summarization works with BGP and some of the design considerations with summarization.
- **Multiprotocol BGP for IPv6 -** This section explains how BGP provides support for IPv6 routing and configuration.

BGP Fundamentals

- Border Gateway Protocol (BGP) is an EGP standardized path vector routing protocol that provides scalability, flexibility, and network stability. When BGP was created, the primary design consideration was for IPv4 inter-organization connectivity on public networks like the Internet and on private dedicated networks.

Autonomous System Numbers

From the perspective of BGP, an autonomous system (AS) is a collection of routers under a single organization's control, using one or more IGPs and common metrics.

- An organization requiring connectivity to the internet must obtain an autonomous system number (ASN) from the Internet Assigned Numbers Authority (IANA).
- RFC 4893 expanded the ASN field from 2 bytes to accommodate 4 bytes (32-bit range). This allows for 4,294,967,295 unique ASNs, providing quite an increase from the original 65,535 ASNs.
- The IANA requires organizations or individuals requesting AS numbers to meet the following conditions:
  - Proof of a publicly allocated network range
  - Proof that Internet connectivity is provided through multiple connections
  - Demonstrated need for a unique routing policy from their providers.

Autonomous System Numbers (Cont.)

Similar to IP Addresses, the AS number assignments include private address ranges that cannot be used for Internet traffic. These ranges are:

- ASNs 64,512–65,535 are private ASNs in the 16-bit ASN range.
- ASNs 4,200,000,000–4,294,967,294 are private ASNs within the extended 32-bit range.
- Use only the ASN assigned by IANA, the ASN assigned by your service provider, or a private ASN. Using another organization's ASN without permission could result in traffic loss and cause havoc on the internet.

Path Attributes

BGP uses path attributes (PAs) associated with each network path. The PAs provide BGP with granularity and control of routing policies within BGP. The BGP prefix PAs are classified as follows:

- Well-known mandatory
- Well-known discretionary
- Optional transitive
- Optional non-transitive

Per RFC 4271, well-known attributes must be recognized by all BGP implementations. Well-known mandatory attributes must be included with every prefix advertisement; well-known discretionary attributes may or may not be included with a prefix advertisement.

Loop Prevention

BGP is a path vector routing protocol and does not contain a complete topology of the network, as link-state routing protocols do. BGP behaves like distance vector protocols, ensuring that a path is loop free.

The Loop Prevention Mechanism:

- AS 100 advertises the 172.16.1.0/24 prefix to AS 200.
- AS 200 advertises the prefix to AS 400, which then advertises the prefix to AS 300.
- AS 300 advertises the prefix back to AS 100 with an AS_Path of 300 400 200 100. AS 100 sees itself in the AS_Path variable and discards the prefix.

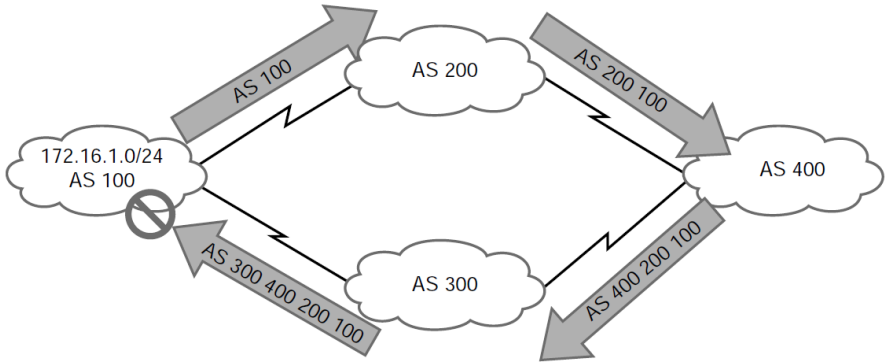

Figure 11-1 *Path Vector Loop Prevention*

Address Families

- RFC 2858 added Multi-Protocol BGP (MP-BGP) capability by adding an extension called the address family identifier (AFI).
- An address family correlates to a specific network protocol, such as IPv4 or IPv6, and additional granularity is provided through a subsequent address-family identifier (SAFI) such as unicast or multicast.
- MBGP achieves this separation by using the BGP path attributes (PAs) MP_REACH_NLRI and MP_UNREACH_NLRI to carry different reachability information for the different address families.
- Every address family maintains a separate database and configuration for each protocol (address family + sub-address family) in BGP.
- Some network engineers refer to Multiprotocol BGP as MP-BGP, and other network engineers use the term MBGP. Both terms refer to the same thing.

Inter-Router Communication

- BGP does not use hello packets to discover neighbors, as do IGP protocols, and it cannot discover neighbors dynamically. BGP neighbors are defined by IP address.
- BGP uses TCP port 179 to communicate with other routers.
- Most recent implementations of BGP set the do-not-fragment (DF) bit to prevent fragmentation and rely on path MTU discovery.
- BGP uses TCP, so unlike IGPs, it can form adjacencies with routers that are multiple hops away.
- A BGP session refers to the established adjacency between two BGP routers. Multi-hop sessions require that the router use an underlying route installed in the RIB (static or from any routing protocol) to establish the TCP session with the remote endpoint.

Inter-Router Communication (Cont.)

BGP can be thought of as a control plane routing protocol or as an application because it allows for the exchange of routes with a peer that is multiple hops away.

