CHAPTER 12

Border Gateway Protocol, Route Manipulation, and IP Multicast

This chapter covers the Border Gateway Protocol (BGP), which is used to exchange routes between autonomous systems. It is most frequently used between enterprises and service providers. The "Route Manipulation" section covers route summarization and redistribution of route information between routing protocols. The CCDA should know where redistribution occurs when required by the network design. This chapter also reviews policy-based routing (PBR) as a method to change the destination IP address based on policies. Finally, this chapter covers IP multicast protocols.

"Do I Know This Already?" Quiz

The purpose of the "Do I Know This Already?" quiz is to help you decide whether you need to read the entire chapter. If you intend to read the entire chapter, you do not necessarily need to answer these questions now.

The eight-question quiz, derived from the major sections in the "Foundation Topics" portion of the chapter, helps you determine how to spend your limited study time.

Table 12-1 outlines the major topics discussed in this chapter and the "Do I Know This Already?" quiz questions that correspond to those topics.

Table 12-1 *"Do I Know This Already?" Foundation Topics Section-to-Question Mapping*

Foundation Topics Section	Questions Covered in This Section
BGP	1, 2, 7, 8
Route Manipulation	3, 4
IP Multicast Review	5, 6

CAUTION The goal of self-assessment is to gauge your mastery of the topics in this chapter. If you do not know the answer to a question or you are only partially sure, you should mark this question wrong for the purposes of the self-assessment. Giving yourself credit for an answer you correctly guess skews your self-assessment results and might give you a false sense of security.

1. What protocol do you use to exchange IP routes between autonomous systems?
   a. IGMP
   b. eBGP
   c. IGRP
   d. OSPF
2. What is the current version of BGP?
   a. BGP Version 2
   b. BGP Version 3
   c. BGP Version 4
   d. BGP Version 1
3. Where should routes be summarized?
   a. On the core routers
   b. On the distribution routers
   c. On the access routers
   d. None of the above
4. What is PBR?
   a. Public-Broadcast Routing
   b. Private-Based Routing
   c. Policy-Broadcast Routing
   d. Policy-Based Routing
5. What is IGMP?
   a. Interior Group Management Protocol
   b. Internet Group Management Protocol
   c. Interior Gateway Routing Protocol
   d. Interior Gateway Media Protocol
6. How many bits are mapped from the Layer 3 IPv4 multicast address to a Layer 2 MAC address?
   a. 16 bits
   b. 23 bits
   c. 24 bits
   d. 32 bits

7. What is the administrative distance of eBGP routes?

- **a.** 20
- **b.** 100
- **c.** 110
- **d.** 200

8. What is CIDR?

- **a.** Classful Intradomain Routing
- **b.** Classful Interior Domain Routing
- **c.** Classless Intradomain Routing
- **d.** Classless Interdomain Routing

The answers to the "Do I Know This Already?" quiz appear in Appendix A, "Answers to Chapter 'Do I Know This Already?' Quizzes and Q&A Sections." The suggested choices for your next step are as follows:

- **6 or less overall score—**Read the entire chapter. It includes the "Foundation Topics," "Foundation Summary," and "Q&A" sections.
- **7 or 8 overall score—**If you want more review on these topics, skip to the "Foundation Summary" section and then go to the "Q&A" section. Otherwise, move to the next chapter.

Foundation Topics

The “Foundation Topics” section includes discussions of BGP, PBR, route redistribution, and IP multicast protocols. The “BGP” section covers the characteristics and design of BGP. eBGP exchanges routes between autonomous systems. eBGP is commonly used between enterprises and their service providers.

The section “Route Manipulation” covers how you use PBR to change packets’ destination addresses based on policies. This section also covers route summarization and redistribution of route information between routing protocols.

The section “IP Multicast Review” covers multicast protocols such as IGMP, Cisco Group Management Protocol (CGMP), and Protocol Independent Multicast (PIM).

BGP

This section covers BGP theory and design concepts. The current version of BGP, Version 4, is defined in RFC 1771 (March 1995). BGP is an interdomain routing protocol. What this means is that you use BGP to exchange routing information between autonomous systems. The primary function of BGP is to provide and exchange network-reachability information between domains or autonomous systems. BGP is a path vector protocol that is suited for setting routing policies between autonomous systems. In the enterprise campus architecture, BGP is used in the Internet connectivity module.

BGP is the de facto standard for routing between service providers on the Internet because of its rich features. You can also use it to exchange routes in large internal networks. The Internet Assigned Numbers Authority (IANA) reserved TCP Port 179 to identify the BGP protocol. BGPv4 was created to provide CIDR, a feature that was not present in the earlier versions of BGP. BGP is a path-vector routing protocol; it is neither a distance-vector nor link-state routing protocol.

NOTE RFC 1519 describes CIDR, which provides the capability to forward packets based on IP prefixes only, with no concern for IP address class boundaries. CIDR was created as a means to constrain the growth of the routing tables in the Internet core through the summarization of IP addresses across network class boundaries. The early 1990s saw an increase in the growth of Internet routing tables and a reduction in Class B address space. CIDR provides a way for service providers to assign address blocks smaller than a Class B network but larger than a Class C network.

BGP Neighbors

BGP is usually configured between two directly connected routers that belong to different autonomous systems. Each autonomous system is under different technical administration. BGP is frequently used to connect the enterprise to service providers and to interconnect service providers, as shown in Figure 12-1. The routing protocol within the enterprise could be any interior gateway protocol (IGP). Common IGP choices include RIPv2, EIGRP, Open Shortest Path First (OSPF), and Intermediate System-to-Intermediate System (IS-IS). BGPv4 is the only deployed exterior gateway protocol (EGP). AS numbers are a managed resource allocated by the American Registry of Internet Numbers (ARIN). In IP, the AS numbers 64,512 through 65,535 are allocated to IANA and are designated for private use.

Before two BGP routers can exchange routing updates, they must become established neighbors. After BGP routers establish a TCP connection, exchange information, and accept the information, they become established neighbors and start exchanging routing updates. If the neighbors do not reach an established state, they do not exchange BGP updates. The information exchanged before the neighbors are established includes the BGP version number, AS number, BGP router ID, and BGP capabilities.

Figure 12-1 *BGP Neighbors*

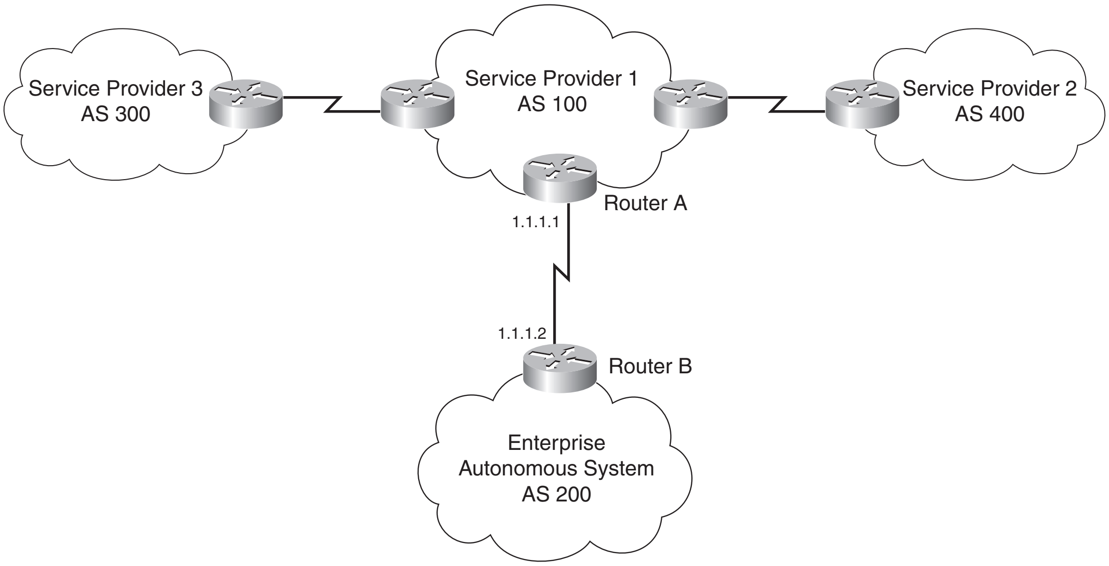

eBGP

eBGP is the term used to describe BGP peering between neighbors in different autonomous systems. As required by RFC 1771, the eBGP peers share a common subnet. In Figure 12-2, all routers speak eBGP with routers in other autonomous systems. Within AS 500, the routers communicate using iBGP, which is covered next.