- A BGP session refers to the established adjacency between two BGP routers.
- BGP neighbors connected to the same network use the ARP table to locate the IP address of the peer. Multi-hop BGP sessions require routing table information for finding the IP address of the peer.
- A default route is not sufficient to establish a multi-hop BGP session.

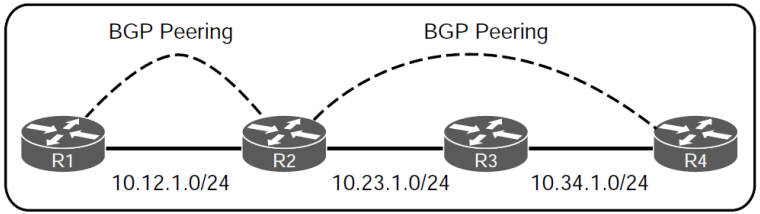

Figure 11-2 *BGP Single- and Multi-Hop Sessions*

BGP Session Types

BGP sessions are categorized into two types:

Internal BGP (iBGP) - Sessions established with an iBGP router that are in the same AS or that participate in the same BGP confederation. iBGP prefixes are assigned an administrative distance (AD) of 200 upon installation in the router's RIB.

External BGP (eBGP) - Sessions established with a BGP router that are in a different AS. eBGP prefixes are assigned an AD of 20 upon installation in the router's RIB.

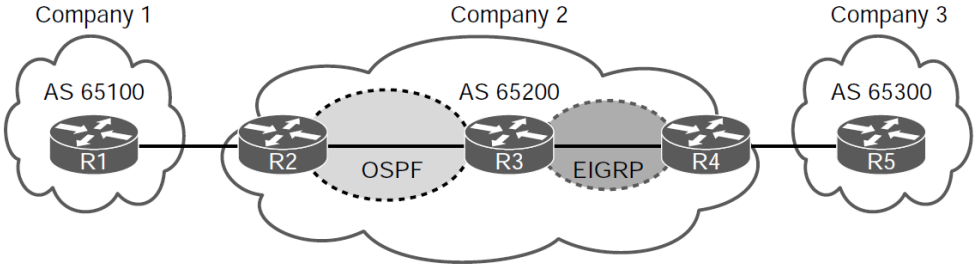

Figure 11-3 *AS 65200 Providing Transit Connectivity*

BGP Session Types (Cont.)

iBGP

Service providers provide transit connectivity. Enterprises should not. Preventing an AS from becoming a transit AS by redistributing BGP routes into the IGP is not a viable solution for the following reasons:

- **Scalability -** The internet has over 780,000 IPv4 network prefixes, IGPs cannot scale to that level of routes.
- **Custom routing -** IGPs use metrics as the primary method of route selection. BGP uses multiple steps to identify the best path or to manipulate the path for a specific network prefix.
- **Path attributes -** All of the BGP path attributes cannot be maintained within IGP protocols.

Establishing iBGP sessions between IGP routers in a full mesh allows for proper forwarding between autonomous systems.

BGP Session Types (Cont.)

eBGP

eBGP is the core component of BGP on the internet. The difference in behavior for eBGP sessions and iBGP sessions are:

- iBGP allows for multi-hop sessions. eBGP packets set the TTL to one, so packets are dropped if multi-hop connections are attempted.
- The advertising router modifies the BGP next-hop address to the IP address sourcing the BGP connection.
- The advertising router prepends its ASN to the existing AS_Path variable.
- The receiving router verifies that the AS_Path variable does not contain an ASN that matches the local routers. BGP discards the NLRI if it fails the AS_Path loop prevention check.

BGP Messages

Table 11-2 BGP Packet Types

TYPE	NAME	FUNCTIONAL OVERVIEW
1	OPEN	Sets up and establishes BGP adjacency
2	UPDATE	Advertises, updates, or withdraws routes
3	NOTIFICATION	Indicates an error condition to a BGP neighbor
4	KEEPALIVE	Ensures that BGP neighbors are still alive

BGP Neighbor States

The BGP session may report the following states:

- Idle
- Connect
- Active
- OpenSent
- OpenConfirm
- Established

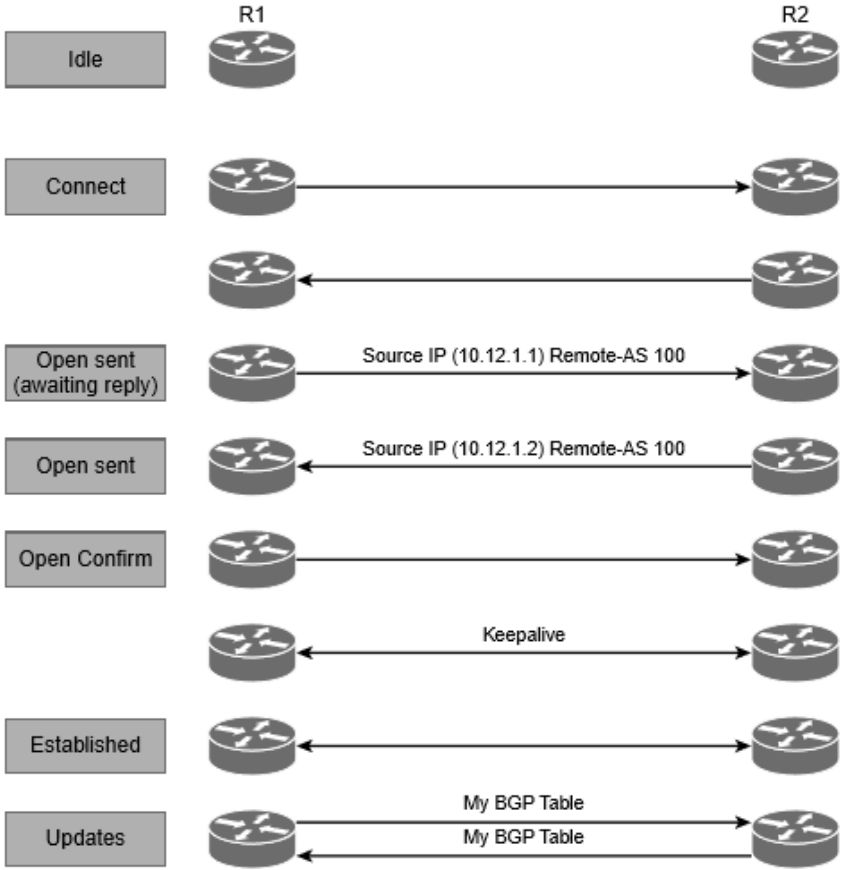

Figure 11-6 *BGP Neighbor States with Session Establishment*

BGP Neighbor States (Cont.)

BGP forms a TCP session with neighbor routers called peers. BGP uses the finite-state machine (FSM) to maintain a table of all BGP peers and their operational status.

- **Idle -** The first stage of the BGP FSM. BGP detects a start event and attempts to initiate a TCP session with the BGP peer.
- **Connect -** BGP initiates the TCP connection. When the TCP handshake is completed, the ConnectRetryTimer is reset, an Open message is sent to the neighbor, and the state changes to OpenSent. During this stage, the router with the higher IP address manages the connection.
- **Active -** BGP starts another TCP three-way handshake. If the connection is established an Open message is sent and the state moves to OpenSent.
- **OpenSent -** If the OPEN message that is received by the originating router contains no errors, the connection state moves to OpenConfirm
- **OpenConfirm -** In the OpenConfirm state, BGP waits for KEEPALIVE or NOTIFICATION messages. Upon receipt of a neighbor's KEEPALIVE message, the state is moved to Established.
- **Established -** BGP neighbors exchange routes using UPDATE messages.

Basic BGP Configuration

When configuring BGP, it is best to think of the configuration from a modular perspective. BGP router configuration requires the following components:

- BGP session parameters
- Address family initialization
- Activate the address family on the BGP peer

BGP Router Configuration Components

- **BGP session parameters -** BGP session parameters provide settings that involve estab- lishing communication to the remote BGP neighbor. Session settings include the ASN of the BGP peer, authentication, and keepalive timers.
- **Address family initialization -** The address family is initialized under the BGP router configuration mode. Network advertisement and summarization occur within the address family.
- **Activate the address family on the BGP peer -** In order for a session to initiate, one address family for a neighbor must be activated. The router's IP address is added to the neighbor table, and BGP attempts to establish a BGP session or accepts a BGP session initiated from the peer router.

Modular Configuration

Steps to Configure Basic BGP are as follows:

Step 1. Initialize the BGP routing process with the global command **router bgp** *as-number***.**

Step 2. (Optional) Statically define the BGP router ID (RID). The dynamic RID allocation logic uses the highest IP address of any of the up loopback interfaces.

Step 3. Identify the BGP neighbor's IP address and autonomous system number with the BGP router configuration command **neighbor** *ip-address* **remote-as** *as-number*.

Step 4. Initialize the address family with the BGP router configuration command **address-family** *afi safi***.**

Step 5. Activate the address family for the BGP neighbor with the BGP address family configuration command **neighbor** *ip-address* **activate.**

Configuring Basic BGP on IOS

R1 (Default IPv4 Address-Family Enabled)

```
router bgp 65100
neighbor 10.12.1.2 remote-as 65200
```

R2 (Default IPv4 Address-Family Disabled)

```
router bgp 65200
 no bgp default ipv4-unicast
neighbor 10.12.1.1 remote-as 65100
!
address-family ipv4
 neighbor 10.12.1.1 activate
exit address-family
```

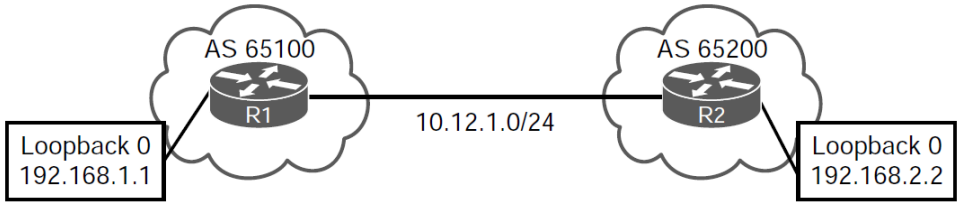

Figure 11-7 *Simple BGP Topology*

Verification of BGP Sessions

- The BGP session is verified with the command **show bgp** *afi safi* **summary.**
- Earlier commands, such as **show ip bgp summary**, came out before MBGP and do not provide a structure for the current multiprotocol capabilities within BGP.

Example 11-3 *Verifying the BGP IPv4 Session Summary*

```
R1# show bgp ipv4 unicast summary
BGP router identifier 192.168.2.2, local AS number 65200
BGP table version is 1, main routing table version 1

Neighbor        V     AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.12.1.2       4  65200       8       9        1    0    0 00:05:23        0
```

Verification of BGP Sessions (Cont.)

Field	Description
Neighbor	IP address of the BGP peer
V	BGP version spoken by the BGP peer
AS	Autonomous system number of the BGP peer
MsgRcvd	Count of messages received from the BGP peer
MsgSent	Count of messages sent to the BGP peer
TblVer	Last version of the BGP database sent to the peer
InQ	Number of messages queued to be processed by the peer
OutQ	Number of messages queued to be sent to the peer
Up/Down	Length of time the BGP session is established or the current status if the session is not in an established state
State/PfxRcd	Current state of the BGP peer or the number of prefixes received from the peer

Table 11-3 BGP Summary Fields

Prefix Advertisement

- BGP network statements do not enable BGP for a specific interface; instead, they identify specific network prefixes to be installed into the BGP table, known as the Loc-RIB table.
- As the BGP prefix is installed into the Loc-RIB table, the following BGP PAs are set, depending on the RIB prefix type:
  - **Connected network -** The next-hop BGP attribute is set to 0.0.0.0, the BGP origin attribute is set to i (IGP), and the BGP weight is set to 32,768.
  - **Static route or routing protocol -** The next-hop BGP attribute is set to the next-hop IP address in the RIB, the BGP origin attribute is set to i (IGP), the BGP weight is set to 32,768, and the MED is set to the IGP metric.

Prefix Advertisement (Cont.)

All routes in the Loc-RIB table use the following process for advertisement to BGP peers:

Step 1. Pass a validity check. Verify that the NRLI is valid and that the next-hop address is resolvable in the global RIB. If the NRLI fails, the NLRI remains but does not process further.

Step 2. Process outbound neighbor route policies. After processing, if a route was not denied by the outbound policies, the route is maintained in the Adj-RIB-Out table for later reference.

Step 3. Advertise the NLRI to BGP peers. If the NLRI's next-hop BGP PA is 0.0.0.0, then the next-hop address is changed to the IP address of the BGP session.

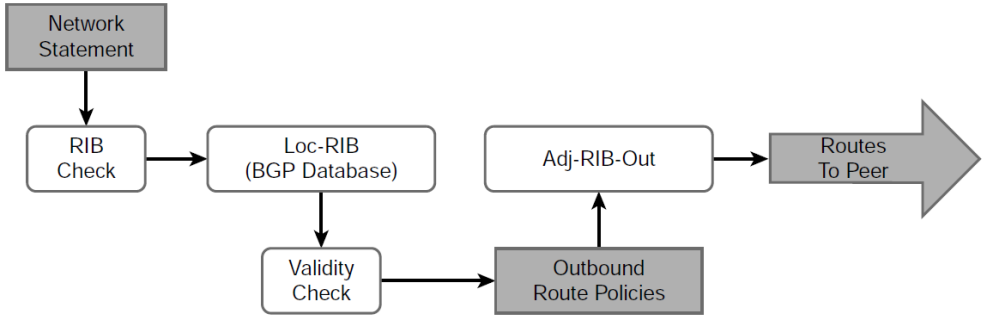

Figure 11-8 *BGP Database Processing of Local Route Advertisements*

Prefix Advertisement (Cont.)

The network statement resides under the appropriate address family within the BGP router configuration. The command **network** *network* **mask** *subnet-mask* [**route-map** *route-map-name*] is used for advertising IPv4 networks. The optional ***route-map*** provides a method of setting specific BGP PAs when the prefix installs into the Loc-RIB table.

```
router bgp 65200
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 10.12.1.1 remote-as 65100
 !
 address-family ipv4
  network 10.12.1.0 mask 255.255.255.0
  network 192.168.2.2 mask 255.255.255.255
  neighbor 10.12.1.1 activate
 exit-address-family
```

Receiving and Viewing Routes

BGP uses three tables for maintaining the network prefix and PAs for a route:

- **Adj-RIB-In -** Contains the NLRIs in original form
- **Loc-RIB -** Contains all the NLRIs that originated locally or were received from other BGP peers.
- **Adj-RIB-Out -** Contains the NLRIs after outbound route policies have been processed.

Processing Steps are shown in Figure 11-9.

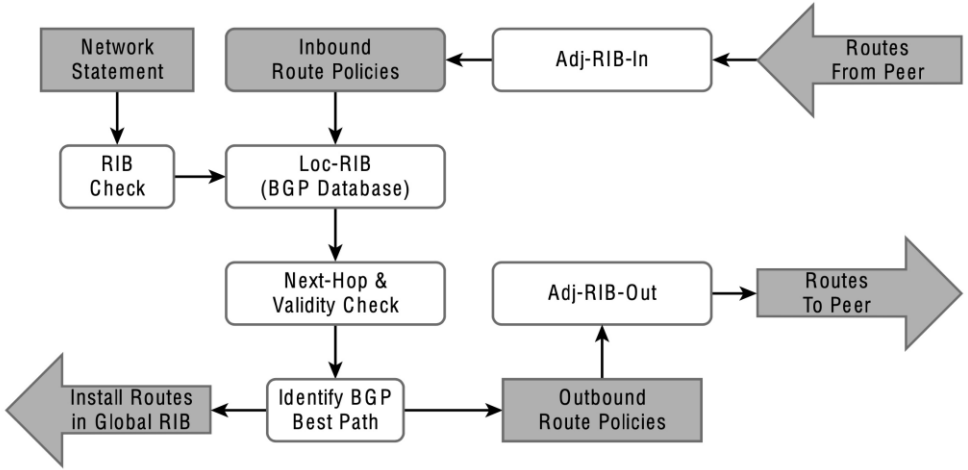

Figure 11-9 *BGP Database Processing*

Receiving and Viewing Routes (Cont.)