Figure 12-2 *eBGP Used Between Autonomous Systems*

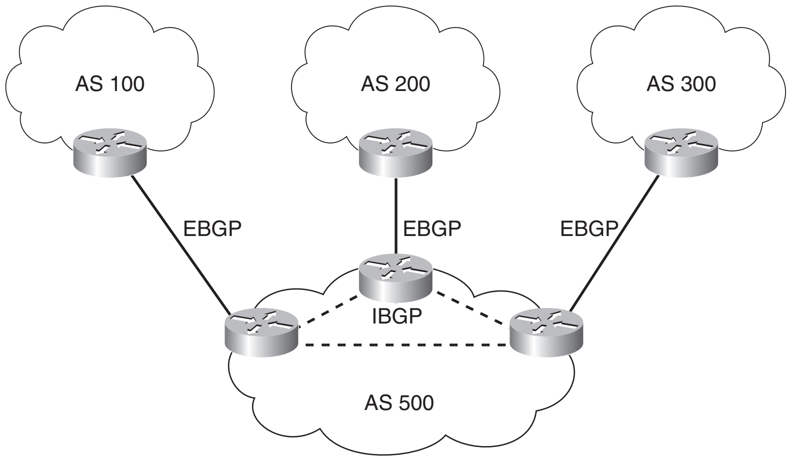

iBGP

iBGP is the term used to describe the peering between BGP neighbors in the same AS. iBGP is used primarily in transit autonomous systems. Transit autonomous systems forward traffic from one external AS to another external AS. If transit autonomous systems did not use iBGP, the eBGP-learned routes would have to be redistributed into an IGP and then redistributed into the BGP process in another eBGP router. Normally the number of eBGP routes is too large for an IGP to handle.

iBGP provides a better way to control the routes within the transit AS. With iBGP, the external route information (attributes) is forwarded. The various IGPs that might be used do not understand or forward BGP attributes, including AS paths, between eBGP routers.

Another use of iBGP is in large corporations where the IGP networks are in smaller independent routing domains along organizational or geographic boundaries. In Figure 12-3, a company has decided to use three independent IGPs: one for the Americas; another for Asia and Australia; and another for Europe, the Middle East, and Africa. Routes are redistributed into an iBGP core.

Figure 12-3 *iBGP in a Large Corporation*

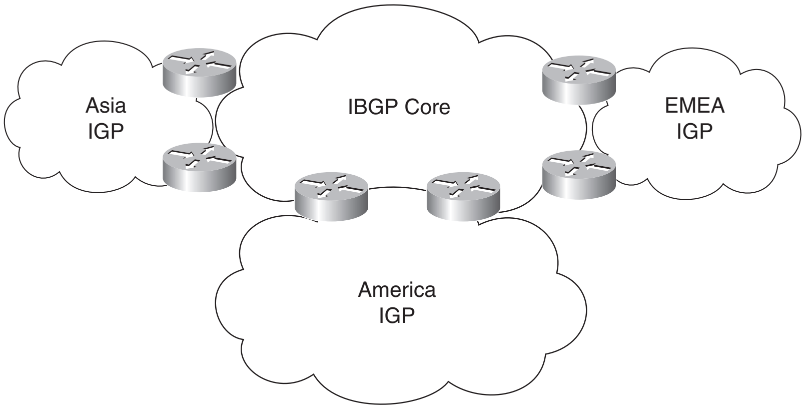

Other Uses of iBGP

The CCDA should know at a high level these other uses for IBGP:

- **Applying policies in the internal AS with the help of BGP path attributes—**BGP path attributes are covered in a later section.
- **QoS Policy Propagation on BGP (QPPB)—**QPPB uses iBGP to spread common QoS parameters from one router to other routers in the network. It classifies packets using IP precedence bits based on BGP community lists, BGP AS paths, and access lists. After packets are classified, QoS features can enforce policies.
- **Multiprotocol BGP peering of Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPN)—**The multiprotocol version of BGP is used to carry MPLS VPN information between all PE routers within a VPN community.

Route Reflectors

iBGP requires that all routers be configured to establish a logical connection with all other iBGP routers. The logical connection is a TCP link between all iBGP-speaking routers. The routers in each TCP link become BGP peers. In large networks, the number of iBGP-meshed peers can become very large. Network administrators can use route reflectors to reduce the number of required mesh links between iBGP peers. Some routers are selected to become the route reflectors to serve several other routers that act as route-reflector clients. Route reflectors allow a router to advertise or reflect routes to clients. The route reflector and its clients form a cluster. All client routers in the cluster peer with the route reflectors within the cluster. The route reflectors also peer with all other route reflectors in the internetwork. A cluster can have more than one route reflector.

In Figure 12-4, without route reflectors, all iBGP routers are configured in an iBGP mesh, as required by the protocol. When Routers A and G become route reflectors, they peer with Routers C and D; Router B becomes a route reflector for Routers E and F. Routers A, B, and G peer among each other.

NOTE The combination of the route reflector and its clients is called a cluster. In Figure 12-4, Routers A, G, C, and D form a cluster. Routers B, E, and F form another cluster.

Figure 12-4 *Route Reflectors*

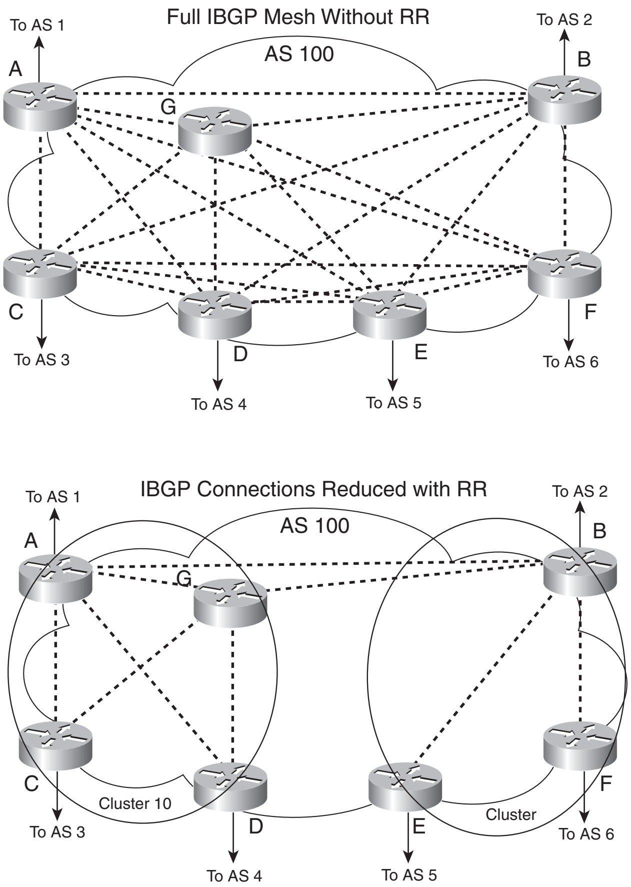

Routers A and G are configured to peer with each other and with Routers B, C, and D. The configuration of Routers C and D is different from the rest; they are configured to peer with Routers A and G only. All route reflectors in the same cluster must have the same cluster ID number.

Router B is the route reflector for the second cluster. Router B peers with Routers A and G and with Routers E and F in its cluster. Routers E and F are route-reflector clients and peer only with

Router B. If Router B goes down, the cluster on the right goes down because no second route reflector is configured.

Confederations

Another method to reduce the iBGP mesh within an AS is BGP confederations. With confederations, the AS is divided into smaller, private autonomous systems, and the whole group is assigned a confederation ID. The private AS numbers or identifiers are not advertised to the Internet but are contained within the iBGP networks. The routers within each private AS are configured with the full iBGP mesh. Each private AS is configured with eBGP to communicate with other semiautonomous systems in the confederation. External autonomous systems see only the AS number of the confederation, and this number is configured with the BGP confederation identifier.

In Figure 12-5, a confederation divides the AS into two.

Figure 12-5 *BGP Confederations*

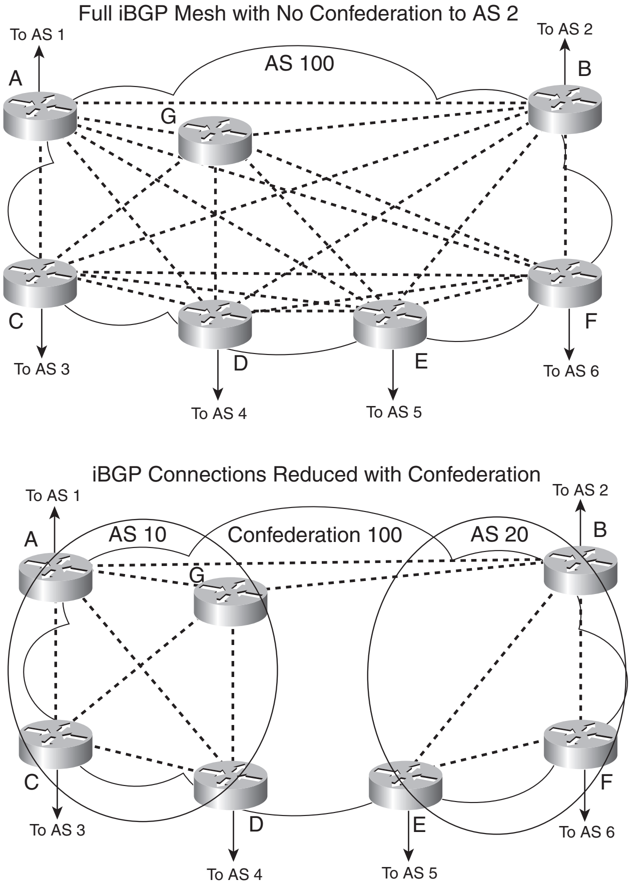

Routers A, B, and G are configured for eBGP between the private autonomous systems. You configure these routers with the **bgp confederation identifier** command. The confederation identifier number is the same for all routers in the network. You use the **bgp confederation peers** command to identify the AS number of other private autonomous systems in the confederation. Because Routers A and G are in AS 10, the peer confederation to Router B is AS 20. Router B is in AS 20, and its peer confederation to Routers A and G is AS 10. Routers C and D are part of AS 10 and peer with each other and with Routers A and G. Routers E and F are part of AS 20 and peer with each other and with Router B.

BGP Administrative Distance

The Cisco IOS Software assigns an administrative distance to eBGP and iBGP routes, as it does with other routing protocols. For the same prefix, the route with the lowest administrative distance is selected for inclusion in the IP forwarding table. Because iBGP-learned routes do not have metrics associated with the route as IGPs (OSPF and EIGRP) do, iBGP-learned routes are less trusted. For BGP, the administrative distances are

- **eBGP routes**—20
- **iBGP routes**—200

BGP Attributes, Weight, and the BGP Decision Process

BGP is a protocol that uses route attributes to select the best path to a destination. This subsection describes BGP attributes, the use of weight to influence path selection, and the BGP decision process.

BGP Path Attributes

BGP uses several attributes for the path-selection process. BGP uses path attributes to communicate routing policies. BGP path attributes include next hop, local preference, AS path, origin, multiexit discriminator (MED), atomic aggregate, and aggregator. Of these, the AS path is one of the most important attributes: It lists the number of AS paths to reach a destination network.

BGP attributes can be categorized as *well-known* or *optional*. Well-known attributes are recognized by all BGP implementations. Optional attributes do not have to be supported by the BGP process; they are used on a test or experimental basis.

Well-known attributes can be further subcategorized as *mandatory* or *discretionary.* Mandatory attributes are always included in BGP update messages. Discretionary attributes might or might not be included in the BGP update message.

Optional attributes can be further subcategorized as *transitive* or *nontransitive*. Routers must advertise the route with transitive attributes to its peers even if it does not support the attribute locally. If the path attribute is nontransitive, the router does not have to advertise the route to its peers.

The following subsections cover each attribute category.

Next-Hop Attribute

The next-hop attribute is the IP address of the next IP hop that will be used to reach the destination. The next-hop attribute is a well-known mandatory attribute. With eBGP, the eBGP peer sets the next hop when it announces the route. Multiaccess networks use the next-hop attribute where there is more than one BGP router.

Local Preference Attribute

The local preference attribute indicates which path to use to exit the AS. It is a well-known discretionary attribute used between iBGP peers and is not passed on to external BGP peers. In Cisco IOS Software, the default local preference is 100. The higher local preference is preferred.

The default local preference is configured on the BGP router with an external path; it then advertises its local preference to internal iBGP peers. Figure 12-6 shows an example of the local preference attribute where Routers B and C are configured with different local preference values. Router A and other iBGP routers then receive routes from both Router B and Router C. Router A prefers using Router C to route Internet packets because it has a higher local preference (400) than Router B (300). The arrows represent the paths taken to go out of the AS.

Figure 12-6 *BGP Local Preference*

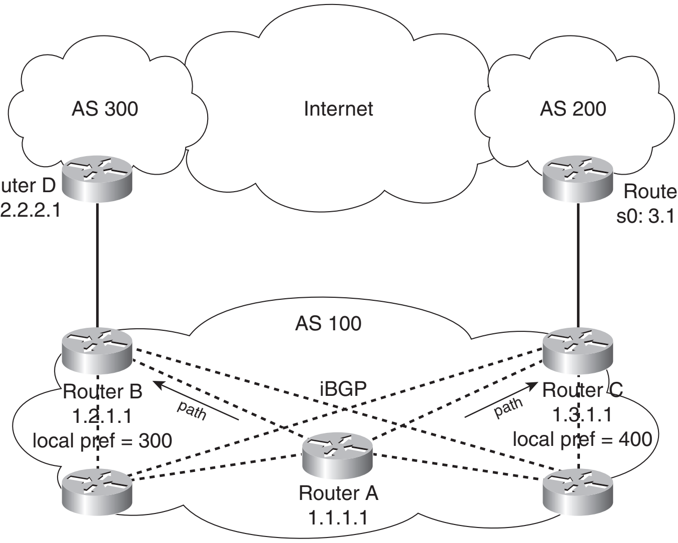

Origin Attribute

Origin is a well-known mandatory attribute that defines the source of the path information. Do not confuse the origin with comparing whether the route is external (eBGP) or internal (iBGP). The origin attribute is received from the source BGP router. There are three types:

- **IGP**—Indicated by an i in the BGP table. Present when the route is learned by way of the **network** statement.
- **EGP**—Indicated by an e in the BGP table. Learned from EGP.
- **Incomplete**—Indicated by a ? in the BGP table. Learned from redistribution of the route.

In terms of choosing a route based on origin, BGP prefers routes that have been verified by an IGP over routes that have been learned from EGP peers, and BGP prefers routes learned from eBGP peers over incomplete paths.

AS Path Attribute

The AS path is a well-known mandatory attribute that contains a list of AS numbers in the path to the destination. Each AS prepends its own AS number to the AS path. The AS path describes all the autonomous systems a packet would have to travel to reach the destination IP network. It is used to ensure that the path is loop-free. When the AS path attribute is used to select a path, the route with the fewest AS hops is preferred. In the case of a tie, other attributes, such as MED, break the tie. Example 12-1 shows the AS path for network 200.50.32.0/19. To reach the destination, a packet must pass autonomous systems 3561, 7004, and 7418. The command **show ip bgp 200.50.32.0** displays the AS path information.

Example 12-1 *AS Path Attribute*

```
Router#show ip bgp 200.50.32.0
BGP routing table entry for 200.50.32.0/19, version 93313535
Paths: (1 available, best #1)
  Not advertised to any peer
  3561 7004 7418
    206.24.241.181 (metric 490201) from 165.117.1.219 (165.117.1.219)
      Origin IGP, metric 4294967294, localpref 100, valid, internal, best
      Community: 2548:182 2548:337 2548:666 3706:153
```

MED Attribute

The MED attribute, also known as a metric, tells external BGP peers the preferred path into the AS when multiple paths into the AS exist. In other words, MED influences which one of many paths a neighboring AS uses to reach destinations within the AS. It is an optional nontransitive attribute carried in eBGP updates. The MED attribute is not used with iBGP peers. The lowest

MED value is preferred, and the default value is 0. Paths received with no MED are assigned a MED of 0. The MED is carried into an AS but does not leave the AS.

Consider the diagram shown in Figure 12-7. With all attributes considered equal, consider that Router C selects Router A as its best path into AS 100 based on Router A's lower router ID (RID). If Router A is configured with a MED of 200, then that will make Router C select Router B as the best path to AS 100. No additional configuration is required on Router B, because the default MED is 0.

Figure 12-7 *MED Attribute*

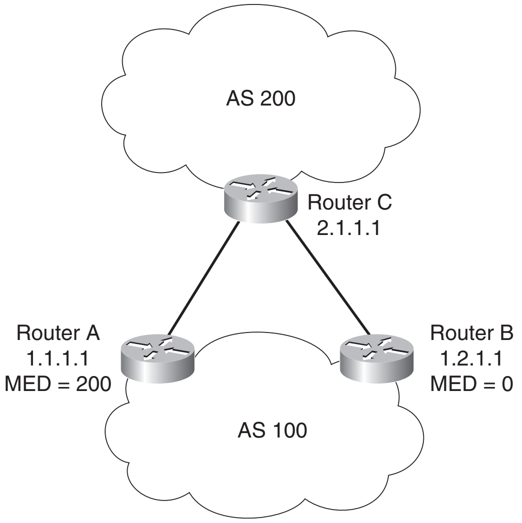

Community Attribute

Although it is not an attribute used in the routing-decision process, the community attribute groups routes and applies policies or decisions (accept, prefer) to those routes. It is a group of destinations that share some common property. The community attribute is an optional transitive attribute of variable length.

Atomic Aggregate and Aggregator Attributes

The atomic aggregate attribute informs BGP peers that the local router used a less specific (aggregated) route to a destination without using a more specific route.