The command **show bgp** *afi safi* displays the contents of the BGP database (Loc-RIB) on the router.

Example 11-6 *Displaying the BGP Table*

```
R1# show bgp ipv4 unicast
BGP table version is 4, local router ID is 192.168.1.1
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
 *   10.12.1.0/24     10.12.1.2                0             0 65200 i
 *>                   0.0.0.0                  0         32768 i
 *>  192.168.1.1/32   0.0.0.0                  0         32768 i
 *>  192.168.2.2/32   10.12.1.2                0             0 65200 i
```

```
R2# show bgp ipv4 unicast | begin Network
     Network          Next Hop            Metric LocPrf Weight Path
 *   10.12.1.0/24     10.12.1.1                0             0 65100 i
 *>                   0.0.0.0                  0         32768 i
 *>  192.168.1.1/32   10.12.1.1                0             0 65100 i
 *>  192.168.2.2/32   0.0.0.0                  0         32768 i
```

Receiving and Viewing Routes (Cont.)

Field	Description
Network	A list of the network prefixes installed in BGP. Valid NLRIs are indicated by the *. The NLRI selected as the best path is indicated by an angle bracket (>).
Next Hop	A well-known mandatory BGP path attribute that defines the IP address for the next hop for that specific NLRI.
Metric	Multiple-exit discriminator (MED): An optional non-transitive BGP path attribute used in BGP for the specific NLRI.
LocPrf	Local Preference: A well-known discretionary BGP path attribute used in the BGP best-path algorithm for the specific NLRI.
Weight	A locally significant Cisco-defined attribute used in the BGP best-path algorithm for the specific NLRI.
Path and Origin	AS_Path: A well-known mandatory BGP path attribute used for loop prevention and in the BGP best-path algorithm for the specific NLRI. Origin: A well-known mandatory BGP path attribute used in the BGP best-path algorithm. A value of i represents an IGP, e indicates EGP, and ? indicates a route that was redistributed into BGP.

Table 11-4 BGP Table Fields

Receiving and Viewing Routes (Cont.)

The command **show bgp** *afi safi network* displays all the paths for a specific route and the BGP path attributes for that route.

Example 11-7 *Viewing Explicit BGP Routes and Path Attributes*

```
R1# show bgp ipv4 unicast 10.12.1.0
BGP routing table entry for 10.12.1.0/24, version 2
Paths: (2 available, best #2, table default)
  Advertised to update-groups:
     2
  Refresh Epoch 1
  65200
    10.12.1.2 from 10.12.1.2 (192.168.2.2)
      Origin IGP, metric 0, localpref 100, valid, external
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  Local
    0.0.0.0 from 0.0.0.0 (192.168.1.1)
      Origin IGP, metric 0, localpref 100, weight 32768, valid, sourced, local, best
      rx pathid: 0, tx pathid: 0x0
```

Receiving and Viewing Routes (Cont.)

The Adj-RIB-Out table is a unique table maintained for each BGP peer. It enables a network engineer to view routes advertised to a specific router. The command **show bgp** *afi safi* **neighbor** *ip-address* **advertised routes** displays the contents of the Adj-RIB-Out table for a neighbor.

Example 11-8 *Neighbor-Specific View of the Adj-RIB-Out Table*

```
R1# show bgp ipv4 unicast neighbors 10.12.1.2 advertised-routes
! Output omitted for brevity
     Network          Next Hop            Metric LocPrf Weight Path
*> 10.12.1.0/24     0.0.0.0                  0         32768 i
*> 192.168.1.1/32   0.0.0.0                  0         32768 i

Total number of prefixes 2
```

```
R2# show bgp ipv4 unicast neighbors 10.12.1.1 advertised-routes
! Output omitted for brevity
     Network          Next Hop            Metric LocPrf Weight Path
*> 10.12.1.0/24     0.0.0.0                  0         32768 i
*> 192.168.2.2/32   0.0.0.0                  0         32768 i

Total number of prefixes 2
```

BGP Route Advertisements from Indirect Sources

BGP should be thought of as a routing application as the BGP session and route advertisement are two separate components. BGP can install multiple routes learned from static routes, EIGRP, and OSPF.

Example 11-11 *R1's Routing Table with Loopbacks for R3, R4, and R5*

```
R1# show ip route
! Output omitted for brevity
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
..
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
C        10.12.1.0/24 is directly connected, GigabitEthernet0/0
C        10.13.1.0/24 is directly connected, GigabitEthernet0/1
C        10.14.1.0/24 is directly connected, GigabitEthernet0/2
C        10.15.1.0/24 is directly connected, GigabitEthernet0/3
C        192.168.1.1 is directly connected, Loopback0
B        192.168.2.2 [20/0] via 10.12.1.2, 00:01:17
D        192.168.3.3 [90/3584] via 10.13.1.3, 00:02:10, GigabitEthernet0/1
S        192.168.4.4 [1/0] via 10.14.1.4
O        192.168.5.5 [110/11] via 10.15.1.5, 00:00:08, GigabitEthernet0/3
```