If a BGP router selects a less specific route when a more specific route is available, it must attach the atomic aggregate attribute when propagating the route. The atomic aggregate attribute lets the BGP peers know that the BGP router used an aggregated route. A more specific route must be in the advertising router's BGP table before it propagates an aggregate route.

When the atomic aggregate attribute is used, the BGP speaker has the option to send the aggregator attribute. The aggregator attribute includes the AS number and the IP address of the router that originated the aggregated route. In Cisco routers, the IP address used is the RID of the router that performs the route aggregation. Atomic aggregate is a well-known discretionary attribute, and aggregator is an optional transitive attribute.

Weight

Weight is assigned locally on a router to specify a preferred path if multiple paths exist out of a router for a destination. Weights can be applied to individual routes or to all routes received from a peer. Weight is specific to Cisco routers and is not propagated to other routers. The weight value ranges from 0 to 65,535. Routes with a higher weight are preferred when multiple routes exist to a destination. Routes that are originated by the local router have a default weight of 32,768.

You can use weight instead of local preference to influence the selected path to external BGP peers. The difference is that weight is configured locally and is not exchanged in BGP updates. On the other hand, the local preference attribute is exchanged between iBGP peers and is configured at the gateway router.

When the same destinations are advertised from both Router B and Router C, as shown in Figure 12-8, Router A prefers the routes from Router C over Router B because the routes received from Router C have a larger weight (600) locally assigned.

Figure 12-8 *BGP Weight*

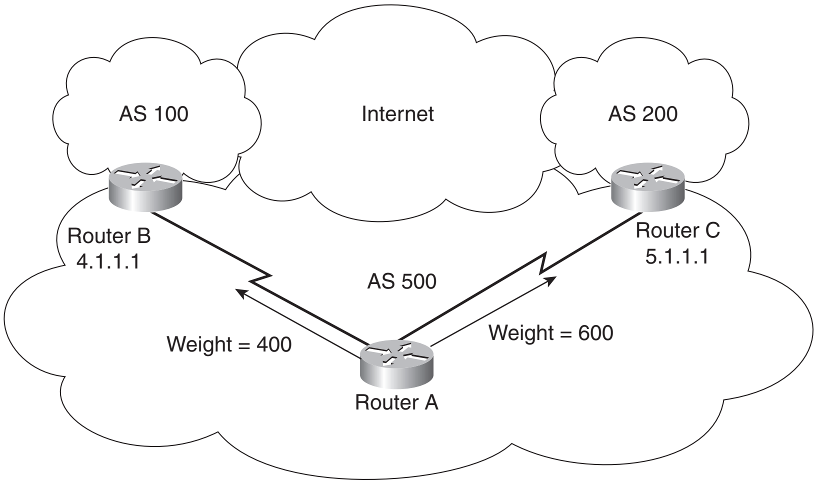

BGP Decision Process

By default, BGP selects only a single path to reach a specific destination (unless you specify maximum paths). The Cisco implementation of BGP uses a simple decision process. When the path is selected, BGP puts the selected path in its routing table and propagates the path to its neighbors.

To select the best path to a destination, Cisco routers running BGP use the following algorithm in the following order:

1. If the specified next hop is inaccessible, drop the path.
2. If the path is internal, synchronization is enabled, and the path is not in the IGP, drop the path.
3. Prefer the path with the largest weight. (This step is Cisco-specific, and weight is localized to the router.)
4. Prefer the path with the largest local preference. iBGP uses this path only to reach the preferred external BGP router.
5. Prefer the path that was locally originated via a **network** or **aggregate** BGP subcommand or through redistribution from an IGP. Local paths sourced by **network** or **redistribute** commands are preferred over local aggregates sourced by the **aggregate-address** command. (This step is Cisco-specific.)
6. If no route was originated, prefer the route that has the shortest AS path. (This step is Cisco-specific.)
7. If all paths have the same AS path length, prefer the path with the lowest origin type. Paths with an origin type of IGP (lower) are preferred over paths originated from an EGP such as BGP, and EGP origin is preferred over a route with an incomplete origin. (This step is Cisco-specific.)
8. If the origin codes are the same, prefer the path with the lowest MED attribute. An eBGP peer uses this attribute to select a best path to the AS. (This step is a tiebreaker, as described in the RFC that defines the BGP.)
9. If the paths have the same MED, prefer the external (eBGP) path over the internal (iBGP) path. (This step is Cisco-specific.)
10. If the paths are still the same, prefer the path through the closest IGP neighbor (best IGP metric). (This step is a tiebreaker, as described in the RFC that defines the BGP.)
11. Prefer the path with the BGP neighbor with the lowest router ID. (The RFC that defines the BGP describes the router ID.)

After BGP decides on a best path, it marks it with a > sign in the **show ip bgp** table and adds it to the IP routing table.

BGP Summary

The characteristics of BGP follow:

- BGP is an exterior gateway protocol (EGP) used in routing in the Internet. It is an interdomain routing protocol.
- BGP is a path vector routing protocol suited for strategic routing policies.
- It uses TCP port 179 to establish connections with neighbors.
- BGPv4 implements CIDR.
- eBGP is used for external neighbors. It is used between different autonomous systems.
- iBGP is used for internal neighbors. It is used within an AS.
- BGP uses several attributes in the routing-decision algorithm.
- It uses confederations and route reflectors to reduce BGP peering overhead.
- The MED (metric) attribute is used between autonomous systems to influence inbound traffic.
- Weight is used to influence the path of outbound traffic from a single router, configured locally.

Route Manipulation

This section covers PBR, route summarization, and route redistribution. You can use PBR to modify the next hop of packets from what is selected by the routing protocol. PBR is useful when the traffic engineering of paths is required. Routes are summarized to reduce the size of routing tables and at network boundaries. Redistribution between routing protocols is required to inject route information from one routing protocol to another. The CCDA must understand the issues with the redistribution of routes.

PBR

You can use PBR to modify the next-hop address of packets or to mark packets to receive differential service. Routing is based on destination addresses; routers look at the routing table to determine the next-hop IP address based on a destination lookup. PBR is commonly used to modify the next-hop IP address based on the source address. You can also use PBR to mark the IP precedence bits in outbound IP packets so that you can apply quality-of-service (QoS) policies. In Figure 12-9, Router A exchanges routing updates with routers in the WAN. The routing protocol might select Serial 0 as the preferred path for all traffic because of the higher bandwidth. The company might have business-critical systems that use the T1 but does not want systems on

Ethernet 1 to affect WAN performance. You can configure PBR on Router A to force traffic from Ethernet 1 out on Serial 1.

Figure 12-9 *Policy-Based Routing*

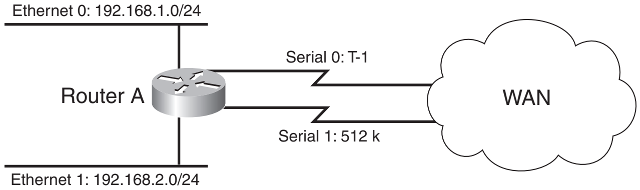

Route Summarization

Large networks can grow very quickly from 500 routes to 1000, to 2000, and so on. Network IP addresses should be allocated to allow for route summarization. Route summarization reduces the amount of route traffic on the network and unnecessary route computation. Route summarization also allows the network to scale as a company grows.

The recommended location for route summarization is to summarize at the distribution layer of the network topology. Figure 12-10 shows a hierarchical network. It has a network core, regional distribution routers, and access routes for sites.

Figure 12-10 *Route Summarization*

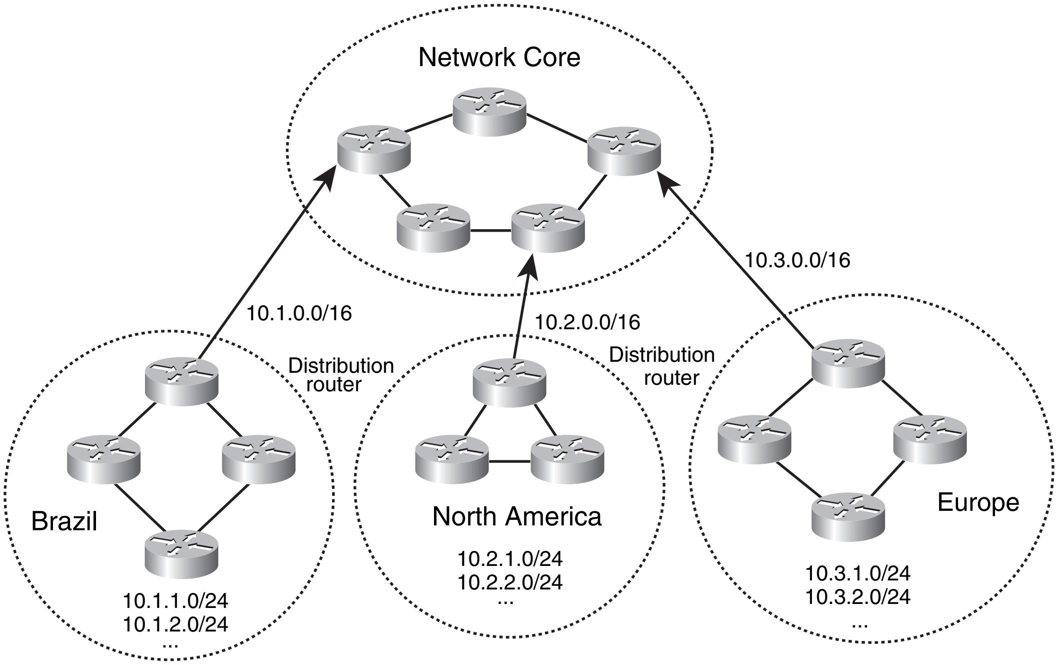

All routes in Brazil are summarized with a single 10.1.0.0/16 route. The North America and European routes are also summarized with 10.2.0.0/16 and 10.3.0.0/16, respectively. Routers in Europe only need to know the summarized route to get to Brazil and North America, and vice versa. Again, design best practices are to summarize at the distribution toward the core. The core only needs to know the summarized route of the regional areas.

Route Redistribution

You configure the redistribution of routing protocols on routers that reside at the Service Provider Edge of the network. These routers exchange routes with other autonomous systems. Redistribution is also done on routers that run more than one routing protocol. Here are some reasons to do redistribution:

- Migration from an older routing protocol to a new routing protocol.
- Mixed-vendor environment in which Cisco routers might be using EIGRP and other vendor routers might be using OSPF.
- Different administrative domain between company departments using different routing protocols.
- Mergers and acquisitions in which the networks initially need to communicate. In this example two different EIGRP processes might exist.

Figure 12-11 shows an example of the exchange of routes between two autonomous systems. Routes from AS 100 are redistributed into BGP on Router A. Routes from AS 200 are redistributed into BGP on Router B. Then, Routers A and B exchange BGP routes. Router A and Router B also implement filters to redistribute only the desired networks.

Figure 12-11 *Redistribution of BGP Routes*

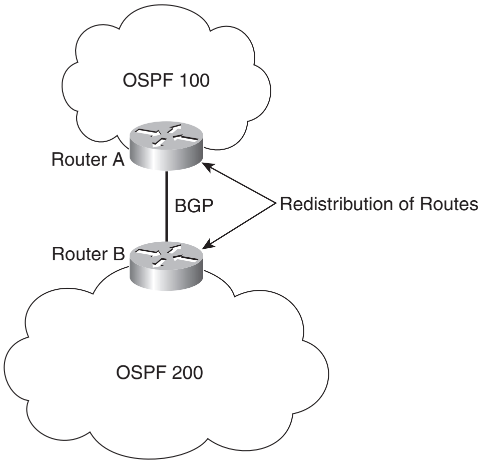

A company might also acquire another company that might be running another routing protocol. Figure 12-12 shows a network that has both OSPF and EIGRP routing protocols. Routers A and B perform redistribution between OSPF and EIGRP. Both routers must filter routes from OSPF before redistributing them into EIGRP and filter routes from EIGRP before redistributing them into OSPF. This setup prevents route feedback.

Figure 12-12 *Redistribution Between IGPs*

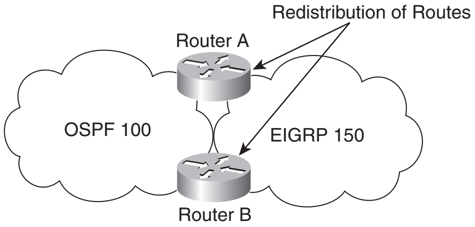

Route feedback occurs when a routing protocol learns routes from another routing protocol and then announces the routes to the other routing protocol. In Figure 12-12, OSPF should not announce the routes it learned from EIGRP, and EIGRP should not announce the routes it learned from OSPF.

You can use access lists, distribution lists, and route maps when redistributing routes. You can use these methods to specify (select) routes for redistribution, to set metrics, or to set other policies for the routes. They are also used to control routes' redistribution direction. Redistribution can be accomplished by two methods:

- Two-way redistribution
- One-way redistribution

In two-way redistribution, routing information is exchanged between both routing protocols. No static routes are used in this exchange. Route filters are used to prevent routing loops. Routing loops can be caused by one route protocol redistributing routes that were learned from a second route protocol back to that second routing protocol.

One-way redistribution only allows redistribution from one routing protocol to another. Normally it is used in conjunction with a default or static route at the edge of a network. Figure 12-13 shows an example of one-way redistribution. The routing information from the WAN routes is redistributed into the campus. But campus routes are not redistributed out to the WAN. The WAN routers use a default gateway to get back to the campus.

Figure 12-13 *One-Way Route Redistribution*

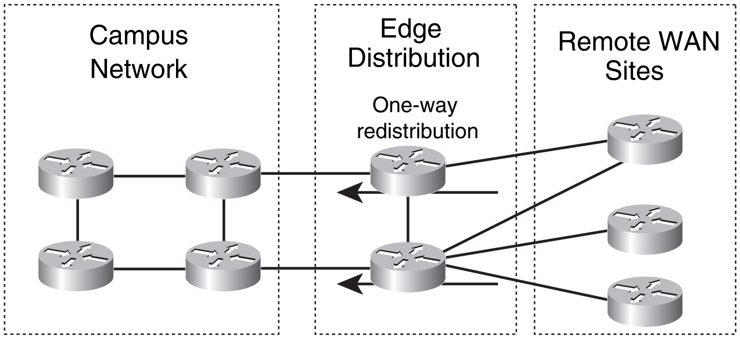

Other locations for one-way redistribution are from building access networks, BGP routes or static routes into the IGP, and from VPN static routes into the IGP.

Default Metric

There is a default metric of 0 when redistributing routes into RIPv2, IS-IS, and EIGRP. You should configure the metric of the redistributed routes to a metric other than 0. You can configure the metric in the **redistribution** command or configure a default metric. You can also use the command in OSPF. IS-IS does not use the **default-metric** command. The **default-metric** command has the following syntax for EIGRP:

```
default-metric bandwidth delay reliability load mtu
```

OSPF Redistribution

This subsection reviews a few things you need to remember when designing a network that will redistribute with OSPF.

When redistributing routes into OSPF, use the **subnets** keyword to permit subnetted routes to be received. If you do not use it, only the major network route is redistributed, without any subnetworks. In other words, OSPF performs automatic summarization to IP classful network values.

By default, redistributed routes are classified as external Type 2 (E2) in OSPF. You can use the **metric-type** keyword to change the external route to an external Type 1 (E1). The network design can take into account the after-redistribution cost (Type 2) or the after-redistribution cost plus the path's cost (Type 1).

In Figure 12-14, Router B is configured to perform mutual redistribution between EIGRP 100 and OSPF process ID 50. In this example, you can use route maps and access lists to prevent routing loops. The route maps permit or deny the networks that are listed in the access lists. The **subnets** keyword redistributes every subnet in EIGRP into OSPF. This book does not cover exact configurations.

Figure 12-14 *OSPF and EIGRP Redistribution*

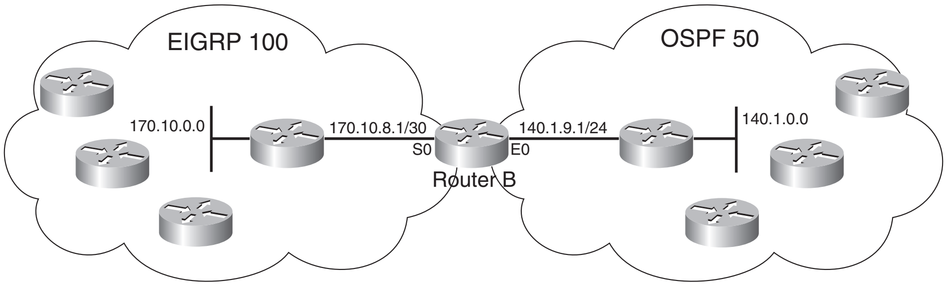

IP Multicast Review

With multicast, packets are sent to a multicast group, which is identified with an IP multicast address. Multicast supports the transmission of IP packets from one source to multiple hosts. Packets with unicast addresses are sent to one device, and broadcast addresses are sent to all hosts; packets with multicast addresses are sent to a group of hosts.

Multicast Addresses

Multicast addressing uses Class D addresses from the IPv4 protocol. Class D addresses range from 224.0.0.0 to 239.255.255.255. IANA manages multicast addresses.

Routing protocols (RIPv2, EIGRP, and OSPF) use multicast addresses to speak to their neighbors. For example, OSPF routers use 224.0.0.6 to speak to the designated router (DR) in a multiaccess network. Class D multicast addresses range from 224.0.0.0 to 239.255.255.255. Multicast addresses in the range of 224.0.0.1 to 224.255.255.255 are reserved for special addresses or network protocol on a multiaccess link. RFC 2365 reserves multicast addresses in the range of 239.192.000.000 to 239.251.255.255 for organization-local scope. Similarly, 239.252.000.000 to 239.252.255.255, 239.254.000.000 to 239.254.255.255, and 239.255.000.000 to 239.255.255.255 are reserved for site-local scope.