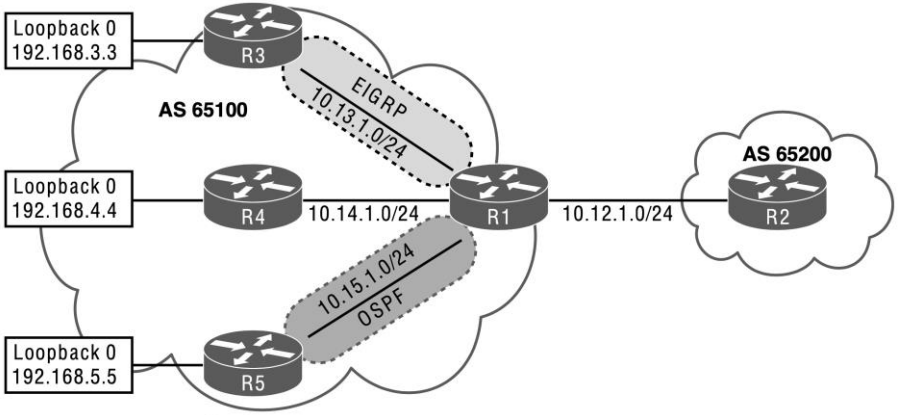

Figure 11-10 *Multiple BGP Route Sources*

BGP Route Advertisements from Indirect Sources (Cont.)

Redistributing routes learned from an IGP into BGP is completely safe. Redistributing routes learned from BGP into an IGP should be done with extreme caution.

Example 11-12 Configuring Advertising Routes for Non-Connected Routes

R1

```
router bgp 65100
 bgp log-neighbor-changes
 network 10.12.1.0 mask 255.255.255.0
 network 192.168.1.1 mask 255.255.255.255
 network 192.168.3.3 mask 255.255.255.255
 network 192.168.4.4 mask 255.255.255.255
 redistribute ospf 1
 neighbor 10.12.1.2 remote-as 65200
```

Route Summarization

Summarizing prefixes conserves router resources and accelerates best-path calculation by reducing the size of the table. Summarization also provides the benefit of stability by hiding route flaps from downstream routers, thereby reducing routing churn.

Summarizing Prefixes

There are two techniques for BGP summarization:

- **Static -** Create a static route to Null0 for the summary network prefix and then advertise the prefix with a network statement. The downfall of this technique is that the summary route is always advertised, even if the networks are not available.
- **Dynamic -** Configure an aggregation network prefix. When viable component routes that match the aggregate network prefix enter the BGP table, then the aggregate prefix is created. The originating router sets the next hop to Null0 as a discard route for the aggregated prefix for loop prevention.

Aggregate Address

Dynamic route summarization is accomplished with the BGP address family configuration command **aggregate-address** *network subnet-mask* [**summary-only**] [**as-set**].

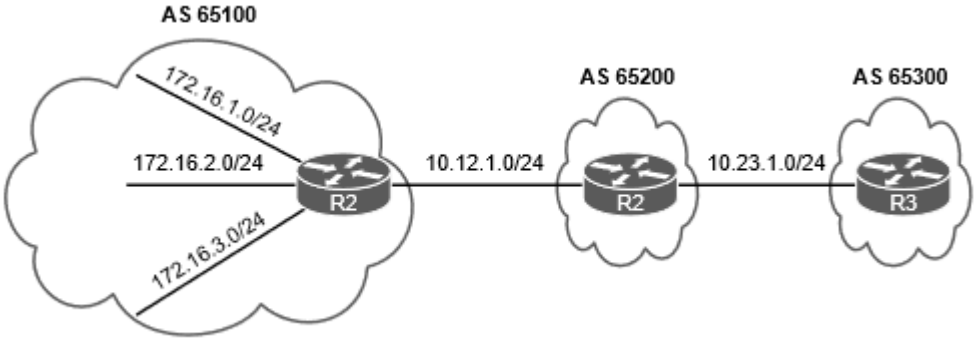

Figure 11-12 *BGP Summarization Topology*

Aggregating Address (Cont.)

R1

```
router bgp 65100
 bgp log-neighbor-changes
 network 10.12.1.0 mask 255.255.255.0
 network 192.168.1.1 mask 255.255.255.255
 network 192.168.3.3 mask 255.255.255.255
 network 192.168.4.4 mask 255.255.255.255
 redistribute ospf 1
 neighbor 10.12.1.2 remote-as 65200
```

R2

```
router bgp 65100
 bgp log-neighbor-changes
 neighbor 10.12.1.1 remote-as 65100
 neighbor 10.23.1.3 remote-as 65300
 !
 address-family ipv4
  aggregate-address 192.168.0.0 255.255.0.0
  aggregate-address 172.16.0.0 255.255.240.0
  redistribute connected
  neighbor 10.12.1.1 activate
  neighbor 10.23.1.3 activate
 exit-address-family
```

Example 11-15 Configuring BGP Route Aggregation

Atomic Aggregate

Aggregated routes act like new BGP routes with a shorter prefix length. When a BGP router summarizes a route, it does not advertise the AS_Path information from before the aggregation. BGP path attributes like AS_Path, MED, and BGP communities are not included in the new BGP advertisement.