Table 12-2 lists some well-known and multicast address blocks.

Table 12-2 *Multicast Addresses*

Multicast Address	Description
224.0.0.0/24	Local network control block
224.0.0.1	All hosts or all systems on this subnet
224.0.0.2	All multicast routers
224.0.0.4	Distance-Vector Multicast Routing Protocol (DVMRP) routers
224.0.0.5	All OSPF routers
224.0.0.6	All OSPF DR routers
224.0.0.9	RIPv2 routers
224.0.0.10	EIGRP routers
224.0.0.13	All PIM routers
224.0.1.0/24	Internetwork control block
224.0.1.39	Rendezvous point (RP) announce
224.0.1.40	RP discovery
224.0.2.0 to 224.0.255.0	Ad hoc block
239.000.000.000 to 239.255.255.255	Administratively scoped
239.192.000.000 to 239.251.255.255	Organization-local scope
239.252.000.000 to 239.254.255.255	Site-local scope

Layer 3 to Layer 2 Mapping

Multicast-aware Ethernet, Token Ring, and Fiber Distributed Data Interface (FDDI) network interface cards use the reserved IEEE 802 address 0100.5e00.0000 for multicast addresses at the MAC layer. This includes Fast Ethernet and Gigabit Ethernet. Notice that for the address, the high-order byte 0x01 has the low-order bit set to 1. This bit is the Individual/Group (I/G) bit. It signifies whether the address is an individual address (0) or a group address (1). Hence, for multicast addresses, this bit is set to 1.

Ethernet interfaces map the lower 23 bits of the IP multicast address to the lower 23 bits of the MAC address 0100.5e00.0000. As an example, the IP multicast address 224.0.0.2 is mapped to the MAC layer as 0100.5e00.0002. Figure 12-15 shows another example looking at the bits of multicast IP 239.192.44.56. The IP address in hexadecimal is EF:C0:2C:38. The lower 23 bits get mapped into the lower 23 bits of the base multicast MAC to produce the multicast MAC address 01:00:5E:40:2C:38.

Figure 12-15 *Mapping of Multicast IP Addressing to MAC Addresses*

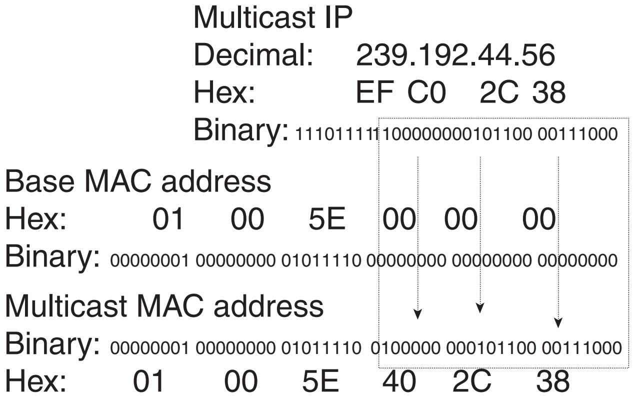
Multicast IP
Decimal: 239.192.44.56
Hex: EF C0 2C 38
Binary: 11101111 10000000 00101100 00111000
Base MAC address
Hex: 01 00 5E 00 00 00
Binary: 00000001 00000000 01011110 00000000 00000000 00000000
Multicast MAC address
Binary: 00000001 00000000 01011110 0100000 000101100 00111000
Hex: 01 00 5E 40 2C 38

IGMP

IGMP is the protocol used in multicast implementations between the end hosts and the local router. RFC 2236 describes IGMP Version 2 (IGMPv2). RFC 3376 describes IGMP Version 3 (IGMPv3). RFC 1112 describes the first version of IGMP.

IP hosts use IGMP to report their multicast group memberships to routers. IGMP messages use IP protocol number 2. IGMP messages are limited to the local interface and are not routed.

IGMPv1

The first RFC describing IGMP (RFC 1112), written in 1989, describes the host extensions for IP multicasting. IGMPv1 provides simple message types for communication between hosts and routers. These messages are

- **Membership query**—Sent by the router to check whether a host wants to join a multicast group
- **Membership report**—Sent by the host to join a multicast group in the segment

The problem with IGMPv1 is the latency involved for a host to leave a group. With IGMPv1, the router sends membership queries periodically; a host must wait for the membership-query message to leave a group. The query interval is 60 seconds, and it takes three query intervals (3 minutes) for a host to leave the group.

IGMPv2

IGMPv2 improves over IGMPv1 by allowing faster termination or leaving of multicast groups.

IGMPv2 has three message types, plus one for backward compatibility:

- **Membership query**—Sent by the router to check whether a host wants to join a group.
- **Version 2 membership report**—A message sent to the group address with the multicast group members (IP addresses). It is sent to by hosts to join and remain in multicast groups on the segment.
- **Version 2 leave group**—Sent by the hosts to indicate that a host will leave a group; it is sent to destination 224.0.0.2. After the host sends the leave group message, the router responds with a group-specific query.
- **Version 1 membership report**—For backward compatibility with IGMPv1 hosts.

You enable IGMP on an interface when you configure a multicast routing protocol, such as PIM. You can configure the interface for IGMPv1, IGMPv2 or IGMPv3.

IGMPv3

IGMPv3 provides the extensions required to support source-specific multicast (SSM). It is designed to be backward-compatible with both prior versions of IGMP.

IGMPv3 has two message types, plus three for backward compatibility:

- **Membership query**—Sent by the router to check that a host wants to join a group.
- **Version 3 membership report**—A message sent to the group address with the multicast group members (IP addresses). It is sent by hosts to request and remain in multicast groups on the segment.
- **Version 2 membership report**—A message sent to the group address with the multicast group members (IP addresses). It is sent by hosts to request and remain in multicast groups on the segment. This message is used for backward compatibility with IGMPv2 hosts.
- **Version 2 leave group**—Sent by the hosts to indicate that a host will leave a group, to destination 224.0.0.2. The message is sent without having to wait for the IGMPv2 membership report message. This message is used for backward compatibility with IGMPv2 hosts.
- **Version 1 membership report**—This message is used for backward compatibility with IGMPv1 hosts.

You enable IGMP on an interface when you enable a multicast routing protocol, such as PIM. You can configure the interface for IGMPv1, IGMPv2, or IGMPv3.

CGMP

CGMP is a Cisco-proprietary protocol implemented to control multicast traffic at Layer 2. Because a Layer 2 switch is unaware of Layer 3 IGMP messages, it cannot keep multicast packets from being sent to all ports.

As shown in Figure 12-16, with CGMP the LAN switch can speak with the IGMP router to find out the MAC addresses of the hosts that want to receive the multicast packets. With CGMP, switches distribute multicast sessions only to the switch ports that have group members.

Figure 12-16 *CGMP*

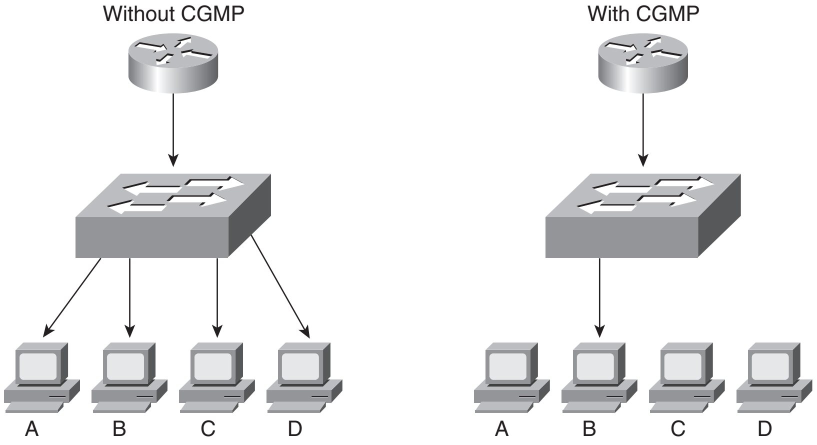

When a router receives an IGMP report, it processes the report and then sends a CGMP message to the switch. The switch can then forward the multicast messages to the port with the host receiving multicast traffic. CGMP fast-leave processing allows the switch to detect IGMP Version 2 leave messages sent by hosts on any of the switch ports. When a host sends the IGMPv2 leave message, the switch can then disable multicasting for the port.

IGMP Snooping

IGMP snooping is another way for switches to control multicast traffic at Layer 2. It listens to IGMP messages between the hosts and routers. If a host sends an IGMP query message to the router, the switch adds the host to the multicast group and permits that port to receive multicast traffic. The port is removed from multicast traffic if the host sends an IGMP leave message to the router. The disadvantage of IGMP snooping is that it has to process every IGMP control message, which can impact the CPU utilization of the switch.

Sparse Versus Dense Multicast Routing Protocols

IP multicast traffic for a particular (source, destination group) multicast pair is transmitted from the source to the receivers using a spanning tree from the source that connects all the hosts in the group. Each destination host registers itself as a member of interesting multicast groups through the use of IGMP. Routers keep track of these groups dynamically and build distribution trees that chart paths from each sender to all receivers. IP multicast routing protocols follow two approaches.

The first approach assumes that the multicast group members are densely distributed throughout the network (many of the subnets contain at least one group member) and that bandwidth is plentiful. The approach with dense multicast routing protocols is to flood the traffic throughout the network and then, at the request of receiving routers, stop the flow of traffic on branches of the network that have no members of the multicast group. Multicast routing protocols that follow this technique of flooding the network include DVMRP, Multicast Open Shortest Path First (MOSPF), and Protocol-Independent Multicast-Dense Mode (PIM-DM).

The second approach to multicast routing assumes that multicast group members are sparsely distributed throughout the network and that bandwidth is not necessarily widely available. Sparse mode does not imply that the group has few members, just that they are widely dispersed. The approach with sparse multicast routing protocols is to not send traffic until it is requested by the receiving routers or hosts. Multicast routing protocols of this type are Core-Based Trees (CBT) and Protocol-Independent Multicast-Sparse Mode (PIM-SM). CBT is not widely deployed and is not discussed in this book.

Multicast Source and Shared Trees

Multicast distribution trees control the path that multicast packets take to the destination hosts. The two types of distribution trees are source and shared. With *source* trees, the tree roots from the source of the multicast group and then expands throughout the network in spanning-tree fashion to the destination hosts. Source trees are also called shortest-path trees (SPT) because they create paths without having to go through a rendezvous point (RP). The drawback is that all routers through the path must use memory resources to maintain a list of all multicast groups. PIM-DM uses a source-based tree.

Shared trees create the distribution tree's root somewhere between the network's source and receivers. The root is called the RP. The tree is created from the RP in spanning-tree fashion with no loops. The advantage of shared trees is that they reduce the memory requirements of routers in the multicast network. The drawback is that initially the multicast packets might not take the best paths to the receivers because they need to pass through the RP. After the data stream begins to flow from sender to RP to receiver, the routers in the path optimize the path automatically to remove any unnecessary hops. The RP function consumes significant memory on the assigned router. PIM-SM uses an RP.

PIM

PIM comes in two flavors: *sparse mode* (PIM-SM) and *dense mode* (PIM-DM). The first uses shared trees and RPs to reach widely dispersed group members with reasonable protocol bandwidth efficiency. The second uses source trees and reverse path forwarding (RPF) to reach relatively close group members with reasonable processor and memory efficiency in the network devices of the distribution trees.

With RPF, received multicast packets are forwarded out all other interfaces, allowing the data stream to reach all segments. If no hosts are members of a multicast group on any of the router's attached or downstream subnets, the router sends a prune message up the distribution tree (the reverse path) to tell the upstream router not to send packets for the multicast group. So, the analogy for PIM-DM is the push method for sending junk mail, and the intermediate router must tell upstream devices to stop sending it.

PIM-SM

PIM-SM is defined in RFC 2362 (experimental). PIM-SM assumes that no hosts want to receive multicast traffic unless specifically requested. In PIM-SM, a router is selected as the RP. The RP gathers the information from senders and makes the information available to receivers. Routers with receivers have to register with the RP. The end-host systems request multicast group membership using IGMP with their local routers. The routers serving the end systems then register as traffic receivers with the RPs for the specified group in the multicast network.

Joining PIM-SM

With PIM-SM, DRs on end segments receive IGMP query messages from hosts wanting to join a multicast group. The router checks whether it is already receiving the group for another interface. If it is receiving the group, the router adds the new interface to the table and sends membership reports periodically on the new interface.

If the multicast group is not in the multicast table, the router adds the interface to the multicast table and sends a join message to the RP with multicast address 224.0.0.13 (all PIM routers) requesting the multicast group.

Pruning PIM-SM

When a PIM-SM does not have any more multicast receiving hosts or receiving routers out any of its interfaces, it sends a prune message to the RP. The prune message includes the group to be pruned or removed.

PIM DR

A designated router is selected in multiaccess segments running PIM. The PIM DR is responsible for sending join, prune, and register messages to the RP. The PIM router with the highest IP address is selected as the DR.

Auto-RP

Another way to configure the RP for the network is to have the RP announce its services to the PIM network. This process is called auto-RP. Candidate RPs send their announcements to RP mapping agents with multicast address 224.0.1.39 (**cisco-rp-announce**). RP mapping agents are also configured. In smaller networks, the RP can be the mapping agent. Configured RP mapping agents listen to the announcements. The RP mapping agent then selects the RP for a group based on the highest IP address of all the candidate RPs. The RP mapping agents then send RP-discovery messages to the rest of the PIM-SM routers in the internetwork with the selected RP-to-group mappings.

PIMv2 Bootstrap Router

Instead of using auto-RP, you can configure a PIMv2 bootstrap router (BSR) to automatically select an RP for the network. The RFC for PIM Version 2, RFC 2362, describes BSR. With BSR, you configure BSR candidates (C-BSR) with priorities from 0 to 255 and a BSR address. C-BSRs exchange bootstrap messages. Bootstrap messages are sent to multicast IP 224.0.0.13 (all PIM routers). If a C-BSR receives a bootstrap message, it compares it with its own. The largest priority C-BSR is selected as the BSR.

After the BSR is selected for the network, it collects a list of candidate RPs. The BSR selects RP-to-group mappings, which is called the RP set, and distributes the selected RPs using bootstrap messages sent to 224.0.0.13 (all PIM routers).

DVMRP

RFC 1075 describes DVMRP. It is the primary multicast routing protocol used in the multicast backbone (MBONE). The MBONE is used in the research community.

DVMRP operates in dense mode using RPF by having routers send a copy of a multicast packet out all paths. Routers that receive the multicast packets then send prune messages back to their upstream neighbor router to stop a data stream if no downstream receivers of the multicast group exist (either receiving routers or hosts on connected segments). DVMRP implements its own unicast routing protocol, similar to RIP, based on hop counts. DVMRP has a 32 hop-count limit. DVMRP does not scale suboptimally. Cisco's support of DVMRP is partial; DVMRP networks are usually implemented on UNIX machines running the **mrouted** process. A DVMRP tunnel is typically used to connect to the MBONE DVMRP network.

IPv6 Multicast Addresses

IPv6 retains the use and function of multicast addresses as a major address class. IPv6 prefix FF00::/8 is allocated for all IPv6 multicast addresses. IPv6 multicast addresses are described in RFC 2373. EIGRP for IPv6, OSPFv3, and RIPng routing protocols use multicast addresses to communicate between router neighbors.

The format of the IPv6 multicast address is described in Chapter 8, "Internet Protocol Version 6." The common multicast addresses are repeated in Table 12-3.

Table 12-3 *Well-Known Multicast Addresses*

Multicast Address	Multicast Group
FF01::1	All nodes (node-local)
FF02::1	All nodes (link-local)
FF01::2	All routers (node-local)
FF02::2	All routers (link-local)
FF02::5	OSPFv3 routers
FF02::6	OSPFv3 designated routers
FF02::9	Routing Information Protocol (RIPng)
FF02::A	EIGRP routers
FF02::B	Mobile agents
FF02::C	DHCP servers/relay agents
FF02::D	All PIM routers

References and Recommended Readings

Border Gateway Protocol. http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm.

Chandra, R., P. Traina, and T. Li. RFC 1997, *BGP Communities Attribute*. Available from http://www.ietf.org/rfc.

Deering, S. RFC 1112, *Host Extensions for IP Multicasting*. Available from http://www.ietf.org/rfc.

Doyle, J. and J. Carroll. *Routing TCP/IP,* Volume I, Second Edition. Indianapolis: Cisco Press, 2005.

Doyle, J. and J. Carroll. *Routing TCP/IP,* Volume II. Indianapolis: Cisco Press, 2001.

Estrin, D., D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, and L. Wei. RFC 2362, *Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification* (experimental). Available from http://www.ietf.org/rfc.

Fenner, W. RFC 2236, *Internet Group Management Protocol, Version 2*. Available from http://www.ietf.org/rfc.

Fuller, V., T. Li, J. Yu, and K. Varadhan. RFC 1519, *Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy.* Available from http://www.ietf.org/rfc.

Halabi, S. *Internet Routing Architectures*. Indianapolis: Cisco Press, 2000.

"Internet Protocol (IP) Multicast Technology Overview" (white paper). Available from http://www.cisco.com/en/US/products/ps5763/products_white_paper0900aecd804d5fe6.shtml.

Meyer, D. RFC 2365, *Administratively Scoped IP Multicast*. Available from http://www.ietf.org/rfc.