R2

```
router bgp 65200
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 10.12.1.1 remote-as 65100
 neighbor 10.23.1.3 remote-as 65300
 !
 address-family ipv4
  aggregate-address 192.168.0.0 255.255.0.0 summary-only
  aggregate-address 172.16.0.0 255.255.240.0 summary-only
  redistribute connected
  neighbor 10.12.1.1 activate
  neighbor 10.23.1.3 activate
 exit-address-family
```

Example 11-21 *Configuring Aggregation for 172.16.0.0/20 and 192.168.0.0/16*

Atomic Aggregate (Cont.)

```
R2#show bgp ipv4 unicast 172.16.0.0
BGP routing table entry for 172.16.0.0/20,
version 6
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 2
  65200, (aggregated by 65200 192.168.2.2)
    10.23.1.2 from 10.23.1.2 (192.168.2.2)
      Origin IGP, metric 0,localpref 100, valid
32768, valid, external, atomic-aggregate, best
      rx pathid: 0, tx pathid: 0x0
```

Example 11-23 Examining the BGP Attribute for the Atomic Aggregate Attribute

Route Aggregation with AS_SET

To keep the BGP path information history, the optional **as-set** keyword may be used with the **aggregate-address** command. As the router generates the aggregate route, BGP attributes from the component aggregate routes are copied over to it.

```
R2#show running-config | section router bgp
router bgp 65200
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 10.12.1.1 remote-as 65100
 neighbor 10.23.1.3 remote-as 65300
 !
 address-family ipv4
  aggregate-address 192.168.0.0 255.255.0.0 as-set summary-only
  aggregate-address 172.16.0.0 255.255.240.0 as-set summary-only
  redistribute connected
  neighbor 10.12.1.1 activate
  neighbor 10.23.1.3 activate
 exit-address-family
```

Example 11-24 Configuring Aggregation While Preserving BGP Attributes

Multiprotocol BGP for IPv6

Multiprotocol BGP (MP-BGP) enables BGP to carry NLRI for multiple protocols, such as IPv4, IPv6, and Multiprotocol Label Switching (MPLS) Layer 3 virtual private networks (L3VPNs).

MP-BGP

RFC 4760 defines the following new features:

- A new address family identifier (AFI) model
- New BGPv4 optional and nontransitive attributes:
  - Multiprotocol reachable NLRI
  - Multiprotocol unreachable NLRI

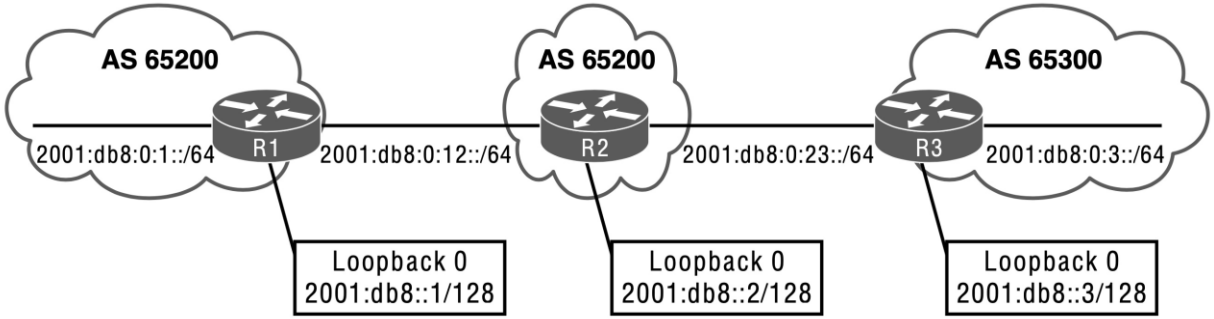

Figure 11-13 *IPv6 Sample Topology*

MP-BGP (Cont.)

- All the same underlying IPv4 path vector routing protocol features and rules also apply to MP-BGP for IPv6.
- MP-BGP for IPv6 continues to use the same well-known TCP port 179 for session peering as BGP uses for IPv4.
- The MP-BGP extensions include an address family identifier (AFI) that describes the supported protocols, along with subsequent address family identifier (SAFI) attribute fields that describe whether the prefix applies to the unicast or multicast routing table:
  - IPv4 unicast: AFI: 1, SAFI: 1
  - IPv6 unicast: AFI: 2, SAFI: 1

Configuring IPv6 BGP

All the BGP configuration rules demonstrated earlier apply with IPv6, except that the IPv6 address family must be initialized, and the neighbor is activated.

R1

```
router bgp 65100
 bgp router-id 192.168.1.1
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 2001:DB8:0:12::2 remote-as 65200
 !
 address-family ipv6
  redistribute connected
  neighbor 2001:DB8:0:12::2 activate
 exit-address-family
```

R2

```
router bgp 65200
 bgp router-id 192.168.2.2
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 2001:DB8:0:12::1 remote-as 65100
 neighbor 2001:DB8:0:23::3 remote-as 65300
 !
 address-family ipv6
  network 2001:DB8::2/128
  network 2001:DB8:0:12::/64
  network 2001:db8:0:23::/64
  neighbor 2001:DB8:0:12::1 activate
  neighbor 2001:DB8:0:23::3 activate
 exit-address-family
```

Example 11-28 *Configuring IPv6 BGP*

Viewing BGP IPv6 Neighbors

- Routers exchange AFI capabilities during the initial BGP session negotiation.
- The command **show bgp ipv6 unicast neighbors** *ip-address* **[detail]** displays detailed information on whether or not the IPv6 capabilities were negotiated successfully.

```
R1# show bgp ipv6 unicast neighbors 2001:DB8:0:12::2
! Output omitted for brevity
BGP neighbor is 2001:DB8:0:12::2, remote AS 65200, external link
BGP version 4, remote router ID 192.168.2.2
BGP state = Established, up for 00:28:25
Last read 00:00:54, last write 00:00:34, 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 IPv6 Unicast: advertised and received
Enhanced Refresh Capability: advertised and received
```

Example 11-29 *Viewing BGP Neighbors for IPv6 Capabilities*

Verifying IPv6 BGP

```
R2# show bgp ipv6 unicast summary
BGP router identifier 192.168.2.2, local AS number 65200
BGP table version is 19, main routing table version 19
7 network entries using 1176 bytes of memory
8 path entries using 832 bytes of memory
3/3 BGP path/bestpath attribute entries using 456 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 2512 total bytes of memory
BGP activity 7/0 prefixes, 8/0 paths, scan interval 60 secs
Neighbor         V AS     MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
2001:DB8:0:12::1 4 65100      35      37    19    0    0 00:25:08        3
2001:DB8:0:23::3 4 65300      32      37    19    0    0 00:25:11        3
```

Example 11-30 *Verifying an IPv6 BGP Session*

Viewing the IPv6 BGP Tables

The unspecified address (::) indicates that the local router is generating the prefix for the BGP table.

```
R2# show bgp ipv6 unicast | begin Network
     Network                Next Hop            Metric LocPrf Weight Path
*>   2001:DB8::1/128        2001:DB8:0:12::1         0             0 65100 ?
*>   2001:DB8::2/128        ::                       0         32768 i
*>   2001:DB8::3/128        2001:DB8:0:23::3         0             0 65300 i
*>   2001:DB8:0:1::/64      2001:DB8:0:12::1         0             0 65100 ?
*>   2001:DB8:0:3::/64      2001:DB8:0:23::3         0             0 65300 i
*>   2001:DB8:0:12::/64     ::                       0         32768 i
*                           2001:DB8:0:12::1         0             0 65100 ?
*>  2001:DB8:0:23::/64      ::                       0         32768 i
                            2001:DB8:0:23::3         0             0 65300 i
```

Example 11-31 *Viewing the IPv6 BGP Tables*

IPv6 Route Summarization

The same process for summarizing or aggregating IPv4 routes occurs with IPv6 routes, and the format is identical except that the configuration is placed under the IPv6 address family using the command **aggregate-address** *prefix/prefix-length* [**summary-only**] [**as-set**].

Bits Needed	Summary Address	Component Networks
2	2001:db8:0:0::/62	2001:db8:0:0::/64 through 2001:db8:0:3::/64
3	2001:db8:0:0::/61	2001:db8:0:0::/64 through 2001:db8:0:7::/64
4	2001:db8:0:0::/60	2001:db8:0:0::/64 through 2001:db8:0:F::/64
5	2001:db8:0:0::/59	2001:db8:0:0::/64 through 2001:db8:0:1F::/64
6	2001:db8:0:0::/58	2001:db8:0:0::/64 through 2001:db8:0:3F::/64

Table 11-5 IPv6 Summarization Table

Prepare for the Exam

Prepare for the Exam

Key Topics for Chapter 11

Description
Autonomous system numbers
Path Attributes
BGP attribute AS_Path
Address family databases and configuration
Inter-router communication
BGP Single- and Multi-Hop Sessions
BGP session types
eBGP

Prepare for the Exam

Key Topics for Chapter 11 (Cont.)

Description
Basic BGP configuration
Verification of BGP Sessions
Prefix advertisement
BGP Database Processing
BGP Table Fields
BGP summarization techniques
Aggregate address
Aggregate address with **summary-only**

Prepare for the Exam

Key Topics for Chapter 11 (Cont.)

Description
Atomic aggregate
Route aggregation with AS_SET
Multiprotocol BGP for IPv6
IPv6 configuration
IPv6 summarization

Prepare for the Exam

Key Terms for Chapter 11

Key Terms	
Address Family	Loc-RIB table
AS_Path	Optional non-transitive
Atomic aggregate	Optional transitive
Autonomous System (AS)	Path vector routing protocol
eBGP session	Well-known discretionary
iBGP session	Well-known mandatory

Command Reference for Chapter 11

Task	Command Syntax
Initialize the BGP router process	**router bgp** *as-number*
Identify a BGP peer to establish a session with	**neighbor** *ip-address* **remote-as** *as-number*
Disable the automatic IPv4 address family configuration mode	**no bgp default ip4-unicast**
Initialize a specific address family and sub-address family	**address-family** *afi safi*
Activate a BGP neighbor for a specific address family	**neighbor** *ip-address* **activate**
Advertise a network to BGP	**network** *network* **mask** *subnet-mask* **[route-map** *route-map-name***]**
Configure a BGP aggregate IPv4 prefix	**aggregate-address** *network subnet-mask* **[summary-only] [as-set]**
Configure a BGP aggregate IPv6 prefix	**aggregate-address** *prefix/prefix-length* **[summary-only] [as-set]**

Command Reference for Chapter 11 (Cont.)

Task	Command Syntax
Display the contents of the BGP database	**show bgp** *afi safi* **[network] [detailed]**
Display a summary of the BGP table and neighbor peering sessions	**show bgp** *afi safi* **summary**
Display the negotiated BGP settings with a specific peer and the number of prefixes exchanged with that peer	**show bgp** *afi safi* **neighbors** *ip-address*
Display the Adj-RIB-Out BGP table for a specific BGP neighbor	**Show bgp** *afi safi* **neighbor** *ip-address* **advertised routes**


CISCO