Rekhter, Y. and T. Li. RFC 1771, *A Border Gateway Protocol 4 (BGP-4)*. Available from http://www.ietf.org/rfc.

Waitzman, D., C. Partride, and S. Deering. RFC 1075, *Distance Vector Multicast Routing Protocol*. Available from http://www.ietf.org/rfc.

Williamson, B. *Developing IP Multicast Networks*. Indianapolis: Cisco Press, 1999.

Foundation Summary

The "Foundation Summary" section of each chapter lists the most important facts from the chapter. Although this section does not list every fact from the chapter that will be on the CCDA exam, a well-prepared CCDA candidate should at a minimum know all the details in each "Foundation Summary" before taking the exam.

This chapter covered the following topics that you need to master for the CCDA exam:

- **BGP**—The characteristics and design of BGP.
- **Route manipulation**—How you use PBR to change the destination address of packets based on policies. This material also covers route summarization and the redistribution of routes between routing protocols.
- **IP multicast protocols**—Multicast protocols such as IGMP, CGMP, and PIM.

The material summarized next can help you review some of these topical areas.

BGP Summary

The characteristics of BGP follow:

- BGP is an exterior gateway protocol (EGP) used in routing in the Internet. It is an interdomain routing protocol.
- BGP is a path vector routing protocol suited for strategic routing policies.
- BGP uses TCP Port 179 to establish connections with neighbors.
- BGPv4 implements CIDR.
- eBGP is for external neighbors. It's used between separate autonomous systems.
- iBGP is for internal neighbors. It's used within an AS.
- BGP uses several attributes in the routing-decision algorithm.
- BGP uses confederations and route reflectors to reduce BGP peering overhead.
- The MED (metric) attribute is used between autonomous systems to influence inbound traffic.
- Weight is used to influence the path of outbound traffic from a single router, configured locally.

Route Redistribution

Route redistribution can occur

- In mixed vendor environments, where Cisco routers might be using EIGRP and other vendor routers using OSPF.
- In migrations from older routing protocol.
- In different administrative domains.
- From static routes and BGP routes into IGP.
- From VPN static routes into IGP.
- Between campus core and WAN routers.
- From selected building access protocols.

IP Multicast

Table 12-4 summarizes IP multicast protocols.

Table 12-4 *IP Multicast Protocols*

Multicast Protocol	Description
IGMP	Internet Group Management Protocol. Used by IP hosts to report their multicast group memberships to routers.
CGMP	Cisco Group Management Protocol. Used to control multicast traffic at Layer 2.
IGMP snooping	Another method used to control multicast traffic at Layer 2.
PIM	Protocol Independent Multicast. IP multicast routing protocol.
DVMRP	Distance-Vector Multicast Routing Protocol. Primary multicast routing protocol used in the MBONE.

Table 12-5 summarizes IP multicast addresses.

Table 12-5 *IP Multicast Addresses*

Multicast Address	Description
224.0.0.0/24	Local network control block
224.0.0.1	All hosts or all systems on this subnet
224.0.0.2	All routers on this subnet
224.0.0.4	DVMRP routers

Table 12-5 *IP Multicast Addresses (Continued)*

Multicast Address	Description
224.0.0.5	All OSPF routers
224.0.0.6	All OSPF DR routers
224.0.0.9	RIPv2 routers
224.0.0.10	EIGRP routers
224.0.0.13	All PIM routers
224.0.1.0/24	Internetwork control block
224.0.1.39	RP announce
224.0.1.40	RP discovery
224.0.2.0 to 224.0.255.0	Ad hoc block
239.000.000.000 to 239.255.255.255	Administratively scoped
239.192.000.000 to 239.251.255.255	Organization-local scope
239.252.000.000 to 239.254.255.255	Site-local scope

Table 12-6 shows IPv6 multicast addresses.

Table 12-6 *IPv6 Multicast Addresses*

Multicast Address	Multicast Group
FF01::1	All nodes (node-local)
FF02::1	All nodes (link-local)
FF01::2	All routers (node-local)
FF02::2	All routers (link-local)
FF02::5	OSPFv3 routers
FF02::6	OSPFv3 designated routers
FF02::9	Routing Information Protocol (RIPng)
FF02::A	EIGRP routers
FF02::B	Mobile agents
FF02::C	DHCP servers/relay agents
FF02::D	All PIM routers

Q&A

As mentioned in the Introduction, you have two choices for review questions: here in the book or the exam questions on the CD-ROM. The answers to these questions appear in Appendix A.

For more practice with exam format questions, use the exam engine on the CD-ROM.

1. True or false: You use iBGP to exchange routes between different autonomous systems.
2. True or false: BGP Version 4 includes support for CIDR.
3. True or false: eBGP and iBGP redistribute automatically on a router if the BGP peers are configured with the same AS number.
4. Use _______ to modify the next hop of packets based on source IP address.
5. eBGP routes have an administrative distance of ____ , and iBGP routes have an administrative distance of ____.
6. True or false: IGMP snooping and CGMP are methods to reduce the multicast traffic at Layer 2.
7. True or false: PIM has a 32 hop-count limit.
8. True or false: PIM-SM routers use the multicast 224.0.0.13 address to request a multicast group to the RP.
9. True or false: AS path is the only attribute BGP uses to determine the best path to the destination.
10. List three IP routing protocols that use multicast addresses to communicate with their neighbors.
11. What IPv6 multicast address does EIGRP use for IPv6?
12. Match the IP multicast address with its description:

    i. 224.0.0.1

    ii. 224.0.0.2

    iii. 224.0.0.5

    iv. 224.0.0.10

    a. All OSPF routers

    b. All routers

    c. EIGRP routers

    d. All hosts

13. Match the BGP attribute with its description:

    i. Local preference

    ii. MED

    iii. AS path

    iv. Next hop

    a. IP address

    b. Indicates the path used to exit the AS

    c. Tells external BGP peers the preferred path into the AS

    d. List of AS numbers

14. Which Cisco feature can you use instead of local preference to influence the selected path to external BGP routers?

15. What is the purpose of route reflectors?

16. When BGP confederations are used, which number do external peers see?

17. With ____________ all routers peer with each other within the private AS, and with __________ client routers peer only with the reflector.

18. Which of the following shows the correct order that BGP uses to select a best path?

    a. Origin, lowest IP, AS path, weight, local preference, MED

    b. Weight, local preference, AS path, origin, MED, lowest IP

    c. Lowest IP, AS path, origin, weight, MED, local preference

    d. Weight, origin, local preference, AS path, MED, lowest IP

19. What feature did BGPv4 implement to provide forwarding of packets based on IP prefixes?

20. What route should be used to summarize the following networks?

    10.150.80.0/23, 10.150.82.0/24, 10.150.83.0/24, 10.150.84.0/22

    a. 10.150.80.0/23, 10.150.82.0/23, and 10.150.84.0/22

    b. 10.150.80.0/22 and 10.150.84/22

    c. 10.150.80.0/21

    d. 10.150.80.0/20

21. Match the IPv6 multicast address with its description:

i. FF02::1

ii. FF02::2

iii. FF02::5

iv. FF02::9

v. FF02::A

a. OSPFv3 routers

b. RIPng routers

c. All routers

d. EIGRP routers

e. All nodes

22. Route summarization and redistribution occur in which layer of the hierarchical model?

- **a.** Building access
- **b.** Distribution
- **c.** Core
- **d.** Server access

23. Which of the following best describes route summarization?

- **a.** Grouping contiguous addresses to advertise a large Class A network
- **b.** Grouping noncontiguous addresses to advertise a larger network
- **c.** Grouping contiguous addresses to advertise a larger network
- **d.** Grouping Internet addresses

Refer to Figure 12-17 to answer the following questions.

Figure 12-17 *Network Scenario*

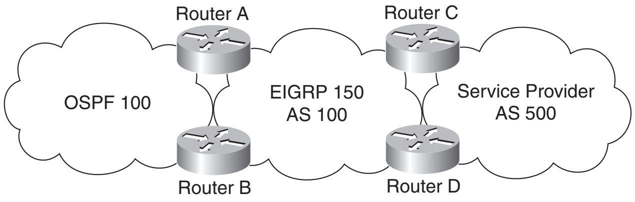

24. Where should you configure BGP?

- **a.** Routers A and B
- **b.** Routers C and D
- **c.** Answers A and B
- **d.** Routers A and C

25. On which router should you configure redistribution for OSPF and EIGRP?

- **a.** Router A only
- **b.** Router B only
- **c.** Routers A and B
- **d.** Redistribution occurs automatically.

26. To announce the networks from AS 100 to AS 500, which routing protocols should you redistribute into BGP?

- **a.** OSPF only
- **b.** EIGRP only
- **c.** OSPF and EIGRP
- **d.** iBGP

27. Where should you use filters?

- **a.** Routers A and B
- **b.** Routers C and D
- **c.** Routers A and C
- **d.** Answers A and